

Güntner, Jochen

Working Paper

How do oil producers respond to oil demand shocks?

Working Paper, No. 1311

Provided in Cooperation with:

Johannes Kepler University of Linz, Department of Economics

Suggested Citation: Güntner, Jochen (2013) : How do oil producers respond to oil demand shocks?, Working Paper, No. 1311, Johannes Kepler University of Linz, Department of Economics, Linz

This Version is available at:

<https://hdl.handle.net/10419/97447>

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.

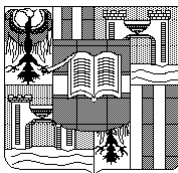
Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

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.



How do oil producers respond to oil demand shocks?

by

Jochen H. F. GÜNTNER^{*)}

Working Paper No. 1311
July 2013

**Johannes Kepler University of Linz
Department of Economics
Altenberger Strasse 69
A-4040 Linz - Auhof, Austria
www.econ.jku.at**

jochen.guentner@jku.at
phone +43 (0)70 2468 -8360, -9679 (fax)

How do oil producers respond to oil demand shocks?*

Jochen H. F. Güntner[†]

Johannes Kepler University Linz

July 18, 2013

Abstract

This paper analyzes the response of international oil producers to demand-induced changes in the real price of oil during 1975–2011. The goal is to disentangle fluctuations in OPEC and non-OPEC production and to derive consistent estimates of the short-run price elasticity of crude oil supply at the country level. I find that oil producers hardly respond to demand shocks within the same month, and that the corresponding impact price elasticities of supply are not statistically different from zero. Although there is little evidence of a systematic response following a typical flow demand shock, the medium-run responses to a speculative demand shock differ between OPEC and non-OPEC producers, i.e., on average over the sample period, OPEC members seem to curtail production, whereas non-OPEC supply expands significantly. Flow and speculative demand shocks account for a non-negligible fraction of the total variability in country-level crude oil production.

Keywords: Oil demand shocks; OPEC; Crude oil production; Price elasticity of crude oil supply

JEL Classification: C32; N50; Q41

*I am grateful to Christiane Baumeister, Lutz Kilian, and Dan Murphy for helpful comments and for sharing their MATLAB code. All remaining mistakes are my own.

[†]Jochen Güntner is Assistant Professor at the Department of Economics, Johannes Kepler University Linz. Address: Altenberger Strasse 69, 4040 Linz-Auhof, Austria. E-mail: jochen.guentner@jku.at, telephone: +43 732 2468-8360, fax: +43 732 2468-9679.

1 Introduction

The present paper analyzes the impulse responses of country-level crude oil production to two types of oil demand shocks, on impact and over time. Monthly estimates of flow and speculative demand shocks for 1975.2–2011.12 are drawn from Kilian and Murphy's (2013) structural VAR model of the global market for crude oil. These orthogonal innovations are used as regressors in a finite order distributed lag model of country-level crude oil production. The focus is on the differences and commonalities in the behavior of OPEC and non-OPEC producers.

Economically, OPEC and non-OPEC oil producers differ in numerous dimensions, e.g. in the public vs. private ownership and extraction of crude oil reserves and the relative importance of oil exports as a source of (government) revenues and foreign exchange. Even though the cartel hypothesis has been overturned, OPEC member states might not only respond to economic incentives but also take political considerations into account, e.g., they might respond differently to an increase in oil demand in order to exploit higher oil prices.¹ Given that they possessed 81% of proven world crude oil reserves below the ground in 2011, while accounting for 42.8% of world production, OPEC members might also react differently, just because they are less capacity-constrained (see, e.g., Hamilton, 2009a, for the historical variability of Saudi Arabian oil production).

A vast empirical literature boils down to one unresolved question: How do oil producers respond to a change in oil prices? For more than twenty years, the standard econometric procedure has been to regress changes in log production levels on changes in the log (real) price of crude oil and other explanatory variables, interpreting the estimate of the price coefficient as the elasticity of oil supply (see, e.g. Griffin, 1985; Jones, 1990; Ramcharran, 2001, 2002). Implicit in this approach is the assumption that causality in the oil market runs from price to production; notwithstanding that in any market price and quantity are determined simultaneously, in general. As shown by Barsky and Kilian (2002, 2004) and Hamilton (2008), the reverse causality from production to price invalidates any estimates of the oil price coefficient and their interpretation as price elasticities of crude oil supply in simple log-log-regressions.

¹This hypothesis is consistent with the finding in Alhajji and Huettner (2000) that *target revenue theory* (TRT) helps explain production decisions in centrally planned as opposed to free-market economies. According to TRT, oil exporters' targeted revenues are effectively determined by their investment needs.

The awareness of this pitfall does not resolve a second, equally important concern. Like any commodity's price, the oil price is determined simultaneously by demand and supply conditions in the market for crude oil (see Barsky and Kilian, 2002, 2004; Kilian, 2008a,b). An observed increase in the price of oil might therefore arise from lower supply, higher demand, or both. Accordingly, spot market oil prices reflect the current composition of the underlying oil demand and supply shocks, at any point in time.

The structural decomposition of oil price fluctuations into oil demand and supply shocks in Kilian (2009) has revealed that "not all oil price shocks are alike". Similarly, different structural innovations may lead to different endogenous responses of oil producers, even though they unanimously imply an increase in the real price of oil. Ex ante, there is no obvious reason why different circumstances should imply the same behavior of oil producers. The only way to address this second concern is by identifying changes in the price of oil which are exogenous from the perspective of crude oil suppliers, such as shifts of the demand curve along the supply curve.

For this purpose, I draw on recent advances by Kilian and Murphy (2013) to disentangle oil supply and demand shocks in a structural VAR model of the global market for crude oil. Structural innovations, which influence the real price of oil, are identified by a combination of sign restrictions on impulse response functions and plausible upper bounds on the impact price elasticities of oil supply and demand. *Flow supply shocks* capture disruptions in physical output measured by world crude oil production. *Flow demand shocks* correspond to unexpected changes in the demand for oil and other industrial commodities due to fluctuations in global real economic activity. Given the fact that oil is a storable commodity, flow supply and demand shocks will not be sufficient to describe the fluctuations in the real price of oil, in general.

Accordingly, Kilian and Murphy (2013) distinguish between shifts in the flow demand and shifts in the demand for oil stocks, which are orthogonal to the global business cycle. *Speculative demand shocks* capture changes in the demand for crude oil inventories above the ground arising from a revision of expectations about future supply relative to future demand or from a precautionary motive, i.e. a shift in the conditional variance as opposed to a shift in the conditional mean of future demand or supply (see also Alquist and Kilian,

2010).² According to this broad definition, any purchases of crude oil in the spot market which are not for current consumption are thus categorized as speculative demand.

A consistent estimate of the behavior of crude oil producers in response to demand-induced fluctuations in the real price of oil can then be obtained by regressing country-level oil production on contemporaneous and lagged values of the orthogonal flow and speculative demand shocks, in a second step.³ Note that the latter are orthogonal to *aggregate* flow supply shocks by construction.

I find that country-level production hardly responds to either demand shock, on impact. Over time, an increase in oil demand associated with the global business cycle leads to a hump-shaped response of total OPEC production and a slight expansion of total non-OPEC production, albeit neither is significant at conventional confidence levels. Both patterns are dominated by a few important producers such as Saudi Arabia in the case of OPEC and Russia in the case of non-OPEC. A typical speculative demand shock is followed by a statistically significant reduction of total OPEC and a significant expansion of total non-OPEC production. The fact that Saudi Arabia curtails supply in response to increased speculative demand suggests that (some) OPEC members decelerate the accumulation of inventories. In contrast, adverse news and higher uncertainty about the future supply relative to the future demand provide a strong motive for relatively advanced economies to increase domestic production levels or to rely on politically more stable suppliers such as, e.g., Mexico and Russia.

On average over the sample period, oil demand shocks explain between 7 and 20% of the overall variance of country-level crude oil production, with speculative demand accounting for relatively larger shares, in general. I also find that total OPEC production is relatively more sensitive to demand-induced oil price fluctuations, for either shock considered separately as well as for both shocks together.

Since both flow demand and speculative demand shocks have a positive effect on the real price of oil, on impact, my finding of a negligible impact response for most producers implies that the short-run supply curve is almost vertical even at the country level. This

²Kilian and Murphy (2013) and Kilian and Lee (2013) are the first to include proxies for crude oil inventories above the ground into a structural model of the global oil market in order to identify *speculative demand shocks* explicitly rather than as part of a residual category (compare, e.g., Kilian, 2009; Kilian and Murphy, 2012; Baumeister and Peersman, 2013a).

³Kilian (2009) applies a distributed lag regression model when assessing the effects of different types of oil shocks on U.S. GDP and inflation.

is fully consistent with the widespread consensus that the price elasticity of *world* crude oil supply is close to zero in the short run (see, e.g., Hamilton, 2009a; Kilian, 2009).⁴

The rest of the paper is organized as follows. Section 2 describes the data used in the SVAR and the distributed lag model. Section 3 outlines the econometric approach in Kilian and Murphy (2013) and derives historical time series of the orthogonal oil demand shocks. Section 4 introduces the distributed lag model and estimates impulse response functions of country-level oil production to flow demand and speculative demand shocks. Based on the impact responses of the real price of oil and country-level oil production, I propose new estimates of the short-run price elasticity of country-level crude oil supply. Section 5 concludes.

2 Data

World crude oil production levels in thousands of barrels per day (tbpd) are from the *U.S. Energy Information Administration* (EIA). I use Kilian’s (2009) index of global real economic activity, based on single voyage bulk dry cargo ocean shipping freight rates.⁵

The measure of the real price of oil is based on the nominal U.S. refiner acquisition cost of imported crude oil reported by the EIA since January 1974. The nominal time series is deflated by the seasonally adjusted U.S. consumer price index (CPI) published by the *U.S. Bureau of Labor Statistics* (BLS). Following Kilian (2009), the real price of crude oil is backcasted from 1974.1 to 1973.1 using the method in Barsky and Kilian (2002). The resulting time series is expressed in log levels in deviations from the mean.

Monthly data on crude oil inventories going back to January 1973 are available for the U.S. only. I therefore follow Hamilton (2009b) and Kilian and Murphy (2013) and scale this data by the ratio of OECD petroleum stocks over U.S. petroleum stocks, provided by the EIA.⁶ All data on crude oil inventories and petroleum stocks are in million barrels.

⁴Disaggregating even further, Kellogg (2011) finds essentially no response of oil production to either the spot or the futures price of oil in monthly well-level oil production data from Texas.

⁵The advantage of this global economic activity index over, e.g., an index of industrial production in the OECD – the closest proxy for monthly world industrial production – is that it neither requires exchange rate weighting nor does it exclude emerging economies such as China and India. For further information on the construction of the index and the underlying raw data, the reader is referred to Kilian (2009). The linearly detrended time series is publicly available from the author’s homepage.

⁶Note that “petroleum stocks” reported by the EIA comprise crude oil (including strategic reserves), unfinished oils, natural gas, plant liquids, and refined products. Since time series on petroleum stocks are provided for OECD countries only, total OECD stocks must serve as a proxy for world petroleum stocks. Moreover, the EIA reports monthly levels of petroleum stocks for all OECD countries since

The VAR model is estimated using monthly data for 1973.2 through 2011.12.

Country-specific data on crude oil production in thousand barrels per day (tbpd) for all countries and groups of countries is obtained from the EIA's *Monthly Energy Review*. Crude oil production levels are available for the following OPEC members: Algeria, Angola, Ecuador, Iran, Iraq, Kuwait, Libya, Nigeria, Qatar, Saudi Arabia, United Arab Emirates, and Venezuela. Non-OPEC production data are available for Canada, China, Egypt, Mexico, Norway, Russia/U.S.S.R., the United Kingdom, and the United States. Moreover, I analyze the aggregates for *Persian Gulf Nations*, *Total OPEC*, and *Total Non-OPEC*. All data are monthly and available since January 1973.

3 Identifying Oil Demand Shocks

In order to analyze the response of crude oil producers to demand-induced changes in the real price of oil, we require a measure of exogenous oil demand shocks. I use the structural VAR model proposed by Kilian and Murphy (2013) to disentangle oil demand and supply shocks in the global market for crude oil.

The vector of endogenous variables contains the percent change in world crude oil production, an index of global real economic activity, the real price of oil, and the change in crude oil inventories above the ground.⁷ The model includes monthly dummies to account for seasonal variation in crude oil inventories and 24 lags in order to capture slow-moving components in the global business cycle and the real price of oil.

3.1 Identification

Following Kilian and Murphy (2012, 2013), orthogonal oil supply and demand shocks are identified based on a combination of sign restrictions on impulse response functions and upper bounds on the impact price elasticities of oil supply and demand.

Flow supply shocks relate to unexpected shifts of the oil supply curve along the oil

December 1987 only. Following Kilian and Murphy (2013), the percent change in OECD inventories is extrapolated backwards using the growth rate of U.S. inventories prior to December 1987. For the overlapping period, the correlation between the growth rates of U.S. and OECD petroleum inventories is close to 80%. The resulting time series is expressed in absolute rather than percent changes, as the latter seems to be non-stationary (compare Kilian and Murphy, 2013).

⁷The fact that oil reserves below the ground are not accessible in the short run, e.g. for consumption smoothing, implies that above-the-ground inventories are predominantly relevant for a spot market equilibrium. For a detailed discussion of the implications of below-ground crude oil reserves, see section 4.3 in Kilian and Murphy (2013).

demand curve, conditional on past realizations of the endogenous variables. They include any unexpected disruptions of the physical supply of crude oil due to, e.g. exogenous political events in oil producing countries or strategic production decisions by OPEC member states. In line with economic theory, a negative supply shock reduces the level of output and raises the real price of oil, thus lowering real economic activity on impact, whereas the response of inventories remains unrestricted. Note that crude oil production and global real economic activity are assumed to decrease, while the real price of oil is assumed to increase, respectively, for at least twelve months.⁸

Flow demand shocks capture unexpected changes in the general demand for industrial commodities that are associated with fluctuations in the global business cycle, i.e. shifts of the oil demand curve along the supply curve, conditional on past realizations of all endogenous variables. Economic theory predicts that a positive flow demand shock raises global real economic activity and the real price of oil within the same month, whereas the impact response of inventories is ambiguous.

A *speculative demand shock* corresponds to an unexpected shift in the demand for crude oil inventories above the ground. This includes precautionary demand shocks, i.e. pure uncertainty shocks that do not affect the expected levels of future demand or supply, as well as changes in the expected future demand for crude oil *relative to* the expected future supply of crude oil. Expectations of political turmoil in oil producing countries such as Libya are a typical example. A positive speculative demand shock raises crude oil inventories and the real price of oil, on impact. The increase in inventories can only be accommodated by a simultaneous expansion of supply and a decrease in demand for consumption, i.e. in global real economic activity.

Table 1 summarizes all impact and dynamic sign restrictions, which are imposed as weak inequality constraints. The fourth shock absorbs any fluctuations in the real price of oil not otherwise accounted for by the model. Hence, it represents a residual category with ambiguous economic interpretation and is not discussed below.

In addition to these sign restrictions, Kilian and Murphy (2012) impose an upper bound on the impact price elasticities of crude oil supply. With regard to the long lead times

⁸By means of these *dynamic sign restrictions*, Kilian and Murphy (2013) exclude candidate models, in which the real price of oil decreases below and real economic activity increases above its initial level, respectively, after the impact period.

Table 1: Impact and Dynamic Sign Restrictions on Impulse Response Functions for the Structural Model in (1)

	Flow supply shock	Flow demand shock	Speculative demand shock
World oil production	$-(^{12})$	+	+
Real economic activity	$-(^{12})$	+	-
Real price of oil	$+(^{12})$	+	+
Crude oil inventories			+

Notes: All shocks are normalized to imply an increase in the real price of oil. All sign restrictions are imposed for the impact period as weak inequality constraints, i.e. $\leq / \geq 0$. $(^{12})$ indicates that the sign restriction is imposed for twelve months including the impact period.

in oil field development, significant adjustment costs in production, and the uncertainty whether an observed change in oil demand represents a transitory shock or a permanent shift, it is widely accepted that the short-run oil supply curve is close to, if not effectively vertical (see also Hamilton, 2009a).⁹ Furthermore, Kilian and Murphy (2013) impose an impact price elasticity of crude oil demand that does not exceed its long-run equivalent. Hausman and Newey (1995) and Yatchew and No (2001) provide empirical estimates of the long-run price elasticity of *gasoline demand* based on U.S. and Canadian household survey data, respectively, arriving at values of -0.8 and -0.9 .¹⁰

3.2 Historical Time Series of Demand Shocks

Based on 5 million random draws for the matrix of contemporaneous coefficients, only 16 models satisfy all identifying assumptions for the sample period in Kilian and Murphy (2013), i.e. 1973.2–2009.8. When extending the sample period to 2011.12, this number increases to 61, although the set of admissible structural models satisfying *only* the impact sign restrictions expands by a much smaller fraction. As we add more recent observations to the sample, a larger number of candidate models satisfies the bounds imposed on the short-run price elasticities of oil supply and demand.

Figure 1 plots the time series of oil supply and demand shocks for the admissible model with a price elasticity of oil demand in use closest to -0.26 , the posterior median of this elasticity in Kilian and Murphy (2013). For the sake of readability, the monthly shocks

⁹By focusing on the change in oil production outside of Kuwait and Iraq in response to the Invasion of Kuwait in August 1990, Kilian and Murphy (2012) derive an upper bound of 0.0258 for the short-run price elasticity of crude oil supply.

¹⁰Appendix B shows how the impact price elasticities of crude oil supply and demand can be expressed in terms of the VAR model's coefficients. For technical details on the empirical implementation of the identifying assumptions, the reader is referred to Kilian and Murphy (2013).

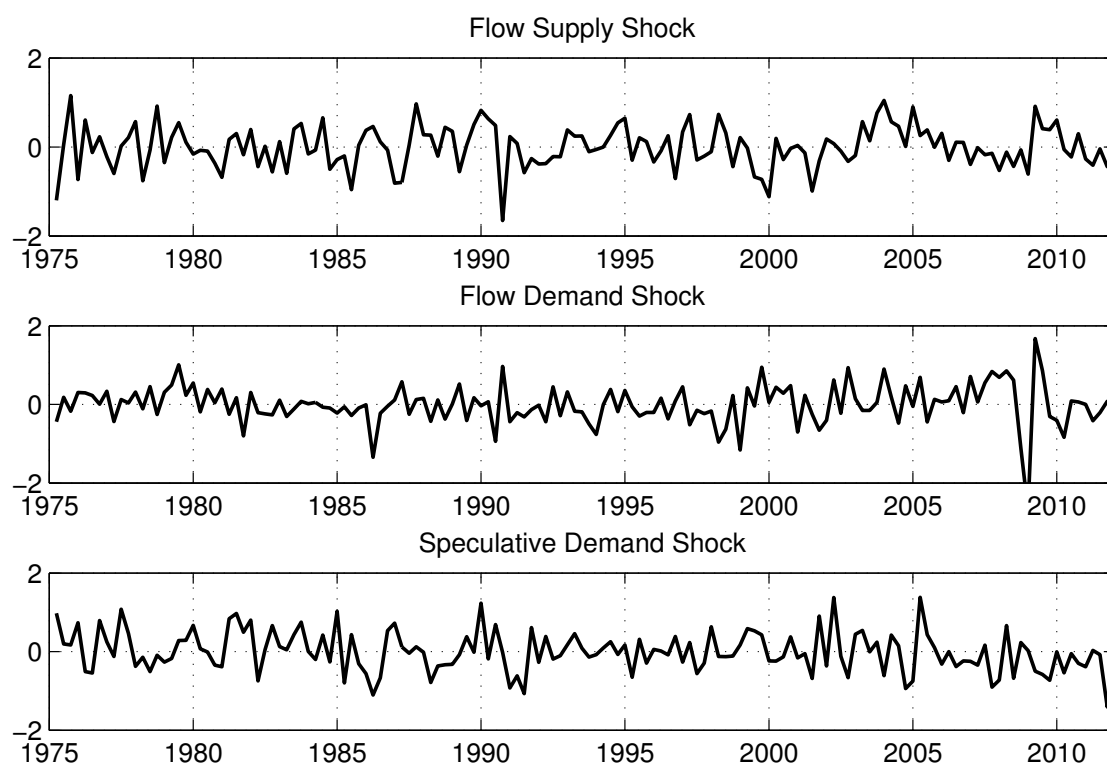


Figure 1: Historical evolution of orthogonal oil supply and demand shocks for 1975Q2–2011Q4

Note: Quarterly averages of monthly structural residuals based on candidate model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013)

are averaged to quarterly frequency. The figure illustrates that the global market for crude oil is subject to different structural shocks at any point in time, which will be reflected in the corresponding market price and volume.

Consider first the episode of the late 1970s. Despite the disruption of crude oil production during the Iranian revolution, there is no sign of a negative flow supply shock until the outbreak of the Iran-Iraq war in late 1980. Moreover, it seems that increased production in the rest of the world more than compensated for this shortfall within one year.

Instead, demand for oil in 1979 was unexpectedly high due to high global real economic activity as well as positive speculative demand shocks. The latter can be attributed to the coincidence of the Iranian Revolution and the Soviet invasion of Afghanistan, which raised concerns about the reliability of future oil supply from these countries (see Kilian, 2009). Demand for inventories remained consistently high throughout 1981, when it became clear that the Iran-Iraq war would not come to an end soon.

The model plausibly identifies a large negative supply shock in the third quarter of 1990,

when Iraqi forces invaded Kuwait, leading to a near halt of production in both countries. Following two spikes in late 1989Q4 and 1990Q2, speculative demand was unexpectedly low during the second half of 1990, reflecting a reduction of inventories and an immediate recovery of world production. Moreover, there are three episodes of sustained negative flow demand shocks after 1995. The first coincides with the Asian financial crisis of 1997–1998 and the Russian financial crisis of 1998, while the second reflects the recessions in developed economies, such as the U.S. and part of the European Union, in the early 2000s. The pronounced dip in 2008, finally, captures the unprecedented drop in global real economic activity and trade in the wake of the *Great Recession* of 2008–2009.

Note that the results shown in Figure 1 represent just one of the 61 admissible structural models, which satisfy all identifying restrictions for the sample period 1973.2–2011.12. Figure A.1 plots the time series of structural shocks for each admissible structural model, conditional on the least squares estimate of the reduced-form, and illustrates that the identified shock series are surprisingly robust to using different candidate models. Most importantly, the historical evolution of flow demand and speculative demand shocks is very similar across all admissible models.

For the sake of brevity, the discussion of structural impulse responses to oil supply and demand shocks in the VAR model of the global oil market is deferred to Appendix B. For the derivation of short-run price elasticities of country-level crude oil supply, it is sufficient to note that, on average over the sample period, a positive flow demand shock raises the real price of oil by about 5 percent, on impact, while a positive speculative demand shock raises the real price of oil by about 2.5 percent.

4 How Do Oil Producers Respond to Demand Shocks?

Based on the monthly time series of structural disturbances identified above, this section investigates the behavior of *individual* oil producing countries in response to flow demand and speculative demand shocks. For this purpose, I use a finite distributed lag model similar to the one in Kilian (2009). The main difference is that, here, all data are available at monthly frequency, sparing any aggregation of the structural residuals. Given that the latter are uncorrelated by construction, each regressor can be considered in a separate

distributed lag model. For oil producing country or group of countries i ,

$$\Delta prod_{i,t} = \alpha_i + \sum_{l=0}^{24} \phi_{i,l} \varepsilon_{t-l}^{flow\ demand} + u_{i,t} \quad (1)$$

and

$$\Delta prod_{i,t} = \beta_i + \sum_{l=0}^{24} \psi_{i,l} \varepsilon_{t-l}^{speculative\ demand} + v_{i,t}, \quad (2)$$

where $\Delta prod_{i,t}$ denotes percent changes in crude oil production. $u_{i,t}$ and $v_{i,t}$ are possibly serially correlated country-specific error terms. From the regression model in (1) and (2), it is straightforward to compute the *level* impulse response of oil production in country i at horizon h to a typical flow and speculative demand shock as $\sum_{j=0}^h \phi_{i,j}$ and $\sum_{j=0}^h \psi_{i,j}$, respectively. As a consequence, the maximum horizon of the impulse response functions, $H = 24$, is determined by the number of lags.

Note that the use of distributed lag regressions has an important advantage over dynamic simultaneous equations models such as a VAR model, where country-level production is substituted for world crude oil production. In (1) and (2), $\Delta prod_{i,t}$ is consistently regressed on the same time series of flow demand and speculative demand shocks, respectively. In contrast, the structural innovations obtained from a series of country-specific VAR models will be different for any two countries, in general, which might invalidate cross-country comparisons.

4.1 Identification

The specification in (1) and (2) rests on the assumption that flow demand and speculative demand shocks are predetermined with respect to changes in country-specific oil production, i.e., neither demand shock responds to $\Delta prod_{i,t}$ *within the same month*. Although predeterminedness is inherently untestable, there are strong arguments in favor of this assumption.

Consider first the case of $\varepsilon_t^{flow\ demand}$. To the extent that a positive flow demand shock is associated with a sustained increase in the real price of oil, it provides an incentive for oil producers to raise their output levels. If instead causality runs from country-specific supply to the demand for industrial commodities, the only plausible conclusion is that an unexpected increase in supply stimulates global real economic activity.¹¹ In any case,

¹¹Due to the fact that providers of shipping services hold large buffer stocks of bunker fuels, one might even claim that a minor oil price change arising from the shift in an individual country's production

Table 2: Contemporaneous Correlations of Oil Demand Shocks with Unexpected Innovations in Country-Level Crude Oil Production

	Flow demand shock	Speculative demand shock
Algeria	-.0126	.0019
Angola	-.0464	.0352
Ecuador	.0045	.0084
Iran	.0271	.0043
Iraq	.0530	.0551
Kuwait	-.0576	-.0253
Libya	.0595	.0253
Nigeria	.0032	-.0076
Qatar	.0173	-.0289
Saudi Arabia	.0165	-.0707
United Arab Emirates	-.0458	-.0173
Venezuela	-.0490	-.1199**
Canada	.0617	.0753
China	-.1025**	-.0235
Egypt	.1004**	.0223
Mexico	.0629	.0338
Norway	-.0206	.0312
Russia/U.S.S.R.	-.0114	.0014
United Kingdom	.0975**	.0170
United States	.0453	-.0165
Persian Gulf Nations	.0270	-.0291
Total OPEC	.0206	-.0322
Total Non-OPEC	.0482	.1123**

Notes: Orthogonal flow and speculative demand shocks based on candidate VAR model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). Unexpected innovations to oil production in country i based on a univariate autoregression in $\Delta prod_{i,t}$ with lag order $q = 24$ (reported results are robust to $q = 12$). *, **, *** indicate statistical significance at a level of 10, 5, and 1%.

the flow demand shock should be positively correlated with unexpected innovations to country-specific crude oil supply. The fact that the contemporaneous correlations, reported in the second column of Table 2, are usually very close to zero and statistically insignificant implies that neither direction of causality is relevant at monthly frequency.¹² Consider next the case of $\varepsilon_t^{speculative\ demand}$. Since a positive speculative demand shock raises the price of crude oil on impact, it entails an incentive for crude oil producers to raise output levels. Accordingly, one would expect the two to be positively correlated. At the same time, a negative innovation to country-level supply has an ambiguous effect on demand for crude oil inventories. The fact that a negative correlation might offset the level has zero impact on real economic activity within the same month (compare Kilian, 2009).

¹²The only countries for which the contemporaneous correlation is positive and statistically significant are Egypt and the U.K., whereas the correlation is negative and statistically significant for China.

positive relationship inferred above invalidates conclusions about the predeterminedness of speculative demand shocks with respect to country-level supply based on a statistically insignificant contemporaneous correlation alone.

However, to the extent that a disruption of crude oil production in any given country is *not* associated with a negative shock to world supply, it cannot affect crude oil inventories unless also contemporaneously affecting oil consumption. Since the flow supply and speculative demand shock series are orthogonal by construction and there is no feedback from country-specific production to flow demand shocks, as I have shown before, $\varepsilon_t^{\text{speculative demand}}$ can be treated as predetermined with respect to $\Delta prod_{i,t}$.¹³ For the sake of completeness, the last column of Table 2 reports the contemporaneous correlations between speculative demand shocks and unexpected innovations to production in country or group i . Note that only Venezuela displays a statistically significant negative correlation.

4.2 Responses to a Flow Oil Demand Shock

Figure 2 plots the impulse response functions of crude oil production in twelve OPEC and eight non-OPEC countries following a typical positive flow demand shock based on the distributed lag model in (1). For each country i , the point estimate of the impulse response is shown for the structural model with a price elasticity of oil demand in use closest to -0.26 , the posterior median of this elasticity in Kilian and Murphy (2013). One- and two-standard-error confidence intervals are based on 10,000 replications of a block bootstrap with block size 12 that accounts for the possibility of serial correlation in the regression residuals.

Note that these confidence intervals do not reflect the fact that the flow demand shocks in (1) are generated regressors. Due to the fact that the structural VAR model is identified by a combination of sign restrictions and upper bounds on the impact price elasticities of oil demand and oil supply, this would be computationally excessive, given a sufficient number of random draws for the inverse matrix of contemporaneous coefficients and a sufficient number of bootstrap replications for the distributed lag model.

¹³Following Kilian (2009), one could argue that speculative demand shocks reflect market participants' expectations or uncertainty about future supply and demand conditions in the *medium or long run*, which are hardly affected by the monthly fluctuations in crude oil production at the country level.

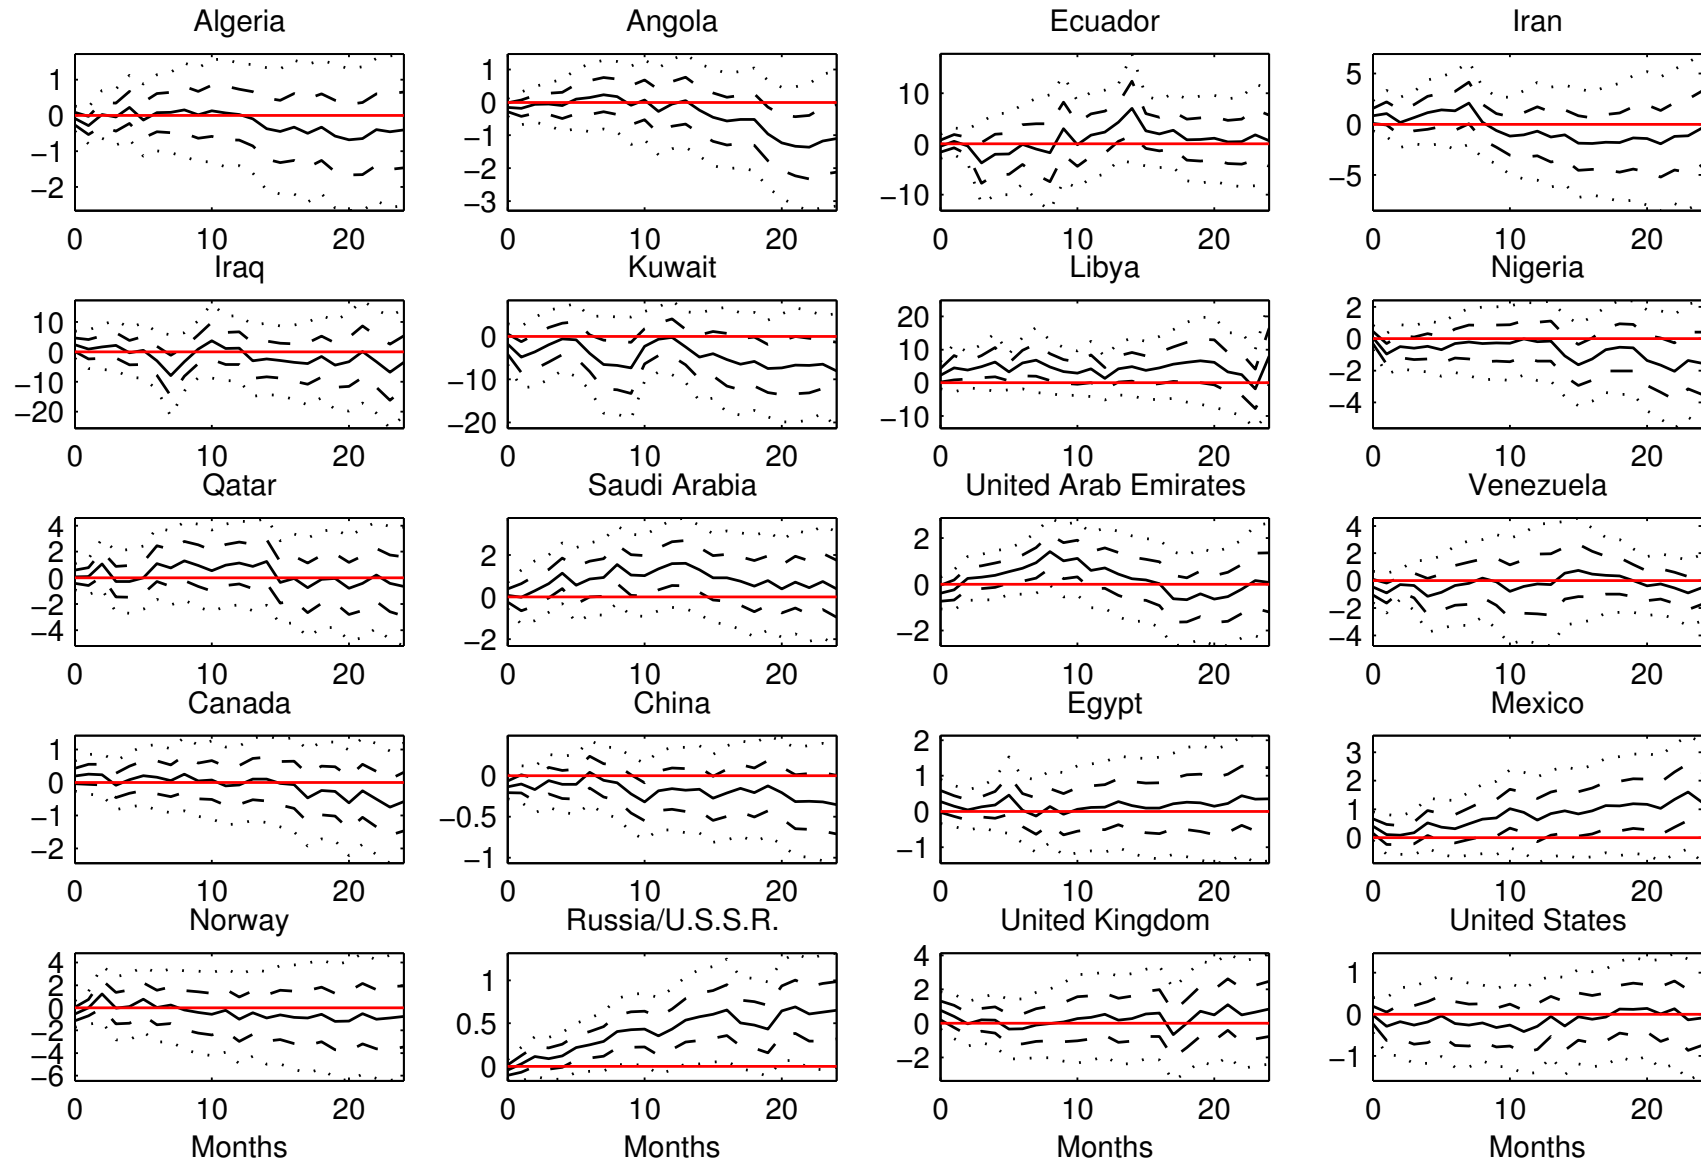


Figure 2: Cumulated impulse responses of country-level crude oil production to a flow demand shock based on the model in (1) (Point estimates with one- and two-standard-error confidence intervals)

Notes: Demand shock series based on admissible structural model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). Confidence intervals based on 10,000 replications of a block bootstrap with block size 12

Recall that a positive flow demand shock is associated with increased aggregate demand for industrial commodities, in general, due to the global business cycle. On average over the sample period, it leads to an immediate and sustained increase in the real price of oil and a transitory reduction in crude oil inventories. Nevertheless, the impact response of country-level production is not statistically different from zero at the approximate 95% confidence level for any of the twenty countries in the sample, indicating that oil producers do not respond to oil price changes arising from a flow demand shock in the very short run. This finding is consistent with the common perception that investments in the oil sector involve significant lead times (see, e.g., Hamilton, 2009a; Smith, 2009) and that the majority of oil producers is effectively producing at full capacity most of the time (see, e.g., Kilian, 2008a; Baumeister and Peersman, 2013). Even in the presence of spare capacity, however, high adjustment costs render it unprofitable to raise extraction volumes in the short run (compare Kilian, 2009; Kellogg, 2011).

Considering the dynamic behavior of crude oil producers, two strategies seem to prevail among OPEC member states. The majority of countries maintains a constant production level for the subsequent 24 months, whereas a few countries such as Saudi Arabia and the United Arab Emirates seem to expand supply. While this finding of a *price-stabilizing* response of Saudi Arabia to an increase in oil demand for consumption is consistent with the theoretical predictions in De Santis (2003)¹⁴, it might equally reflect the flexibility of producers with spare capacity to respond to economic incentives. Finally, there is some evidence of a lower production level in the medium run for Angola and Kuwait.

Similarly, the cumulated impulse response functions of most non-OPEC producers are virtually flat, with two noticeable exceptions. On average over the sample period, Mexico and Russia seem to expand their production in response to a positive flow demand shock, indicating that these countries might try to benefit from higher demand and the associated increase in the real price of oil. This behavior is weakly significant (at an approximate 68% confidence level) for Mexico and strongly significant (at an approximate 95% confidence level) for Russia.¹⁵

¹⁴Inspired by Adelman (1993), De Santis (2003) develops a computational general equilibrium (CGE) model to explain Saudi Arabian production behavior by adherence to its quota in the short run, while playing the role of a dominant firm in the long run. In the dominant firm model, Saudi Arabia determines the oil price, all other countries supply, and the dominant firm accommodates residual world demand.

¹⁵This finding is robust to controlling for a structural break occurring around the dissolution of the Union of Soviet Socialist Republics (U.S.S.R.).

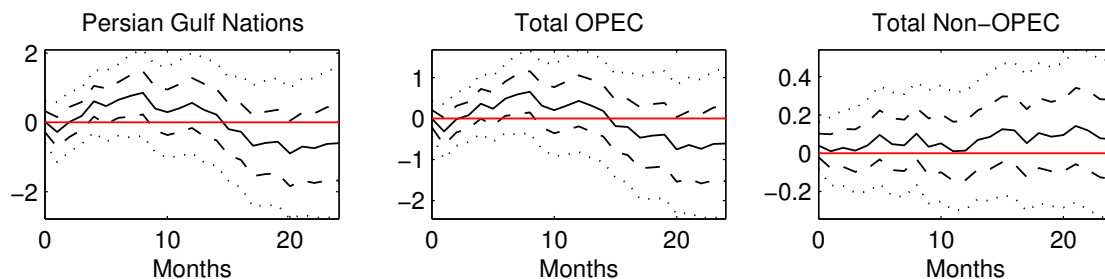


Figure 3: Cumulated impulse responses of crude oil production in a group of countries to a flow demand shock based on the model in (1) (Point estimates with one- and two-standard-error confidence intervals)

Notes: Demand shock series based on admissible structural model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). Confidence intervals based on 10,000 replications of a block bootstrap with block size 12

Figure 3 illustrates that the similarity in behavior of OPEC and non-OPEC oil producers in response to a typical flow demand shock carries over to the cumulated impulse response functions of aggregated country-level production, where I distinguish between *Persian Gulf Nations*¹⁶, *Total OPEC*, and *Total Non-OPEC*.

Note first that the response pattern of total OPEC production is virtually identical to that of Persian Gulf Nations. In line with world crude oil production, both are slightly negative in the very short run before increasing for about 8 months and finally dropping below the initial level after five quarters. While the initial expansion is at least weakly significant (at an approximate 68% confidence level), the subsequent undershooting is not statistically significant.

At the same time, the endogenous response of total non-OPEC production is largely flat with only a slight, statistically insignificant upward trend over the two year horizon. As we have seen in Figure 2, the latter is fully attributable to an expansion of Mexican and Russian oil supply. Note also the different orders of magnitude of the cumulated percent changes of total OPEC and total non-OPEC production levels in response to a typical flow demand shock, reflecting the higher volatility of OPEC production.

4.3 Responses to a Speculative Oil Demand Shock

Consider next the impulse response functions of country-level crude oil production to a typical speculative demand shock, which raises both the demand for crude oil inventories

¹⁶Following the EIA’s classification, the Persian Gulf Nations comprise Bahrain, Iran, Iraq, Kuwait, Qatar, Saudi Arabia, United Arab Emirates, and the Neutral Zone (between Kuwait and Saudi Arabia).

and the real price of oil, on impact. Figure 4 plots the point estimates based on the admissible structural model with a price elasticity in use closest to the posterior median of this elasticity in Kilian and Murphy (2013) together with one- and two-standard-error bootstrap confidence intervals.

Similar to the case of a positive flow demand shocks, on average over the sample period, none of the twenty countries in the sample seems to react strongly, in the short run. With the possible exceptions of Canada and Saudi Arabia, the impact responses to a typical speculative demand shock are not statistically different from zero at conventional significance levels.

At longer horizons, OPEC member states have a tendency to cut back on their physical supply of crude oil, albeit this reduction is strongly statistically significant (at an approximate 95% confidence level) only for Saudi Arabia. Intuitively, the finding that a country reduces supply in response to a positive demand shock seems to contradict the notion of Saudi Arabia trying to stabilize the oil price (see, e.g., De Santis, 2003). Here, it is important to keep in mind the distinct nature of the two structural innovations. Flow demand shocks reflect shifts in the demand *for consumption*, whereas speculative demand shocks reflect shifts in the demand *for inventories*. In principle, the latter might even amplify oil price fluctuations, if the temporary increase in crude oil inventories is quickly released to the market at a later point in time. Consistent with the “dominant firm” interpretation, Smith (2009) argues that the sluggish response of oil production in recent years reflects OPEC’s effort to support cartel discipline by not expanding production capacity despite ample reserves. According to Gately (2004), oil producers might deliberately refrain from expanding capacity in order to preserve higher revenues, when the price elasticity of oil demand is sufficiently low, in line with TRT.

Turning to non-OPEC oil producers, Egypt expands its production during the first year, whereas output levels in Canada, Norway, and the U.S. are broadly constant following a positive speculative demand shock. On average over the sample period, China, Mexico, Russia, and the U.K. respond by raising production levels between twelve and 24 months after the demand shock occurred, although the corresponding impulse response functions are statistically significant at an approximate 68% confidence level only.

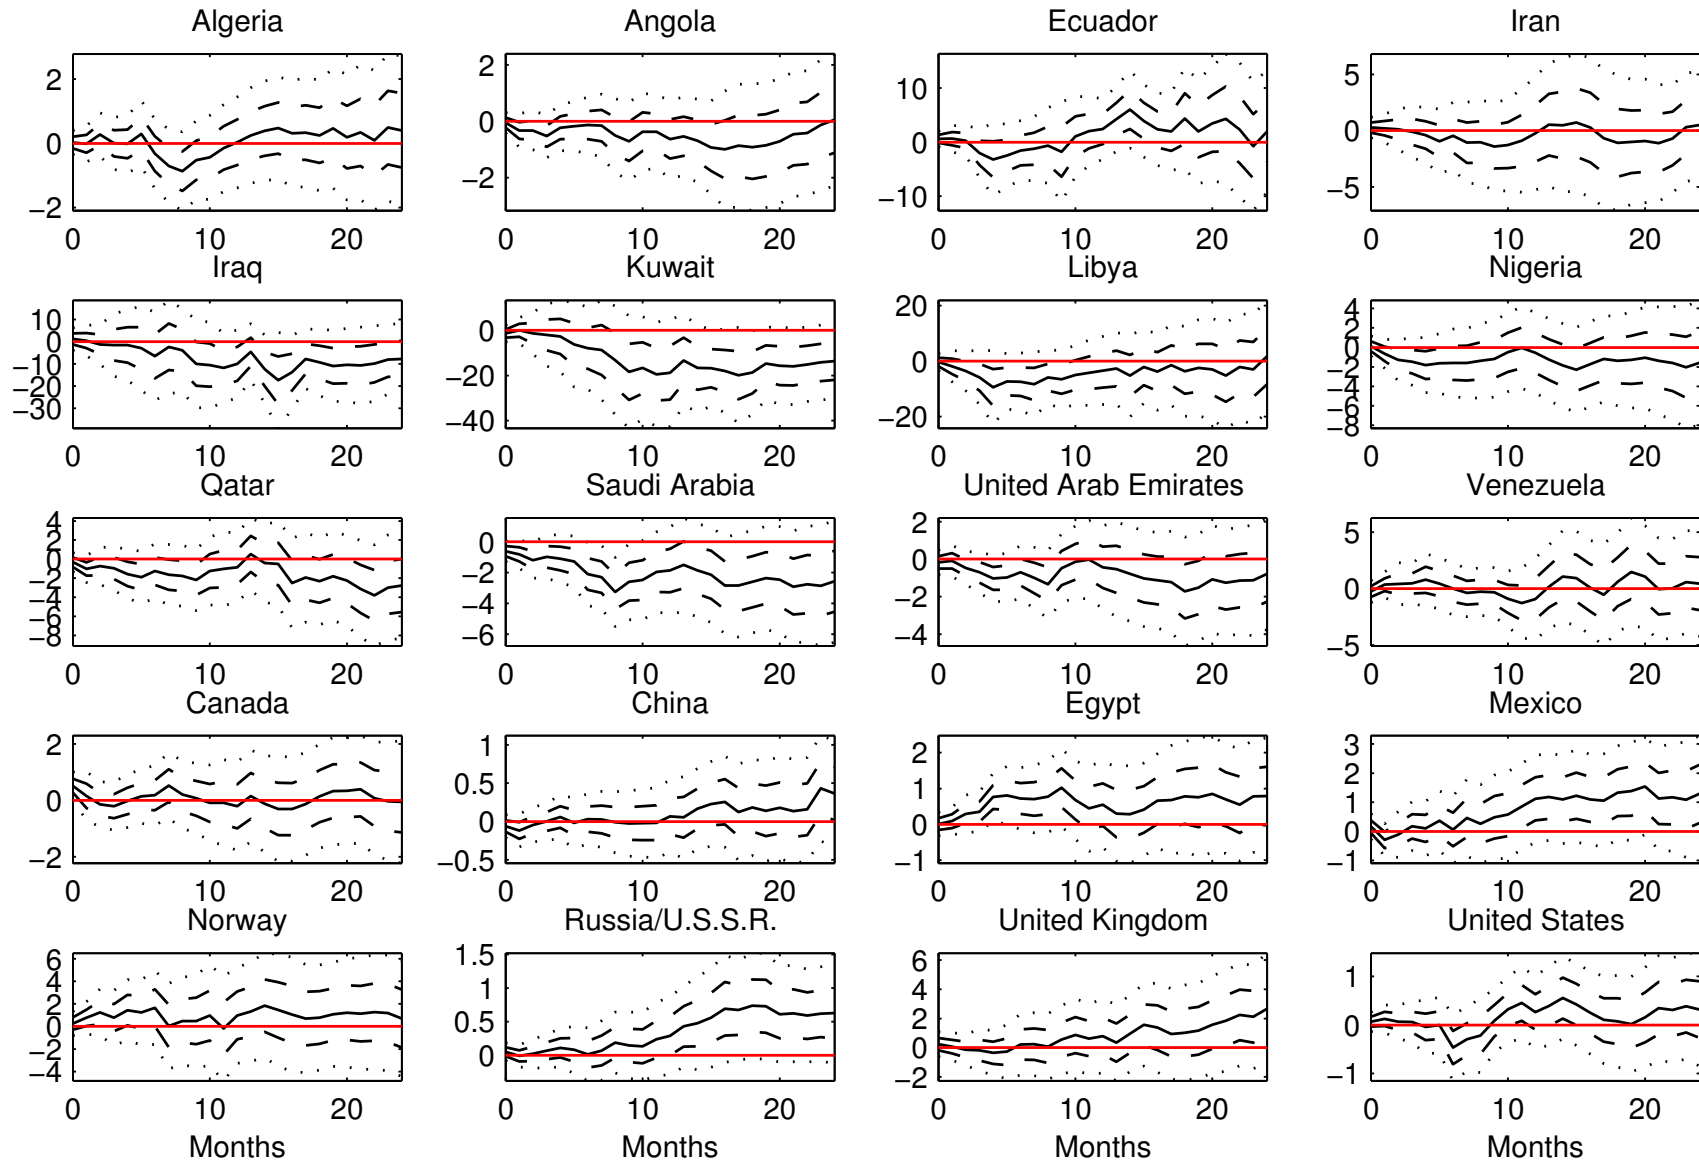


Figure 4: Cumulated impulse responses of country-level crude oil production to a speculative demand shock based on the model in (2) (Point estimates with one- and two-standard-error confidence intervals)

Notes: Demand shock series based on admissible structural model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). Confidence intervals based on 10,000 replications of a block bootstrap with block size 12




Figure 5: Cumulated impulse responses of crude oil production in a group of countries to a speculative demand shock based on the model in (2) (Point estimates with one- and two-standard-error confidence intervals)

Notes: Demand shock series based on admissible structural model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). Confidence intervals based on 10,000 replications of a block bootstrap with block size 12

The cross-country comparison suggests that OPEC and non-OPEC oil producers respond differently to a typical speculative demand shock. An interesting question is whether the impulse responses of aggregate production for these two groups of countries are different, as well. Figure 5 plots the cumulated percent change in oil production in response to a demand shock associated with an unexpected increase in inventories. Again, the behavior of Persian Gulf Nations' and total OPEC production is broadly identical. While neither responds on impact, the subsequent reduction in output levels is statistically significant (at an approximate 95% confidence level) for most of the forecast horizon.

On average over the sample period, the aggregate response of total non-OPEC is exactly the opposite. Following a marginally significant increase on impact, crude oil production expands in 22 of the subsequent 24 months and becomes statistically significant after about one year. As we have seen in Figure 4, the behavior across non-OPEC producers is relatively homogeneous. Most countries contribute to the expansion of total non-OPEC supply in response to a speculative demand shock.

In order to understand the differences in the responses of crude oil producers, recall that unexpected innovations in the demand for inventories above the ground arise, e.g., from a change in expectations about future oil supply relative to future oil demand or pure uncertainty about future supply or demand conditions without a change in the expected fundamentals. In many cases, speculative demand shocks can therefore be attributed to the actual or expected occurrence of exogenous political events such as the Invasion of Kuwait by Iraqi forces in August 1990.

On the one hand, a speculative demand shock might thus be followed by a disruption of physical supply in the countries involved, if (and only if) such expectations are correct. This explains, e.g., the sizeable negative impulse response of production in Iraq and Kuwait. On the other hand, countries such as Saudi Arabia or the United Arab Emirates have not been subject to similar historical events during 1975.2–2011.12. Why is it then that production is also falling in these countries, which account for a significant share of OPEC and world crude oil supply during the sample period?

By definition, a positive speculative demand shock induces higher *current demand* for oil without raising *current consumption*. Hamilton (2009a) emphasizes that “an ongoing speculative price bubble would have to result in continuous inventory accumulation, or else be ratified by cuts in production. The former is clearly unsustainable, [...]”. Figures 4 and 5 suggest that both occur simultaneously. On average over the sample period, (some) OPEC members reduce output to slow down inventory buildup, whereas non-OPEC producers respond by expanding production. When confronted with higher uncertainty about the future availability of crude oil, industrial economies such as the U.S. and the U.K. have a strong incentive to draw on their own reserves for building buffer stocks, in order to cushion possible future fluctuations in price and supply. Alternatively, they can rely on politically more stable suppliers such as Mexico and Russia. The latter, in turn, might be tempted to raise their output levels and pocket the windfall revenues from a temporary increase in the real price of oil.

Figures A.2 to A.5 in Appendix A illustrate that the above conclusions are qualitatively robust to deriving the underlying flow and speculative demand shocks from any of the 61 admissible structural models, conditional on the least squares estimate of the reduced form for 1973.2–2011.12. To improve the readability of the plot, all point estimates are shown without the corresponding bootstrap confidence intervals.

4.4 Are Demand Shocks Quantitatively Relevant?

The previous subsections suggest that oil producers respond to demand shocks, especially if they arise from a surge in demand for crude oil inventories. However, this finding is not informative about whether oil demand shocks are an important driver of fluctuations in country-level crude oil production. In order to address this question, Table 3 reports

Table 3: Percent Share of Total Variation in Country-Level Crude Oil Production Explained by Oil Demand Shocks

	Flow demand shock	Speculative demand shock	Total
Algeria	4.16	8.03	12.19
Angola	4.55	4.73	9.28
Ecuador	9.73	6.29	15.66
Iran	4.92	2.28	7.20
Iraq	6.47	9.44	15.91
Kuwait	5.95	10.73	16.69
Libya	12.27	7.17	19.45
Nigeria	5.06	4.52	9.58
Qatar	5.91	7.27	13.18
Saudi Arabia	5.38	8.21	13.60
United Arab Emirates	6.49	7.20	13.68
Venezuela	4.41	9.89	14.30
Canada	3.13	4.21	7.35
China	6.18	6.03	12.21
Egypt	5.35	5.82	11.17
Mexico	4.46	5.22	9.68
Norway	3.76	4.16	7.92
Russia/U.S.S.R.	6.30	4.74	11.05
United Kingdom	7.88	4.29	12.17
United States	6.83	8.39	15.21
Persian Gulf Nations	5.70	9.85	15.55
Total OPEC	6.27	8.92	15.19
Total Non-OPEC	3.51	5.87	9.38

Notes: R^2 of distributed lag regressions in (1) and (2). Orthogonal flow and speculative demand shocks based on candidate VAR model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013).

the R^2 of the distributed lag regression in (1) and (2), respectively.¹⁷

The contribution of flow demand shocks ranges from 3.13% for Canada to more than 12% for Libya, with no obvious difference between individual OPEC and non-OPEC producers. Nevertheless, flow demand shocks explain almost twice as much of the monthly fluctuations in total OPEC relative to total non-OPEC output, indicating countervailing responses among the latter. Note also that flow demand shocks seem to be less important for aggregate production in the Persian Gulf region than for OPEC as a whole.

A comparison of columns two and three suggests that speculative demand shocks have relatively greater explanatory power, in general, ranging from 4.16% for Norway to almost 11% for Kuwait. This finding holds for OPEC and non-OPEC producers alike and carries over to the corresponding aggregates.

¹⁷The two-step estimation approach does not allow for a conventional forecast error variance decomposition (FEVD) of the percent change in country-level crude oil production.

Due to the fact that flow and speculative demand shocks are orthogonal by construction, columns two and three add up to the total share of demand-induced oil price changes in the variation of $\Delta prod_{i,t}$. Such shifts in the oil demand curve along the supply curve explain between 7.35 and 19.45% of the fluctuations in crude oil production. The corresponding shares for total OPEC and Persian Gulf Nations, in particular, exceed 15%, indicating a nontrivial role for oil demand shocks in the variance of supply.

4.5 Country-Level Price Elasticities of Oil Supply

In this section, I combine the results from the structural VAR model of the global market for crude oil and the distributed lag model in order to derive *country-level* estimates of the impact price elasticity of crude oil supply in response to flow and speculative demand shocks, respectively. For comparable estimates of the price elasticity of *world* crude oil supply, see, e.g., Baumeister and Peersman (2013a) and Kilian and Murphy (2013).

The textbook price elasticity of supply corresponds to the ratio between the percentage change in the quantity supplied and the percentage change in the price of a commodity, conditional on an exogenous shift of the demand curve along the supply curve. Note that the flow and speculative demand shocks from the VAR model are orthogonal to flow supply shocks by construction. Although the model yields thus an estimate of the price elasticity of world crude oil supply, this elasticity is bounded above for the sake of identification. In contrast, no restrictions are imposed on the responses of country-level production in the distributed lag model.

Recall that the real price of oil is expressed in monthly log levels, while the production data are in monthly log differences. Accordingly, estimates of the *impact price elasticity of crude oil supply* for country i can be derived from the relative impact response of oil production and the real price of oil to each of the two oil demand shocks:

$$\eta_i^j = \frac{\Delta prod_{i,0}^j}{rpo_0^j}, \quad j = \text{flow demand, speculative demand}, \quad (3)$$

where subscript 0 denotes the impulse response to shock j within the same month.¹⁸

¹⁸Note that only the impact response of the real price of oil to shock j is exogenous from oil producers' perspective. Over time, other variables adjust due to general equilibrium effects. As a consequence,

$$\frac{\sum_{h=0}^H \Delta prod_{i,h}^j}{rpo_H^j}, \quad j = \text{flow demand, speculative demand},$$

represents a relative impulse response function rather than a price elasticity of supply at horizon H .

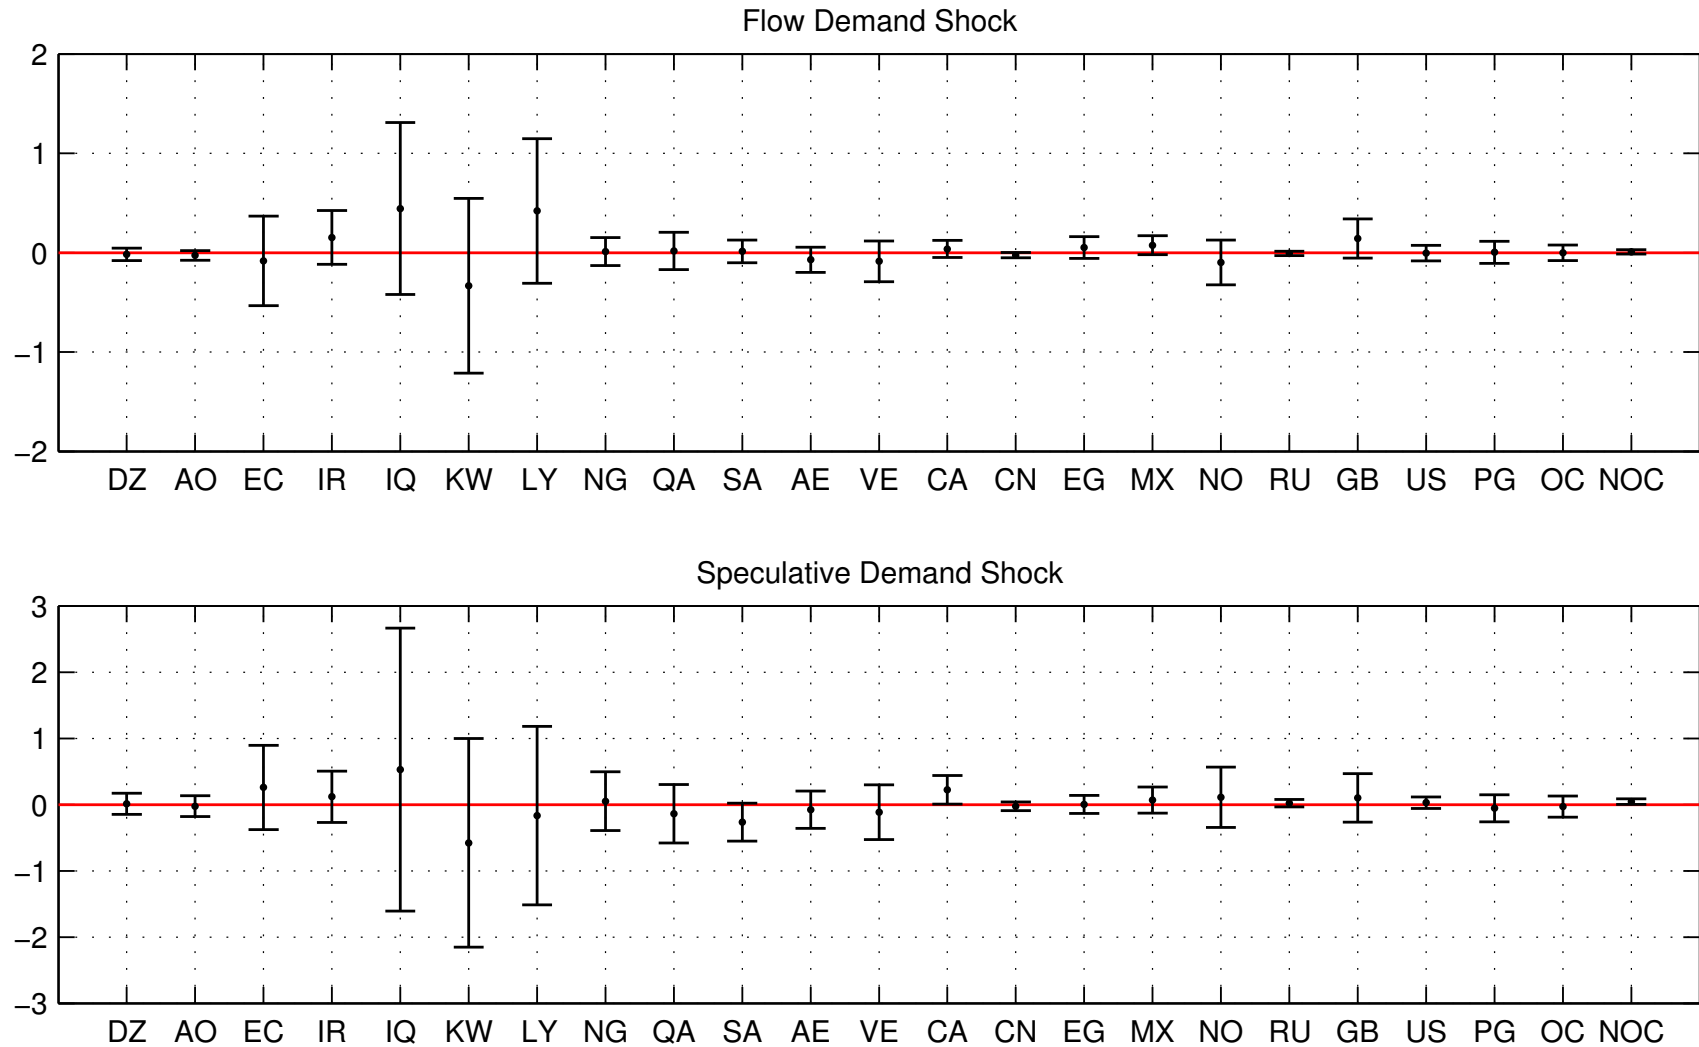


Figure 6: Impact price elasticities of country-level crude oil supply in response to flow and speculative oil demand shocks (Point estimates with approximate 95% confidence intervals)

Notes: Demand shock series based on admissible structural model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). Confidence intervals based on 10,000 replications of a block bootstrap with block size 12. Country names abbreviated using ISO 2-letter codes, except: PG – Persian Gulf Nations, OC – Total OPEC, NOC – Total Non-OPEC

Figure 6 reports the point estimates of the impact price elasticity of crude oil supply for country or group i based on the impact responses of country-level production and the real price of oil to both demand shocks for the admissible structural model used in the previous analysis. Approximate 95% confidence intervals are inferred from the bootstrap confidence intervals around $\Delta prod_{i,0}^j$.¹⁹

Several points are worth stressing. First, the impact price elasticity of country-level crude oil supply is estimated with strongly varying precision. Table 4 shows that only three out of 46 point estimates are statistically different from 0 at the 5% level. Accordingly, a negligible response of world crude oil production to oil demand shocks within the same month does *not* cover large countervailing responses at the country level. This finding is in line with the consensus in the literature that oil producers do not react to noisy price signals, unless there is significant excess capacity (see, e.g., Hamilton, 2009a). Kilian (2009) argues that this short-run elasticity is close to zero even in the presence of spare capacity, due to sizeable adjustment costs in oil production (see also Kellogg, 2011).

Second, countries with a relatively imprecise estimate have experienced exogenous disruptions of production due to natural disasters such as the 1987 Ecuador earthquakes or geopolitical events such as the Iran-Iraq War, the Persian Gulf Wars, or the Venezuela Oil Strike of 2002–2003. Accordingly, their extraction levels are inherently more volatile, on average over the sample period. Aggregating country-level data increases the precision of point estimates, indicating that fluctuations in crude oil production are largely idiosyncratic. Note that total OPEC production seems to be less volatile than the aggregate of Persian Gulf Nations, while non-OPEC production is even less volatile.

Third, although most point estimates are not statistically different from zero, Table 4 reveals a qualitative pattern for the case of speculative demand shocks. While the short-run price elasticities of crude oil supply are equally split between positive and negative point estimates for OPEC member states, six out of eight non-OPEC producers respond by expanding production levels within the same month. Notwithstanding capacity constraints and adjustment costs at the country level, the aggregate impact price elasticity

¹⁹It is important to note that there is no straightforward way of accounting for the uncertainty about the impact response of the real price of oil and the generated regressor issue, i.e., the bootstrap confidence intervals could understate the true variance of the point estimates. Given that the impact price elasticity of oil supply is not statistically different from zero for most countries i , however, my conclusions are fully robust to wider confidence intervals.

Table 4: Impact Price Elasticities of Country-Level Crude Oil Supply in Response to Oil Demand Shocks

	Flow demand shock	Speculative demand shock
Algeria	-.0169	.0120
Angola	-.0279	-.0250
Ecuador	-.0820	.2586
Iran	.1532	.1182
Iraq	.4444	.5273
Kuwait	-.3342	-.5781
Libya	.4196	-.1666
Nigeria	.0120	.0508
Qatar	.0165	-.1389
Saudi Arabia	.0130	-.2648*
United Arab Emirates	-.0710	-.0771
Venezuela	-.0869	-.1146
Canada	.0370	.2234**
China	-.0252**	-.0263
Egypt	.0519	.0022
Mexico	.0749*	.0701
Norway	-.0996	.1105
Russia/U.S.S.R.	-.0079	.0223
United Kingdom	-.1412	.1000
United States	-.0042	.0293
Persian Gulf Nations	.0035	-.0548
Total OPEC	-.0004	-.0310
Total Non-OPEC	.0075	.0448**

Notes: Point estimates of impact price elasticities of crude oil supply. Orthogonal flow and speculative demand shocks based on candidate VAR model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). *, **, *** indicate statistical significance at a level of 10, 5, and 1%.

of non-OPEC supply is thus positive and statistically significant at the 5% level.

Figure A.6 illustrates that the presented elasticities are representative for most of the 61 admissible structural models, conditional on the least squares estimate of the reduced-form model for 1973.2–2011.12. As a consequence, it seems robust to conclude that the short-run price elasticity of country-level crude oil supply is statistically indistinguishable from zero at conventional significance levels, i.e., most *individual* oil producing countries do not respond to demand-induced oil price changes within the same month.

5 Concluding Remarks

The present paper decomposes world crude oil production by country in order to investigate whether oil producers respond to demand-induced changes in the real price of oil. I

focus on the potential differences in the behavior of OPEC and non-OPEC or emerging and developed oil producing economies during 1975.2–2011.12. Monthly time series of flow demand and speculative demand shocks, which are orthogonal to innovations in *world* crude oil supply shocks by construction, are drawn from Kilian and Murphy’s (2013) model of the global oil market. The resulting demand shocks are then used as regressors in a distributed lag model of country-level production.

I find that oil producers do not respond to demand-driven price signals in the very short run, consistent with the notion of high adjustment costs in crude oil production. Given that both flow and speculative demand shocks raise the real price of oil, the finding of an almost vertical short-run supply curve further implies an impact price elasticity of oil producers close to zero, regardless of the underlying demand shock. Also, the consensus view of a near-zero impact price elasticity of *world* crude oil supply in the literature is not due to aggregation bias, i.e. countervailing elasticities at the country level.

While the dynamic impulse responses to a typical flow demand shock are largely flat and statistically insignificant for both OPEC and non-OPEC production, oil producers seem to respond to speculative demand shocks, albeit slowly. On average over the sample period, supply by OPEC as a whole decreases, whereas total non-OPEC production increases. Given that speculative demand shocks are associated with an accumulation of inventories above the ground, dominant producers in the global market for crude oil, such as Saudi Arabia, have an incentive to curtail production in order to avoid future excess supply, although this might destabilize the oil price in the short run. On the other hand, unexpected changes in inventory demand were frequently associated with the actual or expected occurrence of exogenous political events such as the Iranian Revolution. From the point of view of non-OPEC producers, the resulting uncertainty about future supply relative to future demand of crude oil entails a powerful motive for countries such as the U.S. or the U.K. to build up buffer stocks by expanding domestic production or by increasingly relying on politically more stable suppliers, e.g. Mexico and Russia.

The present cross-country comparison of oil producers’ behavior and the corresponding impact price elasticities of supply complements two strands of the recent empirical literature on structural models of the global market for crude oil. On the one hand, previous research has focused on the price elasticity of oil demand and the effects of exogenous

supply shocks on real economic activity (see, e.g., Edelstein and Kilian, 2009; Peersman and Van Robays, 2012; Baumeister and Peersman, 2013b). On the other hand, the price elasticity of supply has been considered only at the global level, as in Baumeister and Peersman (2013a), or used as an identifying restriction in a structural VAR model, as in Kilian and Murphy (2012, 2013). In contrast, my focus was on the production behavior and the price elasticities of crude oil supply at the country level in response to flow and speculative oil demand shocks.

While the present paper distinguishes flow demand and speculative demand shocks as well as twelve OPEC and eight non-OPEC countries, accommodating thus heterogeneous impulse responses along two important dimensions, it assumes a time-invariant behavior of individual crude oil producers. In the light of continuous change in the institutional setting of the global market for crude oil throughout the sample period, e.g. due to the OPEC's loss of influence, a promising direction for future research is the extension of the present analysis to allow for time-varying coefficients in the structural VAR and the distributed lag model (compare Baumeister and Peersman, 2013a,b).

References

- Adelman, M. A., 1993. Modelling World Oil Supply. *Energy Journal* 14 (1), 1–32.
- Alhajji, A. F., Huettner, D., 2000. The Target Revenue Model and the World Oil Market: Empirical Evidence from 1971 to 1994. *Energy Journal* 21 (2), 121–144.
- Alquist, R., Kilian, L., 2010. What do we learn from the price of crude oil futures? *Journal of Applied Econometrics* 25 (4), 539–573.
- Barsky, R. B., Kilian, L., 2002. Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative. In: Bernanke, B. S., Rogoff, K. (Eds.), *NBER Macroeconomics Annual 2001*. Vol. 16. MIT Press, Cambridge, MA, pp. 137–183.
- Barsky, R. B., Kilian, L., 2004. Oil and the Macroeconomy since the 1970s. *Journal of Economic Perspectives* 18 (4), 115–134.
- Baumeister, C., Peersman, G., 2013a. The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market. Forthcoming in: *Journal of Applied Econometrics*.
- Baumeister, C., Peersman, G., 2013b. Time-Varying Effects of Oil Supply Shocks on the U.S. Economy. Forthcoming in: *American Economic Journal: Macroeconomics*.
- De Santis, R. A., 2003. Crude oil price fluctuations and Saudi Arabia's behaviour. *Energy Economics* 25 (2), 155–173.
- Edelstein, P., Kilian, L., 2009. How sensitive are consumer expenditures to retail energy prices? *Journal of Monetary Economics* 56 (6), 766–779.
- Gately, D., 2004. OPEC's Incentives for Faster Output Growth. *Energy Journal* 25 (2), 75–96.
- Griffin, J. M., 1985. OPEC Behavior: A Test of Alternative Hypotheses. *American Economic Review* 75 (5), 954–963.
- Hamilton, J. D., 2008. Oil and the Macroeconomy. In: Durlauf, S. N., Blume, L. E. (Eds.), *The New Palgrave Dictionary of Economics*, 2nd Edition. Palgrave Macmillan.
- Hamilton, J. D., 2009a. Understanding Crude Oil Prices. *Energy Journal* 30 (2), 179–206.
- Hamilton, J. D., 2009b. Causes and Consequences of the Oil Shock of 2007-08. *Brookings Papers on Economic Activity*, 1 (Spring), 215–261.
- Hausman, J. A., Newey, W. K., 1995. Nonparametric Estimation of Exact Consumers Surplus and Deadweight Loss. *Econometrica* 63 (6), 1445–1476.

- Jones, C. T., 1990. OPEC Behavior Under Falling Prices: Implications for Cartel Stability. *Energy Journal* 11 (3), 117–129.
- Kellogg, R., 2011. Short-Run Responses of Oil Production to Spot and Future Oil Prices. mimeo, University of Michigan.
- Kilian, L., 2008a. Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the US Economy? *Review of Economics and Statistics* 90 (2), 216–240.
- Kilian, L., 2008b. The Economic Effects of Energy Price Shocks. *Journal of Economic Literature* 46 (4), 871–909.
- Kilian, L., 2009. Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review* 99 (3), 1053–1069.
- Kilian, L., Lee, T. K., 2013. Quantifying the Speculative Component in the Real Price of Oil: The Role of Global Oil Inventories. Forthcoming in: *Journal of International Money and Finance*.
- Kilian, L., Murphy, D. P., 2012. Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models. *Journal of the European Economic Association* 10 (5), 1166–1188.
- Kilian, L., Murphy, D. P., 2013. The Role of Inventories and Speculative Trading in the Global Market for Crude Oil. Forthcoming in: *Journal of Applied Econometrics*.
- Peersman, G., Robays, I. V., 2012. Cross-country differences in the effects of oil shocks. *Energy Economics* 34 (5), 1532–1547.
- Ramcharran, H., 2001. OPEC's production under fluctuating oil prices: further test of the target revenue theory. *Energy Economics* 23 (6), 667–681.
- Ramcharran, H., 2002. Oil production responses to price changes: an empirical application of the competitive model to OPEC and non-OPEC countries. *Energy Economics* 24 (2), 97–106.
- Smith, J. L., 2009. World Oil: Market or Mayhem? *Journal of Economic Perspectives* 23 (3), 145–164.
- Yatchew, A., No, J. A., 2001. Household Gasoline Demand in Canada. *Econometrica* 69 (6), 1697–1709.

Appendix A Robustness of Results




Figure A.1: Historical evolution of orthogonal oil supply and demand shocks for 1975Q2–2011Q4

Note: Quarterly averages of monthly structural residuals based on *all* candidate models satisfying the identifying assumptions

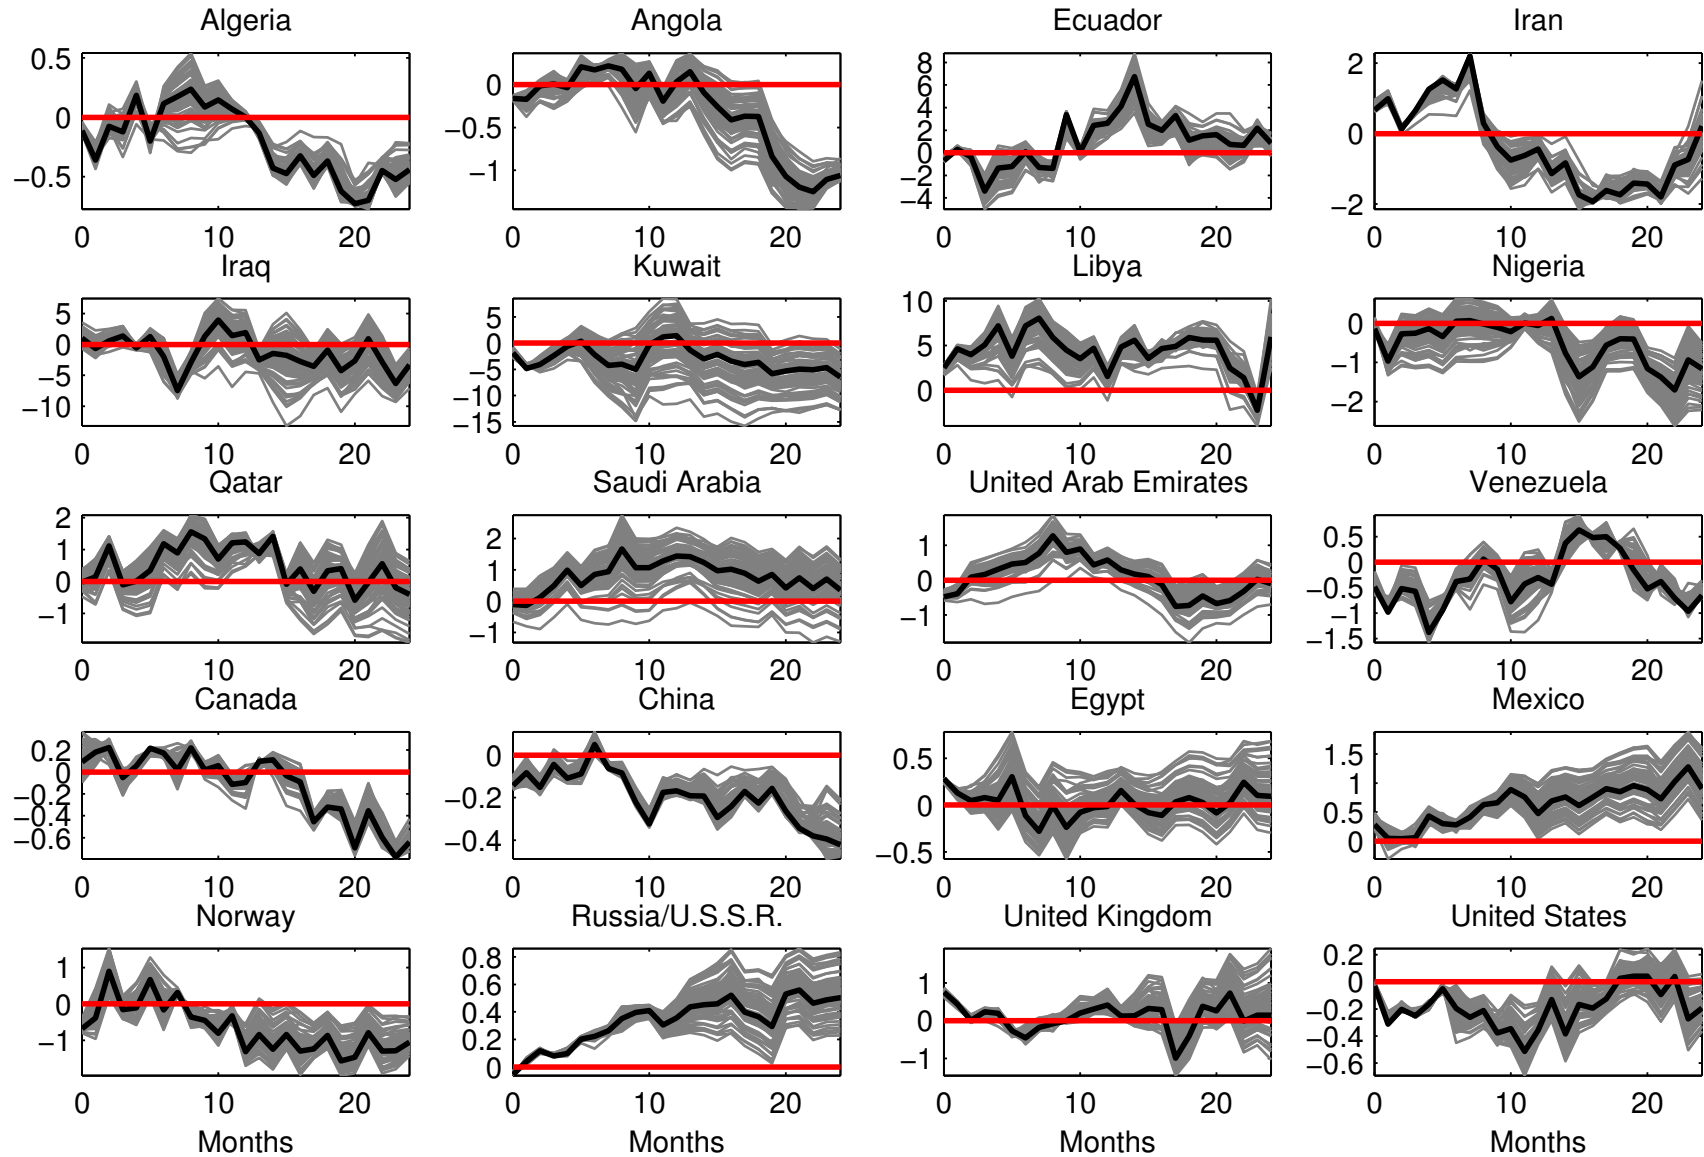


Figure A.2: Cumulated impulse responses of country-level crude oil production to a flow demand shock based on the model in (1) (Point estimates)

Note: Results based on *all* candidate models satisfying the identifying assumptions

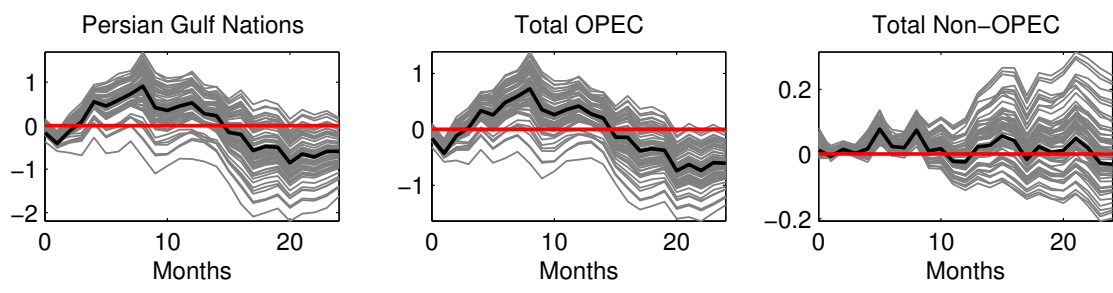


Figure A.3: Cumulated impulse responses of crude oil production in a group of countries to a flow demand shock based on the model in (1) (Point estimates)

Note: Results based on *all* candidate models satisfying the identifying assumptions

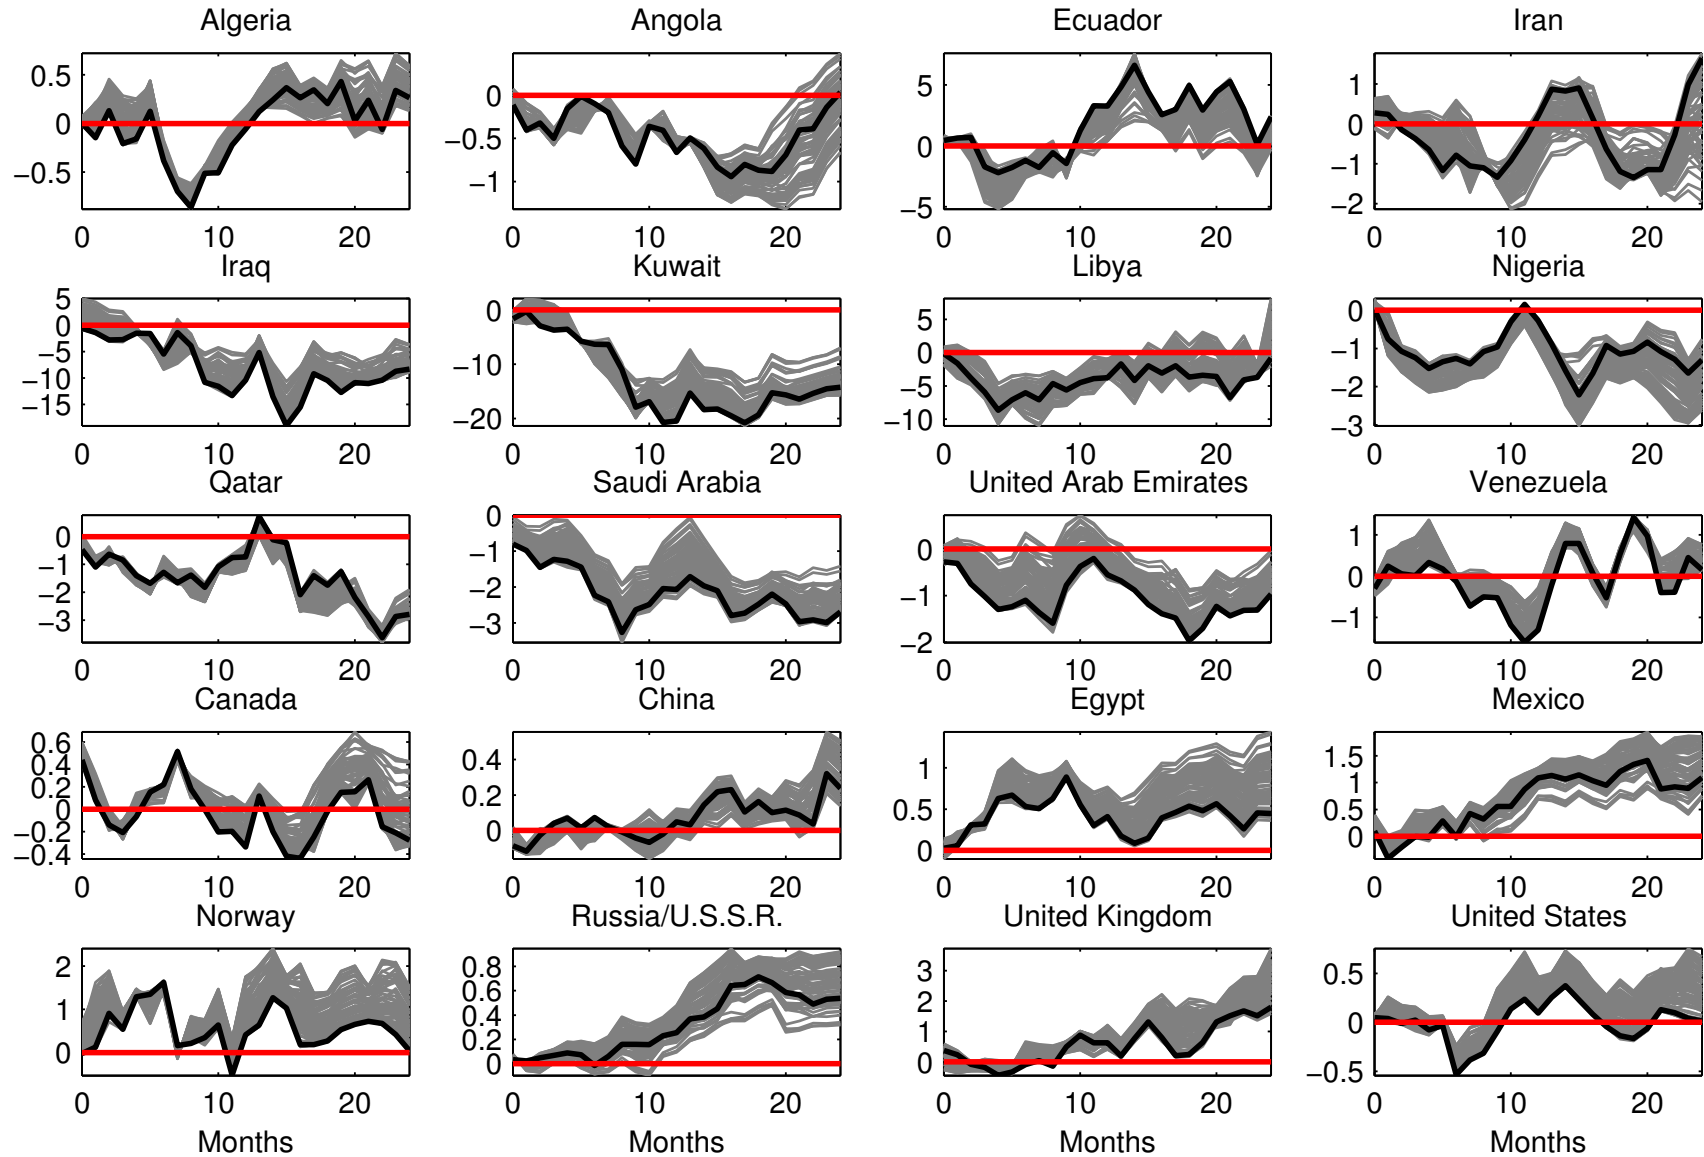


Figure A.4: Cumulated impulse responses of country-level crude oil production to a speculative demand shock based on the model in (2) (Point estimates)

Note: Results based on *all* candidate models satisfying the identifying assumptions




Figure A.5: Cumulated impulse responses of crude oil production in a group of countries to a speculative demand shock based on the model in (2) (Point estimates)

Note: Results based on *all* candidate models satisfying the identifying assumptions

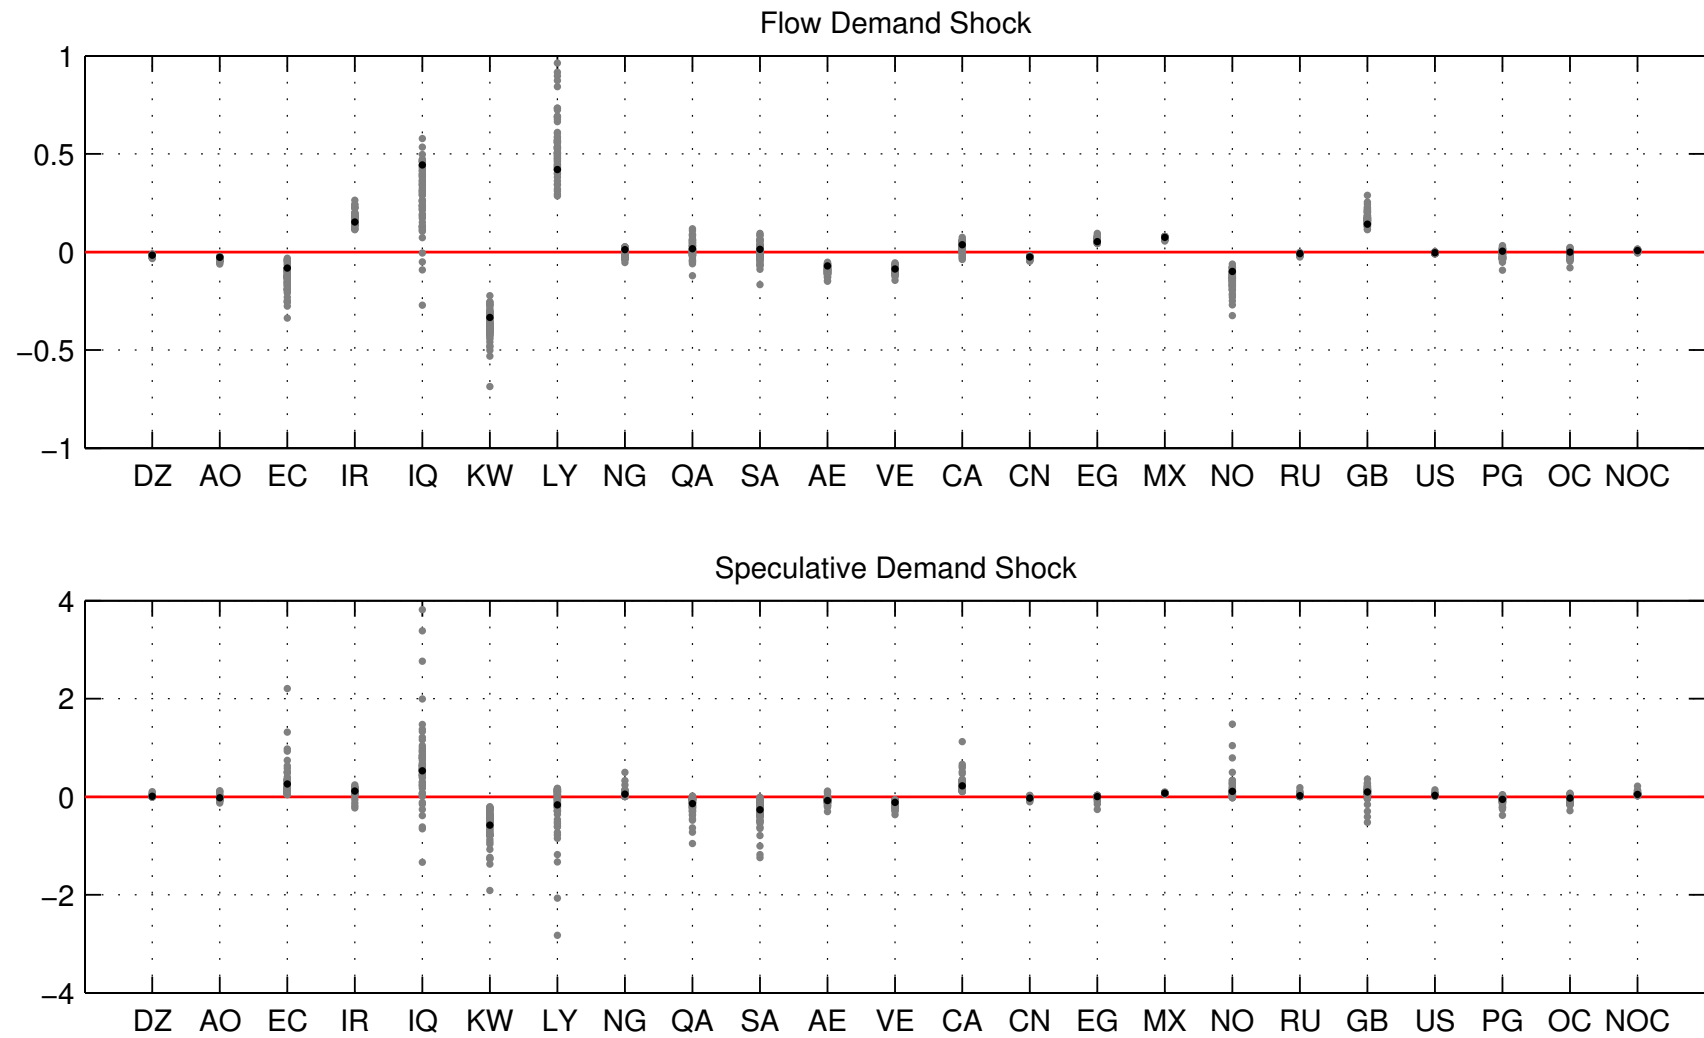


Figure A.6: Impact price elasticities of country-level crude oil supply in response to flow and speculative oil demand shocks (Point estimates)

Notes: Results based on *all* candidate models satisfying the identifying assumptions. Country names abbreviated using ISO 2-letter codes, except: PG – Persian Gulf Nations, OC – Total OPEC, NOC – Total Non-OPEC

Appendix B Kilian and Murphy's (2013) SVAR Model of the Global Oil Market

This appendix summarizes the methodology of the structural VAR model of the global oil market employed in Section 2. For further details and discussion, the reader is referred to Kilian and Murphy (2013). The derivation of orthogonal oil demand shocks is based on a four-variable VAR(24) model of the global market for crude oil,

$$\mathbf{A}_0 \mathbf{y}_t = \alpha + \sum_{l=1}^{24} \mathbf{A}_l \mathbf{y}_{t-l} + \varepsilon_t, \quad (\text{B.1})$$

in the vector $\mathbf{y}_t = (\Delta prod_t, rea_t, rpo_t, \Delta inv_t)'$, where $\Delta prod_t$ denotes the percent change in world crude oil production, rea_t an index of global real economic activity proposed by Kilian (2009), rpo_t the nominal oil price deflated by the U.S. consumer price index, and Δinv_t the change in above-ground inventories of crude oil. ε_t is a vector of orthogonal structural disturbances. A set of monthly dummies which accounts for the seasonality in crude oil inventories is omitted here for notational convenience.

Conditional on \mathbf{A}_0 being invertible, the reduced-form representation of (1) is given by

$$\begin{aligned} \mathbf{y}_t &= \underbrace{\mathbf{A}_0^{-1} \alpha}_{\nu} + \underbrace{\mathbf{A}_0^{-1} \mathbf{A}_1 \mathbf{y}_{t-1}}_{\mathbf{B}_1 \mathbf{y}_{t-1}} + \dots + \underbrace{\mathbf{A}_0^{-1} \mathbf{A}_{24} \mathbf{y}_{t-24}}_{\mathbf{B}_{24} \mathbf{y}_{t-24}} + \underbrace{\mathbf{A}_0^{-1} \varepsilon_t}_{\mathbf{e}_t} \\ &= \nu + \mathbf{B}_1 \mathbf{y}_{t-1} + \dots + \mathbf{B}_{24} \mathbf{y}_{t-24} + \mathbf{e}_t \\ &= \nu + \sum_{l=1}^{24} \mathbf{B}_l \mathbf{y}_{t-l} + \mathbf{e}_t, \end{aligned} \quad (\text{B.2})$$

where \mathbf{e}_t denotes a vector of possibly contemporaneously correlated innovations, while ν_t and \mathbf{B}_l , $l = 1, \dots, 24$, are the intercept and slope coefficients of the reduced-form VAR. Straightforward multivariate least squares estimation of (2) yields consistent estimates of the matrix of coefficients, $\hat{\mathbf{B}} \equiv [\hat{\nu} \hat{\mathbf{B}}_1 \dots \hat{\mathbf{B}}_{24}]$, and the reduced-form disturbances, $\hat{\mathbf{e}}_t$.

Appendix B.1 Data

The VAR(24) in equation (1) is estimated using monthly data for 1973.2 through 2011.12 of the variables in $\mathbf{y}_t = (\Delta prod_t, rea_t, rpo_t, \Delta inv_t)'$. World crude oil production levels in thousand barrels per day (tbpd) are obtained from the *U.S. Energy Information Administration* (EIA). For rea_t , I use Kilian's (2009) index of global real economic activity, based on single voyage bulk dry cargo ocean shipping freight rates.

Following the literature, a global measure of the real price of oil, rpo_t , can be based on

the nominal U.S. refiner acquisition cost of imported crude oil as reported by the EIA since January 1974. The nominal time series is deflated by the seasonally adjusted U.S. consumer price index (CPI) published by the *U.S. Bureau of Labor Statistics* (BLS). In line with Kilian (2009), the real price of crude oil is backcasted from 1974.1 to 1973.1 using the method proposed in Barsky and Kilian (2002). The resulting time series is expressed in log levels in deviations from the mean.

All inventory time series are in million barrels. Since the EIA provides data on “petroleum stocks” for OECD countries only, the total of OECD stocks must serve as a proxy for world crude oil inventories.²⁰ Monthly data on crude oil inventories going back to January 1973 are available for the U.S. only. Hence, I follow Hamilton (2009b) and Kilian and Murphy (2013) and extrapolate total OECD petroleum stocks backwards using the ratio of OECD over U.S. stocks prior to December 1987. The resulting time series is expressed in absolute rather than percent changes, as the percent change in inventories seems to be non-stationary (compare Kilian and Murphy, 2013).

Appendix B.2 Identification

For the subsequent analysis, we are not interested in the reduced-form estimates, but in the structural shocks in the vector ε_t . By definition, $\mathbf{e}_t = \mathbf{A}_0^{-1}\varepsilon_t$. Hence, the covariance matrix of reduced-form innovations is

$$\begin{aligned} \underbrace{E(\mathbf{e}_t\mathbf{e}_t')}_{\Sigma_{\mathbf{e}_t}} &= \mathbf{A}_0^{-1} \underbrace{E(\varepsilon_t\varepsilon_t')}_{\Sigma_{\varepsilon_t}} (\mathbf{A}_0^{-1})' \\ &= \mathbf{A}_0^{-1} \Sigma_{\varepsilon_t} (\mathbf{A}_0^{-1})' = \mathbf{A}_0^{-1} (\mathbf{A}_0^{-1})', \end{aligned} \quad (\text{B.3})$$

where the last step assumes $\Sigma_{\varepsilon_t} = I_4$, i.e., the diagonal entries of the covariance matrix of structural shocks are normalized to 1, while the diagonal of \mathbf{A}_0^{-1} remains unrestricted. Note that we can replace the expression on the left by its consistent least-squares estimate, $\hat{\Sigma}_{\mathbf{e}_t}$.

Following Kilian and Murphy (2012, 2013), the structural shocks in ε_t are identified by a combination of sign restrictions on impulse response functions and bounds on the impact price elasticities of oil supply and oil demand. Note that I adapt the definitions and the terminology from Kilian and Murphy (2013).

²⁰Note that “petroleum stocks” reported by the EIA comprise crude oil (including strategic reserves), unfinished oils, natural gas, plant liquids, and refined products.

Appendix B.2.1 Sign Restrictions

Flow supply shocks relate to unexpected shifts of the oil supply curve along the oil demand curve, conditional on past realizations of \mathbf{y}_t . This incorporates any unexpected disruptions of the physical supply of crude oil due to, e.g., exogenous political events in oil producing countries or strategic production decisions of OPEC member states. In line with economic theory, a negative supply shock reduces the level of output and raises the real price of oil, on impact. Moreover, in the data inventories decrease by less than production in response to a negative flow supply shock, implying imperfect consumption smoothing. As a consequence, the flow demand for consumption of crude oil must also fall, which leads to a decrease in global real economic activity. As in Kilian and Murphy (2013), world crude oil production and global real economic activity are assumed to decrease, while the real price of oil is assumed to increase, respectively, for at least twelve months following a negative oil supply shock. The latter *dynamic sign restrictions* are crucial to exclude candidate models, in which rpo_t decreases below and rea_t increases above its initial level after the impact period.

Flow demand shocks capture unexpected changes in the demand for industrial commodities, in general, that are associated with fluctuations in the global business cycle. Accordingly, they correspond to a shift of the oil demand curve along the supply curve, conditional on all past observations of the endogenous variables. Economic theory predicts that a positive demand shock raises global real economic activity and the real price of oil, on impact. Given that crude oil inventories decrease by less than the increase in consumption, world production must also expand within the same month.

Note that the response of crude oil inventories to both flow supply and flow demand shocks remains unrestricted. In contrast, a *speculative demand shock* is characterized by an unexpected shift in the demand for crude oil inventories above the ground. This includes precautionary demand shocks, i.e. pure uncertainty shocks that do not affect the expected levels of future demand or supply, as well as changes in the expected future demand for crude oil *relative to* the expected future supply of crude oil. Expectations of political turmoil in oil producing countries such as Libya represent a typical example. As a consequence, a positive speculative demand shock raises crude oil inventories and the real price of oil, on impact. In line with Kilian and Murphy (2013), I assume that

the increase in inventories can only be accommodated by a simultaneous expansion of supply and a decrease in consumption, i.e. a reduction in global real economic activity. All impact and dynamic sign restrictions are imposed as weak inequality constraints. No restrictions are imposed on the impulse responses of endogenous variables to the fourth shock in ε_t . Accordingly, the latter represents a residual category, which absorbs other oil demand shocks not otherwise accounted for by the model. Since it lacks a clear economic interpretation, I will not report the associated results below.

Appendix B.2.2 Bounding the Impact Price Elasticity of Oil Supply

The structural shocks obtained from an identification based on (impact and dynamic) sign restrictions are set-identified, by construction. Kilian and Murphy (2012) show that sign restrictions alone are generally not sufficient for drawing unambiguous conclusions, e.g. about impulse responses, as they assign equal prior weight to candidate models with quantitatively implausible implications. For this reason, the authors propose to narrow down the set of admissible structural models by imposing an upper bound on the impact price elasticities of oil supply, which is directly related to the slope of the oil supply curve. A vertical short-run supply curve, e.g., implies an impact price elasticity of zero.

In light of the long lead time in oil field development, significant adjustment costs in crude oil production and the uncertainty about whether an observed change in demand represents a transitory shock or a permanent shift, it is widely accepted that the short-run supply curve is close to, if not effectively vertical (see also Hamilton, 2009a). By focusing on the change in oil production outside of Kuwait and Iraq in response to the Invasion of Kuwait in August 1990, Kilian and Murphy (2012) derive a plausible upper bound of 0.0258 for the short-run price elasticity of crude oil supply.

In terms of the structural model in (1), this translates into an upper bound on the relative impact response of world production and the real price of crude oil to a typical flow demand and speculative demand shock, respectively, i.e.

$$0 \leq \frac{a_{12}}{a_{32}} \leq 0.0258 \quad \text{and} \quad 0 \leq \frac{a_{13}}{a_{33}} \leq 0.0258. \quad (\text{B.4})$$

Note that the constraints in (4) do not restrict the impact response *levels* of the corresponding variables.²¹

²¹Kilian and Murphy (2013) emphasize that the 68% quantiles of the posterior distribution are qual-

Appendix B.2.3 Bounding the Impact Price Elasticity of Oil Demand

Conventional estimates of the price elasticity of oil demand equate crude oil production with crude oil consumption. In the present model, this corresponds to the relative impact response of oil production and the real price of oil to an exogenous flow supply shock. Accordingly, Kilian and Murphy (2013) refer to this as the “oil demand elasticity in production”.

The fact that crude oil is storable, however, implies that the identity between production and consumption breaks down. An appropriate measure of the impact response of oil consumption to an exogenous flow supply shock must take into account the shortfall in production as well as the change in crude oil inventories within the same month. For this reason, Kilian and Murphy (2013) propose the use of an “oil demand elasticity in use”, which is formally derived from the structural VAR model.

Empirical estimates of the long-run price elasticity of oil demand can be used as a plausible upper bound for the corresponding impact elasticity. Hausman and Newey (1995) and Yatchew and No (2001) provide estimates of the long-run price elasticity of *gasoline demand* based on U.S. and Canadian household survey data, respectively, arriving at values of -0.8 and -0.9 . Accordingly, admissible structural models must satisfy

$$-0.8 \leq \eta_t^{use} \leq 0, \quad (\text{B.5})$$

where $\eta_t^{use} \equiv \frac{\% \Delta use_t}{\% \Delta rpot}$ and $\% \Delta use_t = \frac{\Delta prod_t - \Delta(\Delta inv_t)}{prod_{t-1} - \Delta inv_{t-1}}$ (compare Kilian and Murphy, 2013, Appendix A). Note that Δ denotes level changes, while $\% \Delta$ denotes percent changes in response to an exogenous flow supply shock.

This oil demand elasticity in use can be rewritten in terms of the contemporaneous coefficients of the structural model in (1):

$$\eta_t^{use} \equiv \frac{(\text{prod}_{t-1} \cdot a_{11} / 100) - a_{41}}{\text{prod}_{t-1} - \Delta inv} \frac{a_{31} / 100}{a_{31} / 100} \quad (\text{B.6})$$

Since η_t^{use} is a function of $prod_{t-1}$, it is inherently time varying. I therefore follow Kilian and Murphy (2013), who use the “average oil demand elasticity in use”, $\bar{\eta}^{use}$, instead.

itatively very similar to the baseline model even for an unrealistically high elasticity of 0.1.

Appendix B.3 Empirical Results

The empirical implementation of the identifying assumptions is identical with Kilian and Murphy (2013).²² It is well-known that the identification by sign restrictions yields a set of admissible structural models rather than a unique estimate. While the additional constraints imposed on the impact price elasticity of supply and demand help narrow down this set, the matrix \mathbf{A}_0^{-1} and thus both demand shocks I am interested in will not be exactly identified, in general.

All results presented below are for the structural model with an “average price elasticity of oil demand in use” over the sample period closest to the posterior median of this elasticity obtained in Kilian and Murphy (2013). Moreover, I conducted the same analysis for all candidate models satisfying the identifying restrictions conditional on the least squares estimate of the reduced-form in (2).

My first interesting finding concerns the magnitude of the set of admissible structural models. Based on 5 million random draws for the matrix of contemporaneous coefficients, A_0^{-1} , only 16 models satisfy all identifying assumptions for the original sample period in Kilian and Murphy (2013), i.e. for 1973.2 through 2009.8. When extending the sample period to 2010.12 and 2011.12, respectively, this number rises to 58 and 61, although the set of admissible structural models satisfying *only* the impact sign restrictions expands by a much smaller fraction.

As a consequence, there is evidence that, adding more recent observations to the sample, a larger number of candidate models satisfies the upper bounds imposed on the short-run price elasticities of oil supply and oil demand (in use). This is consistent with the finding in Baumeister and Peersman (2013a) that both these elasticities have decreased considerably since the second half of the 1980s.²³

Appendix B.3.1 Impulse Responses in the Global Oil Market

In order to obtain an estimate of the price elasticity of oil supply, we require an estimate of the change in the real price of oil associated with an exogenous shift of the demand curve along the supply curve. Following Kilian (2009), an increase in crude oil demand

²²The reader is referred to Section 3.5 of Kilian and Murphy (2013) for technical details.

²³Note that, using the terminology of Kilian and Murphy (2013), the short-run price elasticity of oil demand in Baumeister and Peersman (2013a) corresponds to an oil demand elasticity “in production” rather than an oil demand elasticity “in use”.

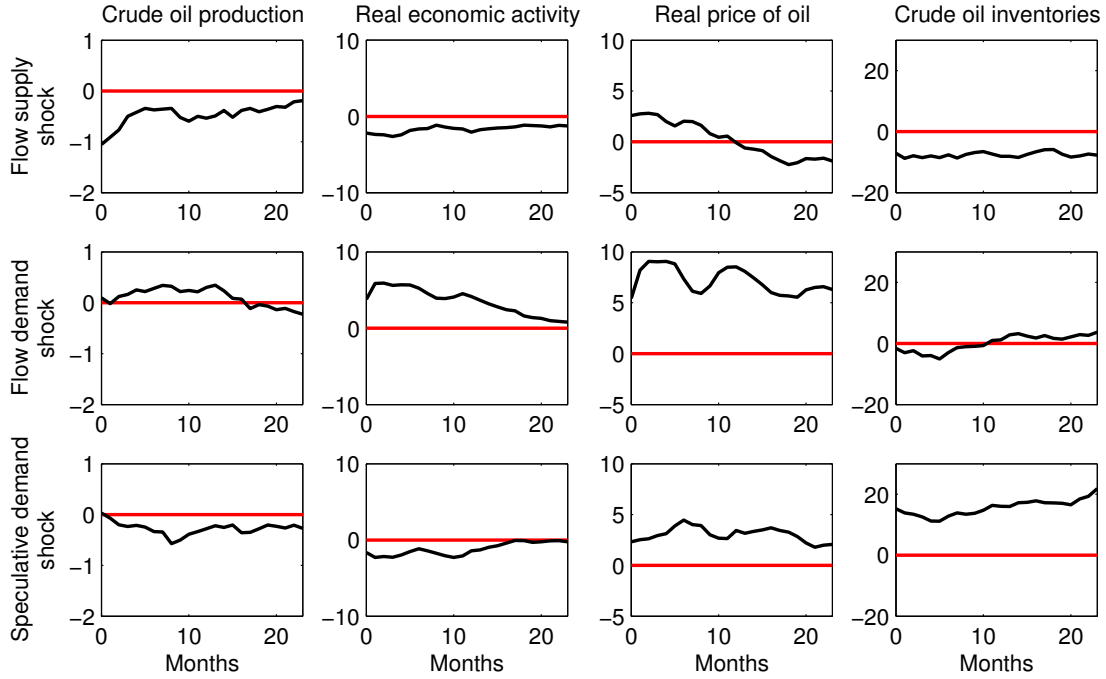


Figure B.1: Impulse responses to orthogonal oil supply and demand shocks based on the VAR model of the global market for crude oil in Kilian and Murphy (2013) (point estimates)

Notes: Results based on candidate model with a *price elasticity of crude oil demand in use* closest to the posterior median in Kilian and Murphy (2013). Crude oil production and inventories corresponds to cumulative impulse responses of $\Delta prod_t$ and Δinv_t .

which is exogenous from the point of view of oil-producing countries can arise from both higher demand *for consumption* or higher demand *for inventories*. Accordingly, the real price of oil might respond differently depending on the underlying structural cause.

Figure B.1 plots the impulse responses to supply and demand shocks in the global market for crude oil based on one of the 61 admissible structural models satisfying all identifying restrictions in Kilian and Murphy (2013). Impulse response functions for the entire set of admissible structural models are plotted in Figure B.2, for comparison. Note that the response of the real price of oil is defined in terms of percent deviations from its sample mean, while the cumulated responses are shown for $\Delta prod_t$ and Δinv_t . For the sake of brevity, we focus here on the dynamic behavior of the real price of oil in response to each of the orthogonal oil demand shocks.

Consider first a positive flow demand shock, which is associated with a higher demand for industrial commodities, in general, due to an unexpected increase in global real economic activity. On average over the sample period, the real price of oil increases by about 5




Figure B.2: Impulse responses to orthogonal oil supply and demand shocks based on the VAR model of the global market for crude oil in Kilian and Murphy (2013) (point estimates)

Notes: Results based on *all* candidate models satisfying the identifying assumptions. Crude oil production and inventories corresponds to cumulative impulse responses of $\Delta prod_t$ and Δinv_t .

percent on impact and remains elevated for at least two years. Moreover, a typical flow demand shock leads to an initial reduction in crude oil inventories which is reversed after approximately twelve months.

A typical speculative demand shock raises the real price of oil by about 2.5 percent on impact and by about 4.5 percent after six months before it starts to decline. Recall that, by construction, positive speculative demand shocks imply an increase in crude oil inventories in the impact period due to the imposed sign restrictions.