

Kilian, Lutz

Working Paper

The Impact of the Fracking Boom on Arab Oil Producers

CESifo Working Paper, No. 5751

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Kilian, Lutz (2016) : The Impact of the Fracking Boom on Arab Oil Producers, CESifo Working Paper, No. 5751, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/128457>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Working Papers

www.cesifo.org/wp

The Impact of the Fracking Boom on Arab Oil Producers

Lutz Kilian

CESIFO WORKING PAPER NO. 5751

CATEGORY 10: ENERGY AND CLIMATE ECONOMICS

FEBRUARY 2016

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

ISSN 2364-1428

CESifo

Center for Economic Studies & Ifo Institute

The Impact of the Fracking Boom on Arab Oil Producers

Abstract

This article contributes to the debate about the impact of the U.S. fracking boom on U.S. oil imports, on Arab oil exports, and on the global price of crude oil. First, I investigate the extent to which this oil boom has caused Arab oil exports to the United States to decline since late 2008. Second, I examine to what extent increased U.S. exports of refined products made from domestically produced crude oil have caused Arab oil exports to the rest of the world to decline. Third, the article quantifies by how much increased U.S. tight oil production has lowered the global price of oil. Using a novel econometric methodology, it is shown that in mid-2014, for example, the Brent price of crude oil was lower by \$10 than it would have been in the absence of the fracking boom. I find no evidence that fracking was a major cause of the \$64 decline in the Brent price of oil from July 2014 to January 2015, however. Fourth, I provide evidence that the decline in Saudi foreign exchange reserves between mid-2014 and August 2015 would have been reduced by 27 percent in the absence of the fracking boom.

JEL-codes: Q430, Q330, F140.

Keywords: Arab oil producers, Saudia Arabia, shale oil, tight oil, oil price, oil imports, oil exports, refined product exports, oil revenue, foreign exchange reserves, oil supply shock.

Lutz Kilian
Department of Economics
University of Michigan
611 Tappan Street
USA – Ann Arbor, MI 48109-1220
lkilian@umich.edu

February 7, 2016

I thank Christiane Baumeister and Bassam Fattouh for helpful comments and discussion. This work has benefited from a financial grant from the Economic Research Forum. The contents and recommendations do not necessarily reflect the views of the Economic Research Forum.

1. Introduction

The use of hydraulic fracturing (or fracking) in conjunction with horizontal drilling and micro-seismic imaging has made it possible to extract crude oil from rock formations characterized by low permeability. Oil extracted by these techniques is commonly referred to as tight oil or shale oil to differentiate it from crude oil extracted by conventional drilling techniques. To date commercial shale oil production has been largely limited to the United States. The U.S. oil fracking boom is an example of a technological change in a single industry in one country affecting international trade worldwide. Increased U.S. shale oil production over time has displaced crude oil exports from Arab oil producing countries, both because the United States no longer relies as heavily on crude oil imports from Arab countries and because U.S. refineries have increasingly exported refined products such as gasoline or diesel made from domestically produced crude oil, causing other countries to cut back on their crude oil imports as well (see Kilian 2016). Whereas the gains to the U.S. economy of the fracking boom are well understood at this point, little is known to date about the losses this development has imposed on foreign oil producers. Understanding the implications of the U.S. fracking boom is important not only for policymakers in Arab economies deciding on how best to respond to the tight oil boom, but it also provides a prime example of a well-identified exogenous shock to the terms of trade of primary commodity exporters.

This article quantifies the impact of the U.S. tight oil boom on U.S. imports of crude oil and exports of refined products as well as Arab exports of crude oil.¹ On the basis of these estimates, I construct an estimate of how the price of oil in global markets would have evolved in the absence of this boom. This counterfactual allows me to quantify the losses in Saudi oil revenue (and hence in Saudi national income) since late 2008. In section 2 I review the evolution of crude oil production in the United States and elsewhere, highlighting the contribution of the U.S. fracking boom to global oil production growth. One of the channels by which the fracking boom has affected Arab oil producers is a reduction in U.S. crude oil imports since late 2008. Section 3 provides estimates of how many barrels of crude oil Arab oil producers have been unable to export to the United States due to the fracking boom. I show that

¹ An excellent nontechnical study of the impact of the tight oil boom on Arab oil producers is Fattouh (2014).

the quantitative importance of this effect differs substantially across Arab oil producers. A second channel is the reduction in Arab crude oil exports to the rest of the world caused by increased U.S. exports of refined products, as U.S. refiners take advantage of access to low-cost domestically produced crude oil. This effect is quantified in section 4.

Measuring the financial implications of the U.S. fracking boom for Arab oil producers is even more challenging than quantifying its physical effects on trade in crude oil. My analysis of this question in section 5 focuses on Saudi Arabia as the leading example. Using a new econometric methodology, I develop two data-based counterfactuals aimed at estimating the implications of these oil export losses for Saudi oil revenues between December 2008 and August 2015. My preferred counterfactual provides evidence that the cumulative losses in Saudi oil revenue caused by the fracking boom by August 2015 had reached 102 billion dollars. It also shows that the decline in Saudi foreign exchange reserves between mid-2014 and August 2015 would have been reduced by 27 percent (or 24 billion U.S. dollars) in the absence of the U.S. fracking boom, with the remaining 66 billion dollars reflecting increased oil production by other oil producers, shifts in oil price expectations, and the weakening of the global economy.

In constructing this counterfactual, I provide explicit estimates of the effect of the U.S. fracking boom on the Brent price of crude oil for every month since late 2008. Based on a recently proposed econometric model of the global market for crude oil I show that the cumulative effect of the fracking boom on the Brent price had been building gradually, reaching a peak in mid-2014, before declining in late 2014. Whereas in mid-2014 the Brent price was lower by \$10 than it would have been in the absence of the fracking boom, by mid-2015 this price differential had fallen to \$5. My analysis also sheds new light on the question of whether the decline in the price of oil after June 2014 was primarily caused by the fracking boom. I demonstrate that a very similar price decline would have occurred between July 2014 and January 2015 even in the absence of increased U.S. shale oil production. This conclusion complements the analysis in Baumeister and Kilian (2016a) based on an alternative methodology. Baumeister and Kilian were unable to pin down the quantitative role of fracking specifically, but provided

strong evidence that a slowdown in the global demand for oil was a major contributor to this specific oil price decline, along with shocks to global oil production and oil price expectations.

The concluding remarks in section 6 relate the estimates of Saudi oil revenue losses to recent reports of a substantial decline in Saudi foreign exchange reserves. I also discuss implications of my findings for Saudi economic policies in the years to come. The technical details of the construction of the counterfactuals are described in an appendix.

2. Recent Changes in Oil Production in the United States and Elsewhere

Only a few years ago, a common view among pundits was that world oil production would no longer be able to keep up with growing oil consumption needs. An extreme version of this skepticism was embodied in the peak oil hypothesis which asserted that global oil production had permanently peaked by 2007 or that the peak was imminent.² Although economists tend to be skeptical of the peak oil hypothesis for reasons discussed in Holland (2008, 2013), it was readily apparent at the time that many traditional oil fields were in decline and that the prospects for discovering crude oil in regions not ridden with civil strife were diminishing. Thus, even granting that high oil prices driven by strong demand for oil tend to provide strong incentives for expanding future oil production, as of 2007, nothing in past experience guaranteed that this supply response to rising oil prices would be adequate or that it would occur in a timely manner.

As the left panel of Figure 1 shows, with the benefit of hindsight, these concerns were unwarranted. Not only did the U.S. fracking boom reverse the long-standing decline in U.S. crude oil production, but Canadian production of unconventional crude oil from oil sands soared and Russian oil production reached unprecedented levels. At the same time, both Saudi Arabian and other Persian-Gulf production growth, which seemed to have levelled off after 2005, accelerated again. This surge in oil production was more than enough to offset the decline in crude oil production in other countries. Table 1 shows an increase in global oil production of 6.17 million barrels per day (mbd) between November 2008

² For example, an IMF study by Benes et al. (2015), using data up to 2009, predicted a near doubling of the price of oil over the coming decade based on the view that geological constraints would win out over technological improvements in conserving oil use and in oil extraction.

and September 2015, corresponding to an average rate of increase of 1.16 percent. Average growth accelerated to 2.63 percent between June 2014 and September 2015. Table 1 dispels the notion that fracking in the United States alone has been responsible for this surge. Although the United States was responsible for the bulk of the oil production increase, accounting for an extra 4.3 mbd over the last seven years, there were also notable production increases by Iraq (2.07 mbd), Saudi Arabia (1.23 mbd), Russia (0.79 mbd) and Canada (0.73 mbd) that were unrelated to the fracking boom.

The evolution of U.S. oil production in particular is remarkable, given the long-standing decline in U.S. oil production that began in the early 1970s and was only briefly reversed by the development of the Alaskan oil fields in the late 1970s. The vertical line in the left panel of Figure 1 refers to the date of November 2008, which marks the reversal of this trend. It can be shown that this reversal is largely due to the U.S. fracking boom. The right panel of Figure 1 plots the sum of crude oil production at the Bakken, Eagle Ford, Haynesville, Marcellus, Niobrara, Permian, and Utica shale oil plays. These seven regions account for 92% of U.S. domestic shale oil production growth. Between November 2008 and September 2015, U.S. shale oil production increased by approximately $3.82 / 0.92 = 4.15$ million barrels/day (mbd), which roughly matches the entry of 4.30 mbd for the change in total U.S. crude oil production in the first column of Table 1. Thus, throughout this article, I will treat November 2008 as the beginning of the fracking boom.

3. Quantifying the Decline in U.S. Crude Oil Imports from Arab Producers

The U.S. oil sector is governed by the following identity, which describes the composition of the total quantity of crude oil used by the economy:

$$Q_t + M_t - X_t = C_t - \Delta I_t,$$

where Q_t is domestic crude oil production, M_t is U.S. crude oil imports, X_t is U.S. crude oil exports, C_t is the consumption of crude oil, which refers to crude oil used by U.S. refineries to produce refined products such as gasoline or diesel for domestic consumption or for export, and ΔI_t denotes the change in crude oil inventories.

The left panel of Figure 2 illustrates how U.S. domestic oil production, crude oil imports and exports of crude oil have evolved since the 1970s. Three facts stand out. First, as already noted, following a secular decline in U.S. oil production, U.S. oil production sharply accelerated in late 2008, reflecting the fracking boom. Second, historically crude oil imports have been rising during economic expansions and falling during economic downturns. This pattern changed in 2005, even before the fracking boom, when the U.S. economy learned to make do with less crude oil during a boom amidst high and rising oil prices. The decline in crude oil imports accelerated during the fracking boom, when increased domestic shale oil production at low prices compared with the global price of oil reduced the need for high-cost crude oil imports. The reason for the comparatively low price of shale oil was a glut of light sweet crude oil in the central United States that reflected transportation bottlenecks within the United States as well as a ban on crude oil exports dating back to the 1970s (see Kilian 2016). There have been some exceptions to this ban, allowing for oil exports to Canada, for example, but overall this oil export ban has been binding. As a result, U.S. crude oil exports have remained negligible throughout most of this sample. Only in 2014 and 2015 U.S. oil exports increased somewhat, as the Obama administration allowed for more flexibility in the interpretation of the law governing U.S. oil exports. As of December 2015, Congress passed a law lifting the oil export ban, but this law has yet to be passed by both houses and signed into law by the president.

The right panel of Figure 2 quantifies the total use of crude oil by the U.S. economy, computed as $Q_t + M_t - X_t$. In quantifying the importance of oil imports for the U.S. economy, it is useful to express U.S. crude oil imports as a share of the amount of crude oil used by the U.S. economy. This approach controls for fluctuations in the size of the U.S. economy and for slow-moving changes in the energy efficiency of the U.S. economy. During a recession, for example, all else equal, both the use of crude oil by the economy and its oil imports would be expected to fall without affecting the oil import share. In contrast, if the fracking boom displaces crude oil imports from Arab oil producers, this should be reflected in a reduction in the share of U.S. imports from Arab countries in the oil used by the U.S.

economy.

The left panel of Figure 3 shows the evolution of the share of U.S. crude oil imports from the rest of the world as well as the share of U.S. crude oil imports from OPEC, from Arab OPEC (defined as Algeria, Iraq, Kuwait, and Saudi Arabia for our purposes), and from West Africa (Nigeria, Angola).³ The overall import share declined from a peak of almost 70 percent in the mid-2000s to about 45% in 2015. This pattern is also found for OPEC oil imports and its components. By 2015 crude oil imports from West Africa fell to near zero. West African crudes are in direct competition with domestically produced light sweet crude oil in the United States. As U.S. refiners along the East Coast found ways of substituting shale oil (or lower priced conventional light sweet crude oil from domestic sources) for these imports, West African crudes were almost entirely displaced. Crude oil imports from Arab OPEC as well fell from about 20 percent of U.S. oil use as recently as 2008 to under 10 percent in 2015. The reason presumably was that these crudes no longer could compete with the low price of domestically produced U.S. crude oil and U.S. imports of heavy Canadian crudes (see Kilian 2016).

The right panel of Figure 3 traces the import shares of four Arab oil producers before and after November 2008, which marks the beginning of the fracking boom (see Figure 1). We are interested in examining whether there are important changes in the dependency of the U.S. economy on oil imports from these countries. Figure 3 shows a striking drop in the share of U.S. oil imports from Saudi Arabia after November 2008. After some fluctuations, the Saudi share reaches near 5% in 2015, compared with about 10% initially. Figure 3 also shows a steady decline in the share of U.S. oil imports from Algeria (another producer of light sweet crude oil not unlike the West African oil producers). The share of U.S. oil imports from Iraq also declines over time. The higher volatility of the latter share is likely explained by political constraints on Iraqi oil production more than the fracking boom. Only the share of oil imports

³ I approximate crude oil imports from these countries by petroleum imports under the assumption that U.S. product imports from Arab countries are zero. This assumption can be verified for selected dates using data in the EIA's database for U.S. crude oil imports at the company level. Figure 3 excludes Libya from the Arab oil producers because Libyan oil exports in recent years are driven primarily by political events in Libya. The UAE and Qatar are excluded, because the EIA does not provide monthly data for U.S. imports from these countries, given their small size.

from Kuwait remains largely stable.

How quantitatively important is the reduction in U.S. export market share for Arab oil producers? A natural counterfactual is how Arab oil exports to the United States would have evolved, if the share of U.S. imports from those countries had remained constant at its average value between September 2004 and November 2008. This benchmark is reasonable in that the Saudi share, for example, remained remarkably stable over this period (see Figure 3). Disregarding earlier data in computing this share helps avoid variation in the share caused by exogenous events such as the 2003 Iraq War. In the absence of alternative explanations of the variation in Arab oil exports, the difference between this counterfactual path of Arab oil exports and the observed path of Arab oil exports to the United States can be interpreted as the loss in Arab oil exports caused by the fracking boom. Figure 4 plots the evolution of these export losses measured in millions of barrels per day between December 2008 and September 2015. In other words, Figure 4 shows estimates of the amount of crude oil that the United States would have imported from Saudi Arabia, Kuwait, Algeria, and Iraq, respectively, if domestically produced U.S. crude oil had not displaced these crude oil imports. Clearly, the assumption that all of the variation in the share of Arab imports can be attributed to the fracking boom may be questioned. There could be alternative explanations of the disproportionate drop in the U.S. import share from Saudi Arabia between 2009 and 2011, for example. For now we put these concerns aside. The realism of this working assumption will be assessed in section 5 by confronting counterfactual oil price estimates based on this assumption with extraneous evidence.

4. Quantifying the Displacement of Arab Oil Exports to the Rest of the World

Given the U.S. oil export ban, very little of the shale oil production generated by the fracking boom was exported. It may seem that therefore this production would not matter for the price of oil beyond the reduction in U.S. crude oil imports. This conclusion would be mistaken because the glut of light sweet crude oil in the central United States that emerged following the fracking boom depressed the price of domestically produced crude oil. U.S. refiners with access to low-priced domestic crude oil therefore

increased exports of refined products such as gasoline or diesel, notably to Europe and to Latin America. These exports of refined products effectively circumvented the ban on crude oil exports. To the extent that countries in the rest of the world that bought these refined products no longer needed to import as much crude oil for their refineries as before, this trade in products reduced the demand for Arab oil exports, adding to the costs of the fracking boom to Arab oil producers.

The left panel of Figure 5 shows a surge in U.S. exports of refined products in recent years. Assuming that the share of U.S. product exports in the oil used by the U.S. economy had stayed constant at its November 2008 value, one can construct a counterfactual for how exports of refined products would have evolved in the absence of the fracking boom starting in late 2008. The difference between the path of actual and counterfactual U.S. product exports provides a rough measure of the magnitude of the implied loss in Arab exports to the rest of the world. The right panel of Figure 5 expresses this loss in terms of its crude oil equivalent. The crude oil equivalent is computed by multiplying the shortfall of product exports computed in the left panel by 1.3548. This conversion factor is based on the fact that U.S. refineries in 2014 on average produced 12 gallons of diesel and 19 gallons of gasoline (for a total of 31 gallons of refined product) from 42 gallons (one barrel) of crude oil, according to the U.S. Energy Information Administration. Assuming an average mix of diesel and gasoline, this implies a conversion factor of $42/31=1.3548$.

This baseline estimate of the Arab crude oil exports lost to U.S. product exports, however, ignores the fact that rising Canadian exports of crude oil allowed the United States to export more refined products than otherwise would have been possible. Imports of heavy crudes from Canada in particular appear to have displaced U.S. imports of heavier crudes from Mexico and Venezuela with the latter countries' import shares after late 2005 steadily falling from 9 percent each to about 4 percent and Canada's share steadily rising from 15 percent to 21 percent over the same period.⁴ Thus, the increase in U.S. oil imports from Canada exceeds the combined reduction in U.S. oil imports from Mexico and in

⁴ These estimates were obtained by multiplying U.S. petroleum imports from these countries by the country-specific share of crude oil in these imports obtained from October 2014 data in the Energy Information Administration's company level reports.

U.S. oil imports from Venezuela. It would clearly be misleading to attribute the resulting product exports to the fracking boom.

Subtracting the net increase in crude oil imports from Canada, Mexico and Venezuela from the baseline export loss series in the right panel of Figure 5 tends to reduce the estimate of the oil export losses of the rest of the world substantially. We take the latter series as the benchmark when assessing the role of the fracking boom in driving up U.S. refined product exports. Figure 6 quantifies the Arab oil export losses caused by the U.S. exports of refined products. For simplicity I assume that the total loss of exports to the rest of the world (measured as the crude oil equivalent of increased U.S. product exports) can be apportioned based on the average share of each country's oil production in oil production in the world excluding the United States. Not surprisingly, by far the largest losses are incurred by Saudi Arabia, given its large market share.

Figure 7 combines the direct losses of section 3 and the indirect losses of section 4 for each Arab oil producer. For Saudi Arabia, Algeria and Kuwait, there is a tendency for the export losses expressed in million barrels of crude oil per day to increase over time. In the case of Saudi Arabia, losses reach one million barrels per day in 2014; Algeria reaches 0.5 million barrels per day in 2013 and Iraq in 2015. Only Kuwait has been largely unaffected with no pronounced tendency for increased export losses over time.

5. How Much Oil Revenue Did Arab Oil Producers Lose?

So far we have estimated the foregone exports of Arab oil producers associated with the surge in U.S. tight oil production. An equally important question is how large the financial losses have been. For concreteness, my discussion in this section focuses on Saudi Arabia. A simple, but arguably unrealistic counterfactual one could entertain would be how much oil revenue Saudi Arabia lost compared with a world in which the additional crude oil produced by U.S. shale oil producers would have been provided by Saudi Arabia instead, holding constant the production of all other oil producers. The advantage of this thought experiment is that we can think of the price of oil under this counterfactual as having been the

same as the observed oil price to a first approximation, allowing us to estimate the loss in Saudi revenue caused by the fracking boom as

$$P_t^{actual} (Q_t^{Saudi,counterfactual} - Q_t^{Saudi,actual}),$$

where $Q_t^{Saudi,counterfactual} - Q_t^{Saudi,actual}$ denotes the Saudi export losses, as shown in Figure 7, and P_t^{actual} denotes the actual dollar price of oil, as measured by the Brent price.⁵

There are several important caveats underlying this proposal. First, this counterfactual abstracts from any stimulating effects increased shale oil production may have on the U.S. economy. Those effects are small enough to be ignored, however (see Kilian 2016). Second, and more importantly, we abstract from differences in the quality of crude oil coming from Saudi Arabia compared with shale oil. These differences matter in that Saudi crude oil exports are not a perfect substitute for the light sweet crude oil that many refineries along the U.S. East Coast and in Europe are designed to process (see Kilian 2016). A third and even bigger concern with this counterfactual is that it is doubtful that Saudi Arabia would have been able to generate a surge in oil production of a magnitude comparable to the surge in U.S. shale oil production. As a result, this simple counterfactual is not credible.

Instead I examine an alternative and presumably more realistic set of counterfactuals, which asks how different the price of crude oil would have been, if all producers but the United States had maintained their observed oil production levels and if the U.S. fracking boom had never happened. In other words, I examine the cumulative effects of lowering world oil production by the amount of Arab oil export losses shown in Figure 7 (or alternatively by the direct estimate of shale oil production in the right panel of Figure 1) on the global price of crude oil, allowing one to compute the implied increase in Saudi revenues as

$$Q_t^{Saudi,actual} (P_t^{counterfactual} - P_t^{actual}),$$

where $Q_t^{Saudi,actual}$ denotes the observed level of Saudi oil production and $P_t^{counterfactual} - P_t^{actual}$ denotes the

⁵ A better measure might be the monthly Saudi f.o.b. price of crude oil, but the latter series is not available on a continuous basis from the U.S. Energy Information Administration.

difference between the actual dollar price of oil and the dollar price of oil obtained under this counterfactual.

The counterfactual price may be estimated with the help of a recently developed structural vector autoregressive (VAR) model of the global market for crude oil (see, e.g., Kilian and Murphy 2014). This model includes four variables: (1) the growth in global crude oil production, (2) a suitably updated measure of cyclical fluctuations in global real economic activity originally proposed by Kilian (2009), (3) the log of the real price of crude oil (obtained by deflating the U.S. refiners' acquisition cost for crude oil imports by the U.S. CPI), and (4) the change in above-ground global crude oil inventories measured by a suitable proxy. We follow the literature in estimating the model at monthly frequency using seasonal dummies and 24 autoregressive lags.

The structural shocks are identified based on a combination of sign restrictions and bounds on the short-run price elasticities of oil demand and oil supply. The key identifying assumptions are restrictions on the signs of the impact responses of the four observables to each structural shock. There are four structural shocks. First, conditional on past data, an unanticipated disruption in the flow supply of oil causes oil production to fall, the real price of oil to increase, and global real activity to fall on impact. Second, an unanticipated increase in the flow demand for oil (defined as an increase in oil demand for current consumption) causes global oil production, global real activity and the real price of oil to increase on impact. Third, a positive speculative demand shock, defined as an increase in inventory demand driven by expectations shifts not already captured by flow demand or flow supply shocks, in equilibrium causes an accumulation of oil inventories and raises the real price of oil. The accumulation of inventories requires oil production to increase and oil consumption to fall (associated with a fall in global real activity). Finally, the model also includes a residual demand shock designed to capture idiosyncratic oil demand shocks driven by a myriad of reasons that cannot be classified as one of the first three structural shocks.

In addition to these static sign restrictions, the model imposes the dynamic sign restriction that structural shocks that raise the price of oil on impact do not lower the real price of oil for the first 12

months following the shock. The rationale for this restriction is that an unexpected flow supply disruption would not be expected to lower the real price of oil within the same year nor would a positive flow demand or speculative demand shock. Finally, the model imposes the restrictions that the impact price elasticity of oil supply is close to zero and that the impact price elasticity of oil demand cannot exceed the long-run price elasticity of oil demand, which conventional estimates put at -0.8. For further discussion of this model, the reader is referred to Kilian and Murphy (2014), Kilian and Lee (2014) and Baumeister and Kilian (2014).

The model is re-estimated on data extending to August 2015, as described in Kilian and Lee (2014). The estimates of the responses of each model variable to each oil demand and supply shock are shown in Appendix A. Our main interest in this article is the sequence of flow supply shocks in this model. Given a pre-specified path of global oil production, obtained by adjusting actual global oil production by its U.S. shale oil component, we may infer the sequence of flow supply shocks required to produce this path of oil production after November 2008, holding constant the remaining structural shocks in the model. Similar techniques are discussed in Waggoner and Zha (1999) and Baumeister and Kilian (2014). Knowledge of this sequence of flow supply shocks allows us to determine how much higher the price of oil would have been under this counterfactual compared with the actual price of oil. The implied sequence of price differentials $P_t^{counterfactual} - P_t^{actual}$ may be used to compute, month by month from December 2008 until August 2015, the additional revenue Saudi Arabia would have received after November 2008 in the absence of the U.S. fracking boom. A more detailed discussion of this procedure can be found in Appendix B.

There are several potential caveats to keep in mind. First, this thought experiment is based on the premise that increased U.S. shale oil production was entirely due to exogenous oil supply shocks rather than being caused by oil demand shocks. This assumption is consistent with the view that in the absence of technological innovation, the shale oil boom would not have taken place, but it makes no allowance for effects working through oil price expectations. Second, we assume that Saudi oil production would have

remained the same under this counterfactual. One could make the case that Saudi Arabia might have chosen to produce more or less crude oil than it did in the absence of the fracking boom. Whereas the simple counterfactual discussed at the beginning of this section postulated that Saudi Arabia would have produced a lot more oil in the absence of the fracking boom, the current counterfactual effectively postulates that Saudi Arabia would have been unable to increase its oil production compared with its observed level.

Figure 8 shows two measures of the counterfactual path of world oil production corresponding to competing views of how large world oil production would have been in the absence of the U.S. fracking boom. The indirect measure in the left panel is derived from the effects of fracking on international trade in crude oil. This measure consists of the reduction in U.S. crude oil imports from the rest of the world and the crude oil equivalent of the increase in U.S. exports of refined products, adjusted for net increases in U.S. crude oil imports from Canada, Mexico, and Venezuela. This measure can be computed analogously to the results in Figure 7. The direct measure in the right panel is based on the change in shale oil production since November 2008, as reported in the right panel of Figure 1. The indirect measure is slightly larger than the indirect measure, but the overall pattern is similar. One potential difference is that increased shale oil production need not translate to increased production of refined products within the same month, given bottlenecks in the shipping of this oil to U.S. refineries. Instead it may be accumulated in inventories and potentially refined at a later point (see Kilian 2016). The indirect measure accounts for this problem and hence may be more reliable. It may involve other approximation errors, however, that do not arise when constructing the direct measure.

Figure 9 shows the sequence of flow supply shocks that would have to be imposed in the global oil market model proposed by Kilian and Murphy (2014) to make global oil production follow the counterfactual path of world oil production in Figure 8. An important concern is whether the flow supply shocks required to implement the pre-specified path of global oil production after November 2008 are too large by historical standards or too predictable to maintain the assumption of a time-invariant structural VAR model. Figure 9 illustrates that neither of these shock sequences involves shocks that are unusually

large by historical standards, addressing the first concern. As to the second concern that the counterfactual shocks must not be predictable, it is useful to focus on evidence of a large number of consecutive flow supply shocks having the same sign (referred to as a “run”). In the upper panel the longest run observed after December 2008 is seven months. Runs of this length are not unprecedented. The longest run prior to December 2008, in fact, lasts nine months. Thus runs in the flow supply shock data are neither particularly unusual nor do they necessarily imply that market participants would have found it easy to anticipate these shocks. Rather their existence is consistent with the view that the rapidly increasing growth rate of shale oil production continued to surprise market participants more often than not. The corresponding shock sequence in the lower panel of Figure 9, in contrast, exhibits one run lasting for fourteen months, which is considerably longer than the longest run in the historical data, suggesting potential caution in interpreting results based on the latter counterfactual, but is not long compared to runs found in other series of prediction errors in the literature.⁶

Figure 10 plots the Brent price of crude oil in U.S. dollars (P_t^{actual}) along with the counterfactual evolution of the Brent price ($P_t^{counterfactual}$) based on these shock sequences. These estimates are obtained based on counterfactual simulations of the log of the real price of oil in the global oil market model discussed earlier, estimated on data until August 2015. The simulated oil price data have been suitably adjusted for differences between the Brent price and the U.S. refiners’ acquisition cost for crude oil imports. The Brent price of crude oil may be interpreted as a proxy for the global price of crude oil. The left panel of Figure 10 shows that under the indirect measure of the counterfactual evolution of world oil production the Brent price of oil would have been higher by as much as \$5 in 2009, but most of the cumulative effects of the fracking boom would have been observed between 2011 and mid-2014, with the counterfactual price exceeding the actual Brent price by as much as \$9 at times. Thereafter, the price differential $P_t^{counterfactual} - P_t^{actual}$ becomes negligible again. In contrast, under the direct measure of the

⁶ In a related context, Kilian and Hicks (2013) documented prediction errors in professional global real GDP growth forecasts of the same sign for nearly five years during 2003-08. As long as such patterns average out in the long run, they do not violate the premise of rational expectations.

counterfactual evolution of world oil production the Brent price of oil would have been largely the same as the actual price until late 2010, as shown in the right panel of Figure 10. Between 2011 and mid-2014, the Brent price of oil would have been higher by as much as \$10 in the absence of the U.S. fracking boom, reaching \$120 repeatedly and at some point even \$130. Unlike in the first counterfactual, even in 2015, the Brent price would have been higher than actually observed price by as much as \$5.

The estimates in Figure 10 also shed light on the extent to which the U.S. fracking boom caused the decline in the global price of crude oil in late 2014 and early 2015. Both counterfactuals agree that in the absence of the fracking boom this oil price decline would have occurred at the same time, as it did, but that it would have started at a higher level. According to Figure 10, the cumulative decline in the Brent price would have been \$70 under both specifications of the counterfactual (and therefore would have been \$6 higher than in the actual data). Thus, there is no support for the view that U.S. fracking was the primary cause of this specific oil price decline. This result complements the analysis in Baumeister and Kilian (2016a) who were unable to pin down the quantitative role of fracking specifically, but who provided evidence that a slowdown in the global demand for oil was a major contributor to this specific oil price decline in addition to a mix of shocks to actual and/or expected global oil supplies prior to July 2014 and a shift in oil price expectations in July 2014. Figure 10 helps refine that interpretation.

The two upper panels of Figure 11 show the price differentials $P_t^{\text{counterfactual}} - P_t^{\text{actual}}$ obtained from Figure 10 under the two alternative counterfactuals. Figure 11 highlights two important differences. Only the indirect measure implies a positive price differential during 2009 and early 2010, whereas only the direct measure implies a positive price differential after late 2014. Which of these price differentials is economically more plausible may be assessed based on the information in the bottom panel of Figure 11, which shows the dollar spread between the Brent price and the West Texas Intermediate (WTI) price of crude oil. Whereas the Brent price refers to the spot price of light sweet crude oil traded in London, UK, the WTI price is the spot price of light sweet crude oil for delivery in Cushing, OK.

As discussed in Kilian (2016), traditionally these crude oil prices have been closely tracking one

another. This pattern only changed when the U.S./ tight oil boom created a glut of crude oil in the central United States, causing WTI oil to trade at a discount relative to the global price for oil, as measured by the Brent price. Thus, the spread in the bottom panel of Figure 11 conveys important information about when the fracking boom began to affect the U.S. market for crude oil. The fact that the Brent-WTI spread only widened in late 2010 suggests that the positive price differential during 2009 and early 2010 in the upper panel cannot be attributed to the fracking boom, casting doubt on this counterfactual. Likewise, the fact that the Brent-WTI price spread remains positive in late 2014 and early 2015 is difficult to reconcile with the price differential in the upper panel being zero or even negative. This extraneous evidence supports the conclusion that the price differential in the second panel based on the direct measure of the counterfactual is a more accurate approximation. The remainder of the analysis therefore focuses exclusively on the latter counterfactual.⁷

As outlined earlier, the price differential $P_t^{counterfactual} - P_t^{actual}$ shown in Figure 11 may be used to infer the losses in Saudi oil revenue, given the historically observed levels of Saudi oil production. The upper panel of Figure 12 shows the evolution of the losses in Saudi oil revenue measured in billions of U.S. dollars under the preferred counterfactual based on the direct measure of tight oil production. There is evidence that the Saudi oil revenue losses were negligible initially, but gradually increased over time, reaching 2 billion dollar per month in late 2012 and peaking at 2.9 billion dollars per month in mid-2014, before levelling off. Of course, this evidence does not speak to the overall loss of Saudi oil revenue following the decline in the price of oil after June 2014, but only to the component of this loss caused by the U.S. fracking boom. It should also be noted that the dip in Saudi oil revenue losses in January 2015 mirrors the temporary tightening of the Brent-WTI spread, adding credibility to this estimate.

⁷ It is important to stress that the Brent-WTI price differential does not by itself represent a proper counterfactual for the global price of oil because an increase in tight oil production all else equal will have a larger effect on the U.S. price of oil than on the global price. All we can say is that if there is no effect on the U.S. price of oil until late 2010, then fracking is unlikely to have affected the global price before this point in time. The evolution of the Brent-WTI spread between 2011 and 2014, in contrast, tells us little about the counterfactual for the global price because it reflects in important part the evolution of transportation and refining bottlenecks in the United States (see Kilian 2016).

A rough measure of the cumulative impact of the fracking boom on Saudi Arabia is obtained by simply cumulating these losses. The cumulative decline in Saudi oil revenue between December 2008 and August 2015 is 102.1 billion dollars. This amounts to an average loss per month of about 1.3 billion dollars. The incremental Saudi oil revenue losses between July 2014, when the price of oil began to fall precipitously, and August 2015 are 24.3 billion dollars. To put these estimates in perspective, the lower panel of Figure 12 expresses the reduction in Saudi oil revenues caused by the fracking boom as a percentage of the Saudi oil revenue in November 2008, when the Brent price of oil stood at \$52.45. It is important to be clear about the interpretation of this estimate. Figure 12 shows that in mid-2014, for example, Saudi Arabia in the absence of the U.S. fracking boom would have earned 120% of the oil revenue it generated in November 2008, if it had produced exactly as much oil as it actually did in mid-2014. In other words, the plot does not measure the actual change in the oil revenue generated by the Saudi economy, but rather the lost revenue relative to what Saudi Arabia would have earned if the United States still relied on conventional crude oil production only.

It should be kept in mind that this estimate is only an approximation, first, because Saudi oil exports that were diverted from the United States to Asian economies, for example, were not necessarily priced based on the Brent benchmark, and, second, because my methodology abstracts not only from changes in Saudi oil production, but also from any Saudi efforts to market its oil more aggressively. My analysis also abstracts from price differences between Saudi exports of heavy and light crude oil. The concern is not so much that Saudi Arabia could have substituted exports of heavy crude to the United States for its exports of light crude. This possibility seems unlikely because, as I showed, imports of lower-priced Canadian heavy crudes increasingly substituted for U.S. imports of heavy crudes from other sources, at the same time as shale oil production displaced U.S. imports of light sweet crude oil. Rather the concern is that by applying the counterfactual price differential indiscriminately to Saudi production of light and heavy crude oil, one overestimates the loss in oil revenue that can be attributed to shale oil specifically, given that shale oil competes with Saudi heavy crude oil only indirectly through U.S. exports of refined products.

6. Concluding Remarks

Saudi Arabia's strategy for dealing with the fracking boom so far has been to preserve its market share and to avoid idle capacity (see, e.g., Fattouh and Sen 2015; Fattouh, Poudineh and Sen 2015). There is no sign of Saudi Arabia accommodating the fracking boom by cutting back its oil production. This strategy comes at the cost of falling oil revenues. There are indications that the Saudi government has dealt with the recent decline in its oil revenue mainly by tapping its financial reserves, as predicted by standard models of precautionary savings (see, e.g., Bems and de Carvalho Filho 2011). Press reports put the reduction in Saudi foreign exchange reserves between mid-2014 and August 2015 near \$90 billion. This estimate is considerably larger than the loss in oil revenue over the same period that is attributable to the fracking boom. The latter we estimated to be about \$24 billion. In other words, the analysis in this article tentatively suggests that in the absence of the fracking boom the reduction in Saudi foreign exchange reserves since mid-2014 would have been less than one third of the actual decline with the remainder reflecting higher crude oil production elsewhere in the world, shifts in oil price expectations affecting demand for oil stocks, and the effects of a slowing global economy (see Baumeister and Kilian 2016a).

A common question is how long Saudi Arabia can sustain such losses in its foreign exchange reserves. If we take the estimates of the decline in the Saudi foreign exchange reserves at face value, at August 2015 prices the Saudi central bank can be expected to lose about \$15 billion in foreign exchange reserves every month. At this rate, one would expect the Saudi foreign exchange reserves to be exhausted by early 2019. Of course, this process would accelerate at the current even lower price of crude oil (or if the price of oil were to fall even further). There are several mitigating factors, however. First, unlike some of its neighbors, Saudi Arabia has very little external debt because it resisted the temptation to leverage its increased oil revenues during the 2003-08 oil price boom by running up external debt (see Tornell and Lane 1998). As a result, for now, the Saudi government has been able to borrow in global financial markets, given the expectation that the price of oil will ultimately recover to levels closer to \$85. Second, the Saudi government so far has been reluctant to impose fiscal retrenchment, because of the political costs of such measures, but in December 2015 it announced a budget plan that envisions cuts in spending

from \$975 billion in 2015 to \$840 billion in 2016 (representing a 14% drop in spending). The finance ministry also announced that it would raise taxes and adjust subsidies for water, electricity and petroleum products over the next five years.

How much fiscal adjustment Saudi Arabia requires will depend ultimately on how long one expects low oil prices to persist. The answer depends in important part on why the price of oil declined and how important these determinants will remain in the future (see Baumeister and Kilian 2016b). For example, to the extent that this price decline was caused by the fracking boom, the relevant question becomes how long this boom can persist at current prices with many commercial tight oil producers already experiencing heavy operating losses and how quickly these firms could resume tight oil production at higher oil prices, if they were forced to close down. The same concern applies to unconventional oil production in Canada. To the extent that low oil prices reflect a sluggish global economy, in contrast, the question becomes at what point emerging Asia, Japan, and Europe, in particular, will recover. A swift and sustained global economic recovery, unlikely as it may seem at this point, would quickly eliminate the current glut of crude oil, all the more so as oil production in many other countries will continue to fall over time. Finally, to the extent that other oil producers including Russia and Iraq have expanded oil production, or, as in the case of Iran, are about to do so, the question becomes one of how long state-owned oil producers with few other options to generate foreign exchange and no accountability to shareholders or financial markets will put pressure on the price of oil.

Although one's view about the persistence of low oil prices may differ, depending on the weight one puts on these factors, it is not unlikely that the Saudi economy could face another two or three lean years, at which point the precautionary savings in the Saudi sovereign wealth fund would likely be exhausted, while the prospects of new external borrowing would likely diminish. Thus, there appears to be no alternative to some measure of fiscal retrenchment in the foreseeable future. The same argument applies even more forcefully to other Arab oil producers faced with more foreign debt and lower foreign exchange reserves.

References

- Baumeister, C., and L. Kilian (2014), “Real-Time Analysis of Oil Price Risks using Forecast Scenarios,” *IMF Economic Review*, 62, 119-145.
- Baumeister, C., and L. Kilian (2016a), “Understanding the Decline in the Price of Oil since June 2014,” forthcoming: *Journal of the Association of Environmental and Resource Economists*.
- Baumeister, C., and L. Kilian (2016b), “Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us,” forthcoming: *Journal of Economic Perspectives*.
- Bems, R., and I. de Carvalho Filho (2011), “The Current Account and Precautionary Savings for Exporters of Exhaustible Resources,” *Journal of International Economics*, 84, 48-64.
- Benes, J., Chauvet, M., Kamenik, O., Kumhof, M., Laxton, D., Murula, S., and J. Selody (2015), “The Future of Oil: Geology versus Technology,” *International Journal of Forecasting*, 31, 207-221.
- Coglianese, J., Davis, L.W., Kilian, L., and J.H. Stock (2016), “Anticipation, Tax Avoidance, and the Price Elasticity of Gasoline Demand,” forthcoming: *Journal of Applied Econometrics*.
- Fattouh, B. (2014), “The U.S. Tight Oil Revolution and its Impact on the Gulf Cooperation Council Countries: Beyond the Supply Shock,” OIES Paper WPM 54, Oxford Institute of Energy Studies.
- Fattouh, B. and A. Sen (2015), “Saudi Arabia Oil Policy: More than Meets the Eye?”, OIES Paper MEP 13, Oxford Institute of Energy Studies.
- Fattouh, B., Poudineh, R., and A. Sen (2015), “The Dynamics of the Revenue Maximization-Market Share Trade-Off: Saudi Arabia’s Oil Policy in the 2014-2015 Price Fall”, OIES Paper WPM 61, Oxford Institute of Energy Studies.
- Holland, S.P. (2008), “Modeling Peak Oil,” *Energy Journal*, 29, 61-80.
- Holland, S.P. (2013), “The Economics of Peak Oil,” in: Shogren, J.F. (ed.), *Encyclopedia of Energy, Natural Resource, and Environmental Economics*, Vol. 1, Amsterdam: Elsevier, 146-150.

- Kilian, L. (2009), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market”, *American Economic Review*, 99, 1053-1069.
- Kilian, L. (2016), “The Impact of the Shale Oil Revolution on U.S. Oil and Gas Prices,” forthcoming: *Review of Environmental Economics and Policy*.
- Kilian, L., and B. Hicks (2013), “Did Unexpectedly Strong Economic Growth Cause the Oil Price Shock of 2003-2008?” *Journal of Forecasting*, 32, 385-394.
- Kilian, L., and D.P. Murphy (2014), “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil,” *Journal of Applied Econometrics*, 29, 454-478.
- Kilian, L., and T.K. Lee (2014), “Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories,” *Journal of International Money and Finance*, 42, 71-87.
- Tornell, A., and P. Lane (1998), “Are Windfalls a Curse? A Non-Representative Agent Model of the Current Account,” *Journal of International Economics*, 44, 83-112.
- U.S. Energy Information Administration (2015), *Monthly Energy Review*, December.
- U.S. Energy Information Administration (2016), *Drilling Productivity Report*, January.
- Waggoner, D.F., and T. Zha, (1999), “Conditional Forecasts in Dynamic Multivariate Models,” *Review of Economics and Statistics*, 81, 639-651.

Table 1. Changes in Crude Oil Production

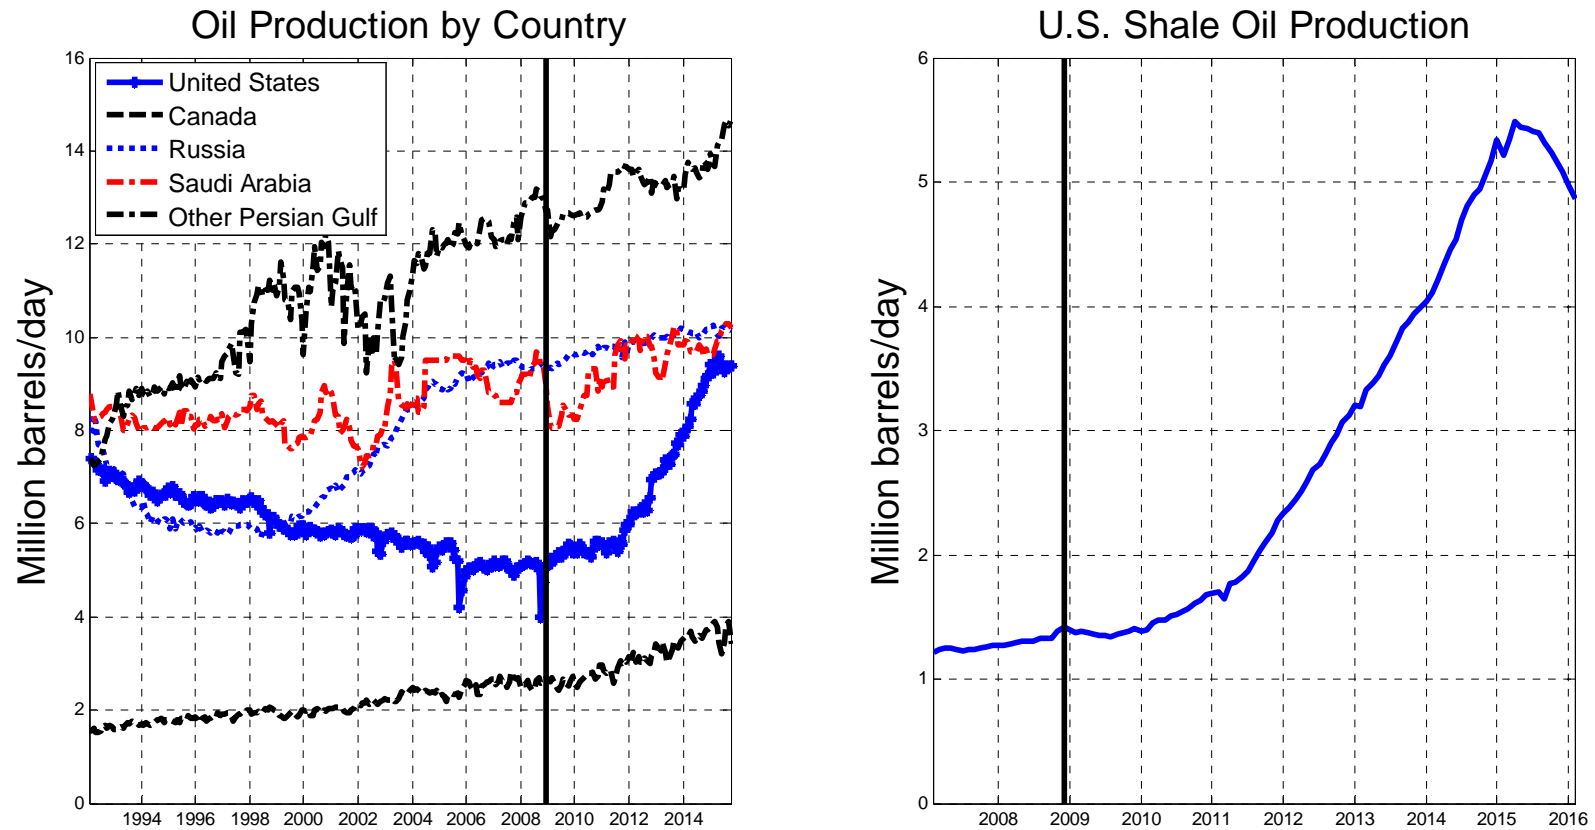
Oil producers	2008.11-2015.9		2014.6-2015.9	
	Change in million barrels/day	Average percent change ¹	Change in million barrels/day	Average percent change ¹
World	6.17	1.16	2.76	2.63
Saudi Arabia	1.23	1.86	0.50	3.78
Other Persian Gulf	1.97	2.08	1.14	6.05
Iran	-0.80	-3.13	0.15	3.49
Iraq	2.07	9.12	1.10	21.63
Kuwait	0.06	0.37	-0.10	-2.88
Qatar	0.36	3.90	-0.02	-0.78
UAE	0.25	1.39	0	0
Russia	0.79	1.17	0.06	0.41
Canada and United States	5.03	7.24	0.57	3.44
Canada	0.73	3.48	-0.14	-2.93
United States	4.30	8.89	0.71	5.91

Source: U.S. Energy Information Administration (2105), *Monthly Energy Review*.

Notes: 2008.11 marks the reversal in the long-term decline of U.S. oil production. 2014.6 marks the end of a long period of relative price stability in the market of crude oil, before the recent collapse of the price of crude oil (see Baumeister and Kilian 2016).

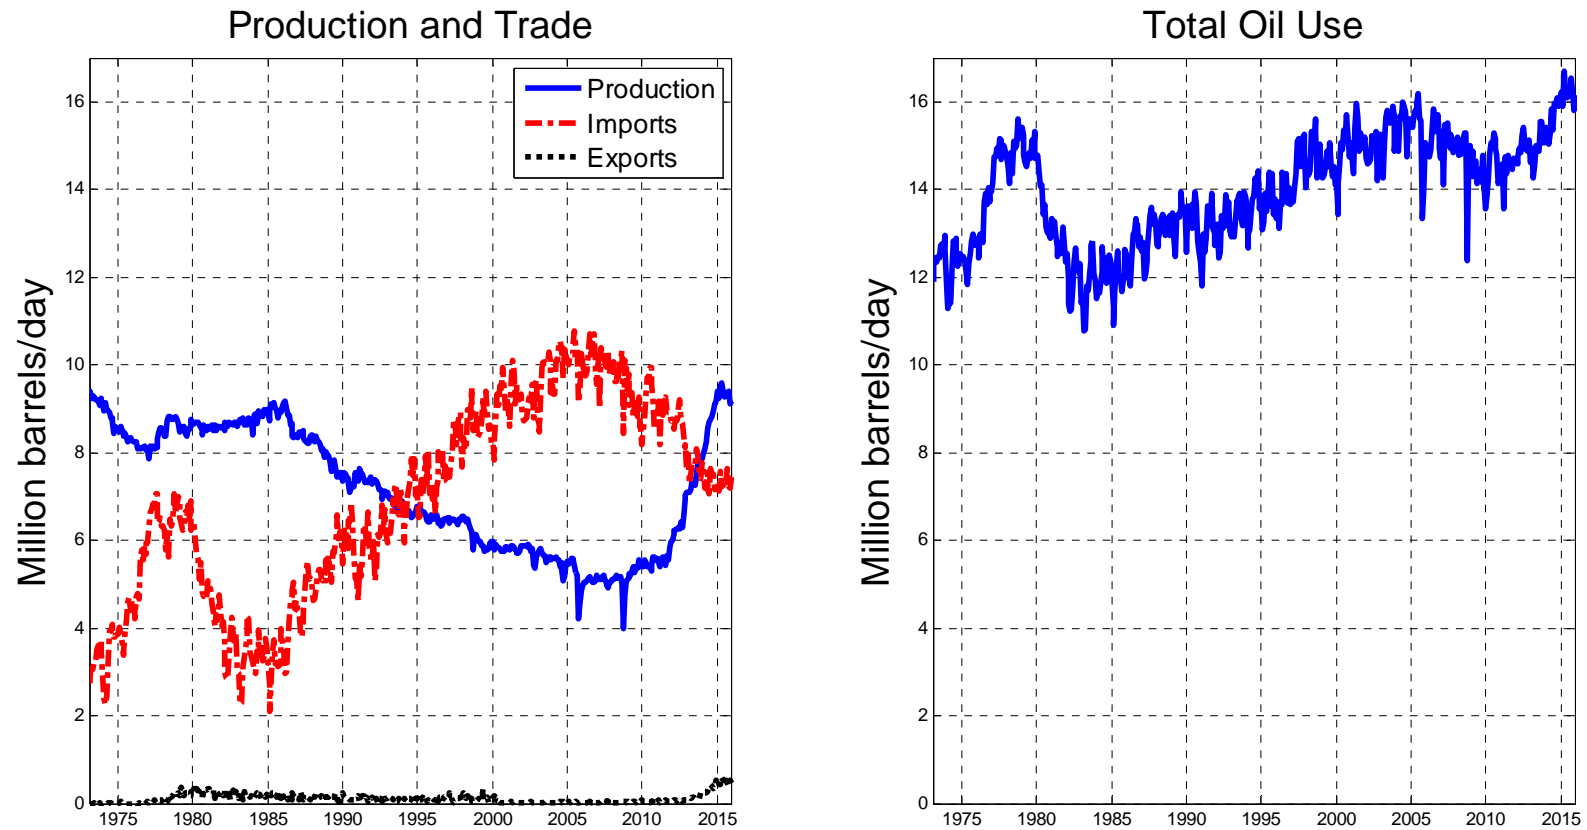
¹ at annualized rates

Figure 1. Crude Oil Production Disaggregated



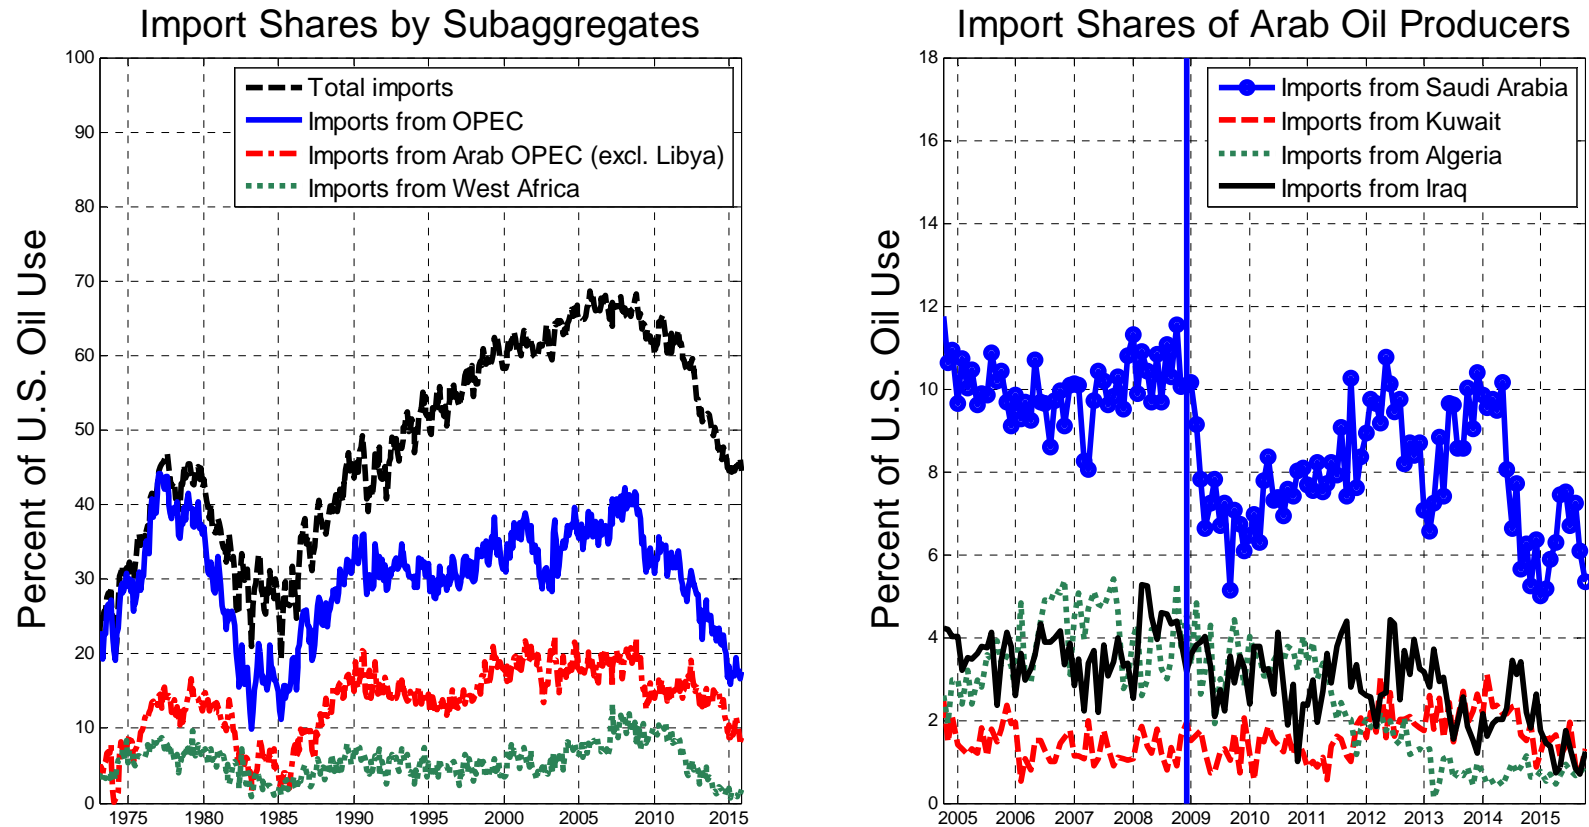
NOTES: The international data are from U.S. Energy Information Administration (2015). The production data are for crude oil and lease condensate, but exclude natural gas plant liquids. The vertical line marks the reversal in the long-term decline of U.S. oil production in November 2008. The shale oil production data are obtained from U.S. Energy Information Administration (2016) by adding crude oil production in the seven largest U.S. shale oil plays, which account for the bulk of U.S. shale oil production.

Figure 2. The U.S. Crude Oil Sector, 1973.1-2015.11




NOTES: The raw data are from U.S. Energy Information Administration's (2015). Total oil use is defined as: Domestic Production of Crude Oil + Crude oil imports – Crude oil exports.

Figure 3. Crude Oil Imports from Selected Countries as a Percent Share of U.S. Oil Use




NOTES: The raw data are from U.S. Energy Information Administration (2015). The vertical line marks the reversal in the long-term decline of U.S. oil production in November 2008. It also coincides with the drop of the share of Saudi oil imports in the U.S. use of oil.

Figure 4. Export Losses of Arab Oil Producers Based on Counterfactual of a Constant U.S. Import Share after 2008.11




NOTES: Based on data from U.S. Energy Information Administration (2015). The counterfactual involves assuming that the share of U.S. crude oil imports from each country remains at its average value between November 2004 and November 2008 for the remainder of the sample. This reference period is chosen to reduce the impact of past political events in some of these countries on their share.

Figure 5. Quantifying the Displacement of Other Crude Oil Producers' Exports Caused by U.S. Exports of Refined Products




NOTES: Based on data from U.S. Energy Information Administration (2015). The vertical line marks the reversal in the long-term decline of U.S. oil production in November 2008. The counterfactual is constructed by assuming that the share of U.S. exports of refined products in total use of oil by the U.S. economy remains constant at its level in November 2008. The crude oil equivalent is computed by multiplying the shortfall of exports computed in the left panel by 1.3548. This conversion factor is based on the fact that U.S. refineries in 2014 on average produced 12 gallons of diesel and 19 gallons of gasoline (for a total of 31 gallons of refined product) from 42 gallons (one barrel) of crude oil, according to the EIA. The Canada effect refers to increased imports of crude oil from Canada in excess of the reduction of U.S. imports of crude oil from Mexico and from Venezuela. To isolate U.S. exports of refined products made from domestic crude oil, this effect must be corrected for.

Figure 6. Export Losses of Arab Oil Producers Based on Counterfactual of No U.S. Exports of Refined Products Made from Shale Oil after 2008.11




NOTES: The estimates are obtained by apportioning the crude oil equivalent in Figure 5 to each country according to the share of its crude oil production in non-U.S. world crude oil production between 2004.9 and 2008.11. This reference period is chosen to reduce the impact of past political events in some of these countries on their share.

Figure 7. Direct and Indirect Crude Oil Export Losses Combined




NOTES: The estimates are obtained by summing the direct loss arising from the fact that the United States reduced its imports from Arab oil producers given the increased availability of shale oil, as quantified in Figure 4, and the indirect loss arising from the loss Arab crude of oil export markets, as U.S. refineries increase their exports of refined products, as quantified in Figure 6, including the adjustment for net gains in crude oil imports from Canada, Mexico and Venezuela.

Figure 8. Alternative Counterfactuals for World Oil Production in the Absence of the U.S. Fracking Boom




NOTES: The indirect counterfactual is based on the losses in crude oil exports of the rest of the world caused by a reduction in U.S. crude oil imports and the expansion of U.S. exports of refined products, adjusted for the net increase in U.S. crude oil imports from Canada, Mexico, and Venezuela, representing primarily an influx of heavy crudes unrelated to fracking. The direct counterfactual is derived on an estimate of the change in U.S. shale oil production since 2008.11, as shown in the right panel of Figure 1.

Figure 9. Sequence of Flow Supply Shocks Required to Implement the Alternative Counterfactuals in Figure 8




NOTES: Computations by the author based on the algorithm described in Appendix B. The vertical line marks the beginning of the counterfactual in December 2008.

Figure 10. The Impact of Shale Oil Production on the Brent Price of Crude Oil




NOTES: The counterfactual path for the nominal Brent price of crude oil inferred from the two counterfactuals in Figure 8, as discussed in Appendix B.

Figure 11. Brent Oil Price Differentials




NOTES: The first two panels are based on computations of the author reflecting the counterfactual price paths for the Brent price in Figure 10. The third panel shows the spread between the Brent and WTI price.

Figure 12. Saudi Oil Revenue Losses Under the Direct Measure of the Counterfactual



NOTES: The plot in the lower panel expresses the reduction in Saudi oil revenues caused by the fracking boom as a fraction of Saudi oil revenue in November 2008, when the Brent price of oil stood at \$52.45. For example, at its peak in 2014 the oil revenue loss amounted to 20% of the Saudi oil revenue of November 2008.

Appendix A: Responses of the Global Oil Market VAR Model Variables to Each of the Structural Shocks



NOTES: The underlying model is constructed as in Kilian and Lee (2014), except the estimation sample has been updated to August 2015. All responses but the inventory responses are expressed in percentages. The inventory responses are in million barrels.

Appendix B: Constructing the Oil Price Series under Counterfactual 2

Let $y_t = (y_{1t}, y_{2t}, y_{3t}, y_{4t})'$ denote the vector of variables in the global oil market model of Kilian and Murphy (2014), where y_{1t} stands for the growth in world oil production and y_{3t} for the log of the real price of oil, in particular. Consider the structural VAR model

$$B_0 y_t = B_1 y_{t-1} + \dots + B_{24} y_{t-24} + w_t, \quad (1)$$

where the 4×1 vector y_t is assumed to be stationary and the deterministic regressors have been suppressed for expository purposes. The dimension of $B_i, i = 0, \dots, 24$, is 4×4 . The 4×1 vector of structural shocks, $w_t = (w_{1t}, w_{2t}, w_{3t}, w_{4t})'$, where w_{1t} denotes a shock to the flow supply of crude oil, is assumed to be white noise with a diagonal 4×4 variance-covariance matrix Σ_w that is of full rank and without loss of generality can be normalized to equal the identity matrix. This structural model can be expressed in its reduced form as

$$y_t = A_1 y_{t-1} + \dots + A_{24} y_{t-24} + u_t, \quad (2)$$

where $A_i = B_0^{-1} B_i, i = 1, \dots, 24$, and $u_t = B_0^{-1} w_t$ with variance-covariance matrix $\Sigma_u = B_0^{-1} (B_0^{-1})'$. Let the 4×4 matrix $\partial y_{t+i} / \partial w_t' = \Theta_i$ denote the responses of the model variables to each of the structural shocks at horizon $i = 0, 1, \dots, H$. The matrix $\Theta_i = [\theta_{jk,i}]$ consists of elements $\theta_{jk,i} = \partial y_{j,t+i} / \partial w_{kt}$ that denote the response of variable j to structural shock k at horizon i . These responses may be computed as $\Theta_i = J A^i J' B_0^{-1}$, where $J = \begin{bmatrix} I_4 & 0_{4 \times 4(24-1)} \end{bmatrix}$ and A denotes the matrix of slope parameters obtained by expressing the VAR(24) reduced-form model in its VAR(1) companion format.

In constructing the counterfactual in question, we make use of the fact that after removing the deterministic terms

$$y_t \approx \sum_{i=0}^{t-1} \Theta_i w_{t-i}.$$

As a result, the fitted value for the log real price of oil can be written as

$$\hat{y}_{3t} = \sum_{j=1}^4 \hat{y}_{3t}^{(j)}, \quad (3)$$

where $\hat{y}_{3t}^{(j)}$ denotes the cumulative contribution since the beginning of the sample of structural shock j to the third model variable at time t , defined as

$$\begin{aligned} \hat{y}_{3t}^{(1)} &= \sum_{i=0}^{t-1} \theta_{31,i} w_{1,t-i}, \\ \hat{y}_{3t}^{(2)} &= \sum_{i=0}^{t-1} \theta_{32,i} w_{2,t-i}, \\ \hat{y}_{3t}^{(3)} &= \sum_{i=0}^{t-1} \theta_{33,i} w_{3,t-i}, \\ \hat{y}_{3t}^{(4)} &= \sum_{i=0}^{t-1} \theta_{34,i} w_{4,t-i}. \end{aligned}$$

Having simulated the posterior of the reduced-form model (2) obtained based on the full sample from February 1973 to August 2015, we recover estimates of B_0^{-1} , $\theta_{jk,i}$ and of $w_t = B_0 u_t$, $t = 1, \dots, T$, as in Kilian and Lee (2014) by selecting the posterior model with a short-run price elasticity of oil demand closest to the benchmark estimate of -0.26 in the literature (see Kilian and Murphy 2014; Coglianese et al. 2016).

In constructing the counterfactual, we retain the realizations of the flow supply shock $w_{1,t}$ up to November 2008, but we replace the estimates of $w_{1,t}$ for December 2008.12 through August 2015 by counterfactual values chosen to ensure that the path of world oil production corresponds to the counterfactual path shown in Figure 9. Given the suitably demeaned growth rate of world oil production implied by this counterfactual, these values may be easily computed by an iterative procedure. Having computed what the expected value of y_{1t} given $\hat{y}_{1t}^{(j)}$, $j = 2, 3, 4$, would have been next month, if this month's supply shock had been zero, the difference between that model prediction and the target level of oil production growth, scaled by the impact response of world oil production to a flow supply shock, will be the magnitude of the flow supply shock, w_{1t} , required to reach the growth target.

We then recompute $\hat{y}_{3t}^{(1)}$ under the alternative sequence of flow supply shocks, $w_{1,t}$, while retaining the other three structural shock sequences as originally estimated in computing $\hat{y}_{3t}^{(2)}$, $\hat{y}_{3t}^{(3)}$ and $\hat{y}_{3t}^{(4)}$. Given the implied sequence of \hat{y}_{3t} , the corresponding sequence $P_t^{\text{counterfactual}}$ is constructed from the

real price of Brent crude oil, scaled up by the percent deviation between the counterfactual real price of oil and the actual real price of oil in the model. The counterfactual real Brent price is converted back to dollars using the U.S. consumer price index. The price differential $P_t^{counterfactual} - P_t^{actual}$ for the remainder of the sample is obtained by subtracting the actual dollar price of Brent crude oil, P_t^{actual} , from $P_t^{counterfactual}$.