

The Effect of Oil Supply Shocks on U.S. Economic Activity: What Have We Learned?*

Ana María Herrera[†]

Sandeep Kumar Rangaraju[‡]

Abstract

Estimated responses of real oil prices and US GDP to oil supply disruptions vary widely. We show that most variation is attributable to differences in identification assumptions and in the model specification. Models that allow for a large short-run price elasticity of oil supply imply a larger response of oil prices and a larger, longer-lived contraction in U.S. real GDP. We find that if we condition on a range of supply elasticity values supported by microeconomic estimates, the differences in the oil price responses diminishes. We also examine the role of lag length, of using pre-1973 data, alternative measures of real economic activity and using the median response function instead of the modal structural model.

Key words: oil prices, oil supply shocks, economic activity, vector autoregressions, inventories, identification, Bayesian inference, sign restrictions.

JEL codes: C32, E32, Q41, Q43

*This material is partially based upon work conducted for the *North American Energy Harmonization Workshops* lead by Resources for the Future and supported by the National Science Foundation under Grant 1559339. We are grateful to Christiane Baumeister, Martin Bodenstein, Jordi Galí, James Hamilton, Lutz Kilian, Alan Krupnik, David Montgomery, Barbara Rossi, James Stock and participants at the Resources for the Future meetings on "New Approaches to Estimating the Macroeconomic Implications of Oil Price Shocks" for very helpful suggestions.

[†]Corresponding author: Department of Economics, Gatton College of Business and Economics, University of Kentucky, Lexington, KY, 40506-0034; e-mail: amherrera@uky.edu; phone: (859)257-1119.

[‡]Department of Economics, Goddard School of Business & Economics, 1337 Edvalson St., Dept 3807, Weber State University, Ogden, UT 84408; email: srangaraju@weber.edu; phone: (801) 626-6094.

1 Introduction

The use of structural vector autoregressions to estimate the effect of oil supply and demand shocks on economic activity has become common since Kilian (2008). However, whereas there appears to be an agreement regarding the usefulness of such econometric tool, there is considerable disagreement concerning the size of the effect of oil supply shocks on real oil prices and economic activity. For instance, among the studies reviewed in this paper, estimates of the effect on the real oil price of an unexpected decline in world oil supply –scaled to imply a 1% reduction in monthly world oil production– vary from 0.05% to 4.27% with the peak occurring at four and two months after the shock, respectively. In turn, estimates of the cumulative four-quarter loss in U.S. real GDP a year after the shock range from -0.14% to -0.51%.

Consider a researcher who is interested in estimating the effect of oil supply shocks on real oil prices or a policy maker who wants to assess the impact of oil supply disruptions on economic activity. How should he or she model the global market for crude oil? How would different modeling choices affect the estimated response of real oil prices and U.S. real GDP and inform policy decisions? Can recent micro-level estimates of the oil supply elasticity inform the researcher’s prior on the short-run elasticity of supply and, if so, how would relying on different estimates of this elasticity affect his or her inference regarding the effect of oil supply disruptions on the U.S. real GDP?

To answer these questions, we start by estimating four alternative structural vector autoregressive (SVAR) models for the global crude oil market that have shaped the recent literature¹ and have been used by practitioners to inform policy (i.e., Kilian 2009; Kilian and Murphy 2012, 2014; Baumeister and Hamilton 2018). We then evaluate how alternative modelling strategies (e.g., identification assumptions, period spanned by the data, lag length) map into different conclusions regarding the effect of oil supply shocks on real oil prices and U.S. real GDP. We complement our analysis by re-evaluating two additional VAR specifications (Lippi and Nobili 2012; Baumeister and Peersman 2013b) aimed at tackling two additional questions: (a) What do we learn from separating shocks stemming from the U.S. from the rest-of-the-world? (b) What does a time-varying parameter VAR tell us about the responsiveness of U.S. economic activity to oil supply disturbances?

The reader may ask, why focus on the role of oil supply shocks? First, in the last decade

¹When this paper was written, the Google Scholar citation count was 2270 for Kilian (2009), 291 for Kilian and Murphy (2012), 759 for Kilian and Murphy (2014), 314 for Baumeister and Peersman (2013b), 217 for Lippi and Nobili (2012) and 62 for Baumeister and Hamilton (2018).

a consensus had emerged regarding the smaller role of supply relative to demand driven shocks (Baumeister and Kilian, 2016). Nevertheless, recent work by Baumeister and Hamilton (2018) has prompted academics and policy makers to revisit the way prior information is used to conduct inference in structural vector autoregressive models of the oil market. Indeed, they obtain estimates of the effect of oil supply shocks that are considerably larger than suggested by earlier studies (e.g., Kilian 2009, Kilian and Murphy 2012, 2014) and larger than any microeconomic estimates of that elasticity (e.g., Anderson et al. 2018, Bjørnland et al. 2017, Newell and Prest 2018). Hence, our interest in analyzing how alternative modeling assumptions map into different conclusions regarding the importance of oil supply shocks. Second, the surge of U.S. shale oil production –and the corresponding reduction in U.S. net oil imports– led energy pundits and journalists to hypothesize that U.S. oil security had improved and, thus, that the macroeconomic costs of unanticipated oil supply shocks on U.S. economic activity had decreased. Yet, the International Energy Agency has recently warned that U.S. shale oil forecasts might have been too optimistic and that the world’s reliance on the U.S. shale oil potential is risky (Bloomberg, November 12, 2018).²

Our paper makes several contributions. First, we show that imposing a larger bound on –or allowing for a prior that places a greater probability on large values of– the short-run price elasticity of oil supply results in smaller estimates of the short-run price elasticity of demand in production and a larger response of real oil prices to oil supply disruptions.³ Second, we document how alternative identification assumptions map into significant differences in the size and persistence of the response of U.S. real GDP to oil supply disturbances. Third, we show that using larger dimension VARs –without imposing adequate restrictions– might lead to larger, yet imprecise, estimates of the cumulative loss in U.S. real GDP. Fourth, we examine the role of using pre-1973 data, longer lag lengths, using the median response function instead of the modal structural model, and using alternative measures of aggregate economic activity in accounting for the differences in the response of real oil prices and U.S. real GDP to oil supply shocks. Last but not least, we show that allowing for a prior that attaches considerable probability mass to larger values of the short-run elasticity of oil supply leads the researcher to infer a more elastic one-month oil supply, a less elastic short-run demand, a more responsive real oil price, and a larger one-year cumulative loss in U.S. real GDP. That is, if the researcher is willing to condition on a prior that has support on a range of supply elasticities that is backed by microeconomic evidence, the disagreements in the

²See also Kilian (2016, 2017).

³These results confirm the findings of Kilian and Murphy (2012) for their sign-identified model.

literature regarding the response of real oil prices are much smaller than they initially appear and closer to the original conclusions of Kilian (2009) and Kilian and Murphy (2012, 2014). Moreover, use of the same measure of economic activity diminishes the differences in the response of real oil price and U.S. real GDP across models.

This paper unfolds as follows. The next section describes the SVAR models used to analyze the role of oil supply and demand shocks on real oil prices and economic activity. Section 3 briefly describes the data. The following section discusses the effects of oil supply shocks real oil prices. Section 5 analyzes the response of U.S. real GDP. Section 6 explores the role of alternative priors on the short-run price elasticity of oil supply. The next section investigates the sensitivity of the results to ignoring the pre-1973 data, increasing the lag length, summarizing the estimation results under alternative loss functions, and using alternative measures of world economic activity. Section 8 summarizes and concludes.

2 SVAR models for the world crude oil market

This section revisits the SVAR specifications of the world oil market evaluated in this paper.

2.1 Kilian (2009)

Barsky and Kilian (2001, 2004) and Kilian (2008, 2009) first underscored the importance of separately identifying the effect of supply and demand driven shocks on the real price of oil. In the years preceding these studies it was common to estimate responses to oil price shocks that implicitly were a composite of demand and supply driven innovations. Work by Kilian (2009) suggested that this thought experiment needed to be reevaluated.

Kilian (2009) –hereafter *K09*– specified a model for the global oil market as follows. Let y_t represent a vector of observable monthly variables assumed to be governed by the structural vector autoregression:

$$B_0 y_t = \sum_{i=1}^{24} B_i y_{t-i} + \varepsilon_t \quad (1)$$

where $y_t = \begin{bmatrix} q_t & rea_t & p_t \end{bmatrix}'$, q_t is the percentage change in global oil production, rea_t is a suitable index of real global economic activity, p_t is the real price of oil, ε_t is a vector of serially and mutually uncorrelated structural innovations, B_0 is a 3×3 matrix of contemporaneous coefficients and the

B_i are 3×3 matrices of lagged coefficients.⁴

Identification is obtained by assuming that B_0^{-1} has a recursive structure. Crude oil supply is assumed not to respond contemporaneously (within a month) to innovations in the demand for oil. In addition, innovations in real oil prices that are driven by shocks to the oil market are assumed to have no contemporaneous effect on global economic activity. These restrictions imply a structural model of the form

$$u_t \equiv \begin{bmatrix} u_t^{\Delta prod} \\ u_t^{rea} \\ u_t^{rpo} \end{bmatrix} = \begin{bmatrix} b^{11} & 0 & 0 \\ b^{21} & b^{22} & 0 \\ b^{31} & b^{32} & b^{33} \end{bmatrix} \begin{bmatrix} \varepsilon_t^{oil\ supply\ shock} \\ \varepsilon_t^{aggregate\ demand\ shock} \\ \varepsilon_t^{oil-market-specific\ demand\ shock} \end{bmatrix} = B_0^{-1} \varepsilon_t \quad (2)$$

Two modeling choices are key for *K09's* identification scheme. First, having the oil supply shock first in the Wold causal chain implies that the short-run supply curve is vertical. This assumption is supported by (a) anecdotal evidence regarding the response of oil producers to demand-induced price shocks, (b) theoretical work by Anderson, Kellogg, and Salant (2018) who show that oil producers should respond by changing investment, but not current production, and (c) direct estimates of the price elasticity of supply close to zero. We will return to the importance of the short-run price elasticity of supply in section 6. Second, oil-market specific demand shocks are assumed to have no contemporaneous effect on real economic activity, which is consistent with Kilian and Zhou's (2018a) evidence that the real economic activity index is not driven by changes in the real price of oil.

2.2 Kilian and Murphy (2012)

Kilian and Murphy (2012) –hereafter *KM12*– use the tri-variate VAR setup in (1) but, instead of imposing exclusion restrictions as in *K09*, they attain identification via sign restrictions. They show that agnostic sign-restrictions on their own are not enough to infer the response of real oil prices to structural shocks. Instead, they find that sign restrictions combined with bounds on the short-run price elasticity of oil supply are required for identification. Their approach consists of the following steps. First, they impose sign restrictions on the impact responses (see Table A.1 in the online appendix). These imply that: (a) in response to an unanticipated oil supply decrease, oil production falls, the real oil price increases, and global economic activity falls; (b) in response to an unanticipated aggregate demand increase, global oil production, global real economic activity

⁴We exclude the constant for simplicity.

and the real oil price increase, and (c) in response to an unanticipated increase in oil-specific demand, global real activity falls whereas oil production and prices increase. Second, they impose an upper bound of 0.0258 on the short-run oil supply elasticity.⁵ This restriction is motivated by the presence of adjustment costs, which would prevent crude oil production from responding to unexpected demand shifts within the month.

Our identification strategy departs slightly from *KM12* and follows instead Inoue and Kilian (2013) in imposing additional dynamic sign restrictions on the response of the real oil price to the structural shocks. Namely, the response of the real oil price to oil supply disruptions, positive aggregate demand and oil-specific demand shocks is assumed to be non-negative for twelve months.

2.3 Kilian and Murphy (2014)

Kilian and Murphy (2014) –hereafter *KM14*– refined the framework of their earlier study by adding changes in global crude oil inventories to the VAR model and by modifying the identifying restrictions accordingly. These modifications allowed them to identify the effect of storage demand shocks driven by oil price expectations and to quantify the short-run price elasticity of demand, which was not explicitly defined in the previous models. Their identification strategy is summarized as follows. First, they impose the following sign restrictions on the impact multiplier (see Table A.1 in the online appendix): (a) An unanticipated flow supply shock causes oil production to decline, the real price of oil to increase, and global real economic activity to fall; (b) An unanticipated increase in flow demand causes oil production, global real economic activity and the real oil price to increase on impact; (c) A positive speculative (or storage) demand shock raises world oil production, the real price of oil, and crude oil inventories; yet it leads to a decline in real economic activity. Second, they impose an upper bound on the impact price elasticity of oil supply (0.0258 in the baseline model). Third, *KM14* limit the impact elasticity of oil demand in use to lie between -0.8 and 0. This amounts to imposing the restriction that the short-run price elasticity of demand cannot exceed the long-run elasticity as estimated by Hausman and Newey (1995) and Yatchew and No (2001). Fourth, they limit the response of the real oil price to a negative oil supply shock to be non-negative for 12 months and the response of oil production and global real economic activity to the same shock to be non-positive for 12 months.

⁵As mentioned in the previous section, recent estimates of the oil supply elasticity vary from numbers close to zero (Anderson, Kellogg and Salant, 2018) to numbers closer to 0.04. In that sense, Kilian and Murphy’s (2012) bound would seem conservative relative to the first studies but too small for the latter. While we return to this issue later in the paper, it is worth noting here that Kilian and Murphy (2012) show that their results are robust to imposing a less restrictive bound of 0.09.

2.4 Baumeister and Hamilton (2018)

The baseline model of Baumeister and Hamilton (2018) –hereafter *BH18*– is given by:

$$B_0 y_t = \sum_{i=1}^{12} B_i y_{t-i} + \varepsilon_t, \quad (3)$$

where

$$y_t = \begin{bmatrix} q_t & z_t & p_t & \Delta i_t \end{bmatrix}', B_0 = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 \\ 0 & 1 & -\alpha_{yp} & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -\chi^{-1} \\ -\psi_1 & -\psi_2 & -\psi_3 & 1 \end{bmatrix}, \varepsilon_t = \begin{bmatrix} \varepsilon_{1t}^* \\ \varepsilon_{2t}^* \\ \varepsilon_{3t}^* - \chi^{-1} e_t \\ \chi \varepsilon_{4t}^* + e_t \end{bmatrix},$$

$$\Delta i_t = \chi \Delta i_t^* + e_t \quad (4)$$

z_t is the world industrial production index, Δi_t represents a measure of the change in OECD crude-oil inventories as a percentage of the previous month's world oil production, $\chi < 1$, and e_t is a classical additive Gaussian measurement error. Note that the number of lags included in this specification is considerably smaller than in previous models (i.e., 12 instead of 24). *BH18* note that the system in (4) needs to be modified because the elements in ε_t are contemporaneously correlated due to the assumption of the measurement error, e_t . Uncorrelated structural shocks are obtained by premultiplying the system in (4) by

$$\Gamma = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \rho & 1 \end{bmatrix} \text{ where } \rho = \frac{\chi^{-1} \sigma_e^2}{d_{33}^* + \chi^{-2} \sigma_e^2} \text{ and } D = \begin{bmatrix} d_{11}^* & 0 & 0 & 0 \\ 0 & d_{22}^* & 0 & 0 \\ 0 & 0 & d_{33}^* + \chi^{-2} \sigma_e^2 & -\chi^{-1} \sigma_e^2 \\ 0 & 0 & -\chi^{-1} \sigma_e^2 & \chi^2 d_{44}^* + \sigma_e^2 \end{bmatrix}.$$

To attain identification, explicit Bayesian priors are imposed on parameters of B_0 . The researcher's information and/or beliefs about the contemporaneous coefficients are summarized in terms of a prior distribution on the elements of B_0 , $p[B_0]$. Table A.2 in the online appendix summarizes the priors for the main variables of interest (e.g., short-run price elasticity of supply, short-run price elasticity of demand in production). We refer the reader to *BH18* for a complete description of restrictions imposed on the model, which include –but are not limited to– setting $\psi_2 = 0$, and defining prior distributions for ρ , χ , $h_2 = \frac{\det(B_0^{-1}) - \alpha_{yp} \beta_{qy}}{\det(B_0^{-1})}$ and $\det(B_0)$. Note that *BH18*, unlike earlier studies, construct some of their priors based on estimating their model on pre-1973 data. While we follow their strategy and specify the priors in the same manner, it is worth

noting that using pre-1973 data to estimate the effect of oil supply shocks on the U.S. economy could be problematic as prices were controlled by the Texas Railroad Commission until the early 1970s and no integrated global oil market existed.

We conclude this section by highlighting important differences in the identification schemes described above. *K09*, *KM12*, and *KM14* impose exclusion or inequality restrictions on the elements of B_0^{-1} . In contrast, *BH18* impose Bayesian priors and zero restrictions on the elements of B_0 . Therefore, the strategy used by *BH18* does not impose the same identification assumptions as *KM12* or *KM14*; more specifically, it does not impose any cross-equation restrictions or dynamic sign restrictions (see Kilian and Lütkepohl (2017) and Kilian and Zhou (2018) for an in-depth discussion). However, the priors used by *BH18* imply prior distributions and probable signs⁶ for the elements of $B_0^{-1} = \frac{1}{\det(B_0)} \times \text{adjugate}(B_0) = \frac{1}{\det(B_0)} \times C$ where $\det(B_0) = \chi^{-1}(-\psi_3 + \chi\alpha_{pq} - \chi\beta_{pq} - \psi_1\alpha_{pq} - \chi\alpha_{yp}\beta_{qy})$ and

$$C = \begin{bmatrix} -\frac{1}{\chi}(\psi_3 + \chi\beta_{qp} + \chi\alpha_{yp}\beta_{qy}) & \alpha_{qp}\beta_{qy} & \alpha_{qp} & \frac{1}{\chi}\alpha_{qp} \\ -\frac{1}{\chi}(\chi\alpha_{yp} - \psi_1\alpha_{yp}) & -\frac{1}{\chi}(\psi_3 + \psi_1\alpha_{qp}) + (\beta_{pq} - \alpha_{qp}) & \alpha_{yp} & \frac{1}{\chi}\alpha_{yp} \\ -\frac{1}{\chi}(\chi - \psi_1) & \beta_{qy} & 1 & \frac{1}{\chi} \\ -\psi_3 - \psi_1\beta_{pq} - \psi_1\alpha_{yp}\beta_{qy} & \psi_3\beta_{qy} + \psi_1\alpha_{qp}\beta_{qy} & \psi_3 + \psi_1\alpha_{qp} & \alpha_{qp} - \beta_{pq} - \alpha_{yp}\beta_{qy} \end{bmatrix}.$$

In addition, a key difference between *BH18* and previous studies is that they specify a prior for the price elasticity of oil supply that assigns a 94% probability mass to values that exceed the 0.0258 of *KM12* and *KM14* or the 0.04 bound considered by Zhou (2019) when replicating *KM14*'s findings.

2.5 Lippi and Nobili (2012)

Lippi and Nobili's (2012) –hereafter *LN12*– identification strategy is similar to *KM12* in the use of sign restrictions. However, they deviate in three important ways. First, their SVAR includes –in addition to real oil prices and oil production– real prices in the U.S. economy, the U.S. industrial production index, and a measure of the business cycle in the rest of the world (RoW output). Second, in contrast with *KM12* and *KM14*, no bounds or dynamic sign restrictions are imposed (see Table A.1 in the online appendix for impact sign restrictions). Third, they build a theoretical three-country model derived from Backus and Crucini (2000), which allows them to map five fundamental shocks to the observed responses of relative prices and production.

⁶See Table 2 of Baumeister and Hamilton (2018) for prior and posterior probabilities that the impact of a particular structural shock on the VAR variables is positive.

2.6 Baumeister and Peersman (2013)

Baumeister and Peersman (2013b) –hereafter *BP13*– were the first to use a time-varying coefficients SVAR. Three modeling choices are key in *BP13*'s. First, quarterly, instead of monthly, data is used. Second, they employ the rate of growth of U.S. real GDP as a measure of economic activity, real oil prices are expressed in rates of growth and U.S. CPI inflation is included in the SVAR. This model allows for direct estimation of the effect of oil supply shocks on U.S. real GDP. However, it does not capture movements in demand stemming from the rest-of-the-world that are not reflected in the latter. Third, the identification of oil supply shocks is "based on the sign restrictions that (these shocks move) oil prices and oil production in the opposite direction" for four quarters following the shock. No additional sign restrictions are imposed. This SVAR is partially identified as only the supply shock has a structural interpretation.⁷ For this reason, and to economize space, we have opted to report the results in the online appendix.⁸

3 Data

Although the data for the above-described studies are available from the data repository or from the authors, we reconstruct the series using the original sources and apply the same transformations as in the original studies. To facilitate the comparison across alternative specifications, the monthly data used in all SVARs cover the same sample period of January 1973 to December 2016.

World oil production is obtained from the Monthly Energy Review published by the Energy Information Agency (EIA) and is measured in thousands of barrels per day. Kilian's (2009) index of global economic activity is computed as the cumulative average rate of increase in bulk dry cargo ocean freight rates, deflated by the U.S. CPI and linearly detrended.⁹ The log of the real oil price is measured as the log difference between the refiners' acquisition cost (RAC) of imported crude oil and the U.S. CPI. The refiners acquisition cost is provided by the EIA starting in January of 1974. We extrapolate the data from January 1974 until January 1973 following Barsky and Kilian (2002). Data for the CPI are obtained from the St. Louis Fed FRED database. In all SVAR

⁷In a separate paper, Baumeister and Peersman (2013a) identify both supply and demand driven shocks in a model for the world oil market. Given our interest in supply disruptions, and for the sake of brevity, we will not re-visit the latter. However, we note here that the results for both models are qualitatively similar.

⁸See <http://gattonweb.uky.edu/faculty/herrera/documents/HRoilAppendix.pdf>.

⁹The updated time series for this index is available from Lutz Kilian's website at <http://www-personal.umich.edu/~lkilian/reaupdate.txt>. See Kilian (2009) and Kilian and Zhou (2018a) for a detailed description of the methodology employed to compute the index and a comparison with alternative indicators of global real economic activity.

models the world oil production enters as a log growth rate. RoW output for *LN12* is calculated as global exports to the world net of global imports from the U.S. and from oil exporting countries. These data are obtained from the IMF. RoW output is expressed in real terms and measured in logarithms.

Because data for inventories are not available for countries other than the US, we follow *KM14* in using data from the EIA for total US crude oil inventories, as well as US and OECD petroleum stocks to compute a proxy. The inventory proxy is calculated by multiplying US crude oil inventories by the ratio of OECD to US petroleum stocks. This ratio ranges from 2.23 to 2.61 in our sample, which is almost identical to the 2.23 to 2.59 range in the original *KM14* sample (1973:1-2009:8). Although *KM14* and *BH18* use the same proxy for the level of world crude oil inventories, they use a different transformation of the series in their SVAR. We preserve the transformations used in each paper and, thus, include the change in the inventory proxy for *KM14* model and the change in world inventories as a fraction of last month’s world oil production for the *BH18* model. In addition, *BH18* measure global economic activity as the log change in the industrial production index for the OECD and the six major non-member economies (Brazil, China, India, Indonesia, the Russian Federation and South Africa) instead of using Kilian’s index of real economic activity.¹⁰

4 The Effect of Oil Supply Shocks on Real Oil Prices

This section describes the effects of oil supply shocks on real oil prices implied by alternative SVAR specifications.

4.1 Preliminaries

The procedures for estimating the impulse response functions from VAR models with a recursive structure are straightforward. In contrast, inference in sign-identified models warrants some discussion. Until recently, the convention in summarizing the results of sign-identified models was to report the vector of pointwise posterior medians of the impulse responses in conjunction with the pointwise $100(1 - \alpha)\%$ posterior error bands. However, there are two concerns with that practice. First, as noted by Fry and Pagan (2011), *KM12* and Inoue and Kilian (2013), these point estimates do not correspond to one particular admissible model. Hence the median response has no

¹⁰Note that data on industrial production is not available for the whole sample period for the non-member countries. Data become available in January 1975 for Brazil, April 1994 for India, January 1986 for Indonesia, January 1993 for Russia, January 1999 for China, and January 1990 for South Africa.

structural interpretation. Second, the vector of pointwise medians is not an appropriate measure of central tendency for the impulse response functions (see Inoue and Kilian 2013). To address these concerns, Inoue and Kilian (2013) suggest reporting the impulse responses for the modal model among the admissible models. The modal model is defined as the admissible model that maximizes the joint posterior density of the admissible SVAR models. In addition, Inoue and Kilian (2013) propose conducting joint inference by reporting the $100(1 - \alpha)\%$ highest posterior density (HPD) credible sets for the admissible models. These joint credible sets are constructed by ranking the admissible models based on the value of the joint density and taking the set contained in the first $(1 - \alpha)Q$ sorted pairs of response functions and joint densities, where Q denotes the number of admissible models among the total draws used to compute the response functions.

Hence, we deviate from *KM12* who report the impulse response functions for all the admissible models and from *KM14* who depict the model with an impact price elasticity in oil demand use that is closest to the posterior median of the elasticity among admissible models. Instead, we follow Inoue and Kilian (2013) in reporting the responses for the modal model as well as the 95% pointwise HPD credible sets.¹¹ For ease of comparison, we report the posterior median for all models but *K09*, noting that these summary statistics have to be interpreted with caution for the reasons outlined above.

BH18 report the pointwise posterior median response function and the pointwise 95% posterior credible sets.¹² Their methodology does not allow for the construction of joint credible sets as proposed by Inoue and Kilian (2013). Thus, for the Bayesian VAR of *BH18* we follow their estimation procedure and report the same statistics.

4.2 Estimation Results: From exclusion restrictions to Bayesian analysis

Figure 1a reports the responses of oil production and real oil prices to an oil supply disruption normalized to represent the effect of an unexpected 1% decline in monthly world oil production. The top panels of Figure 1a depict the responses for *K09*'s specification and the one and two standard deviation error bands computed via a recursive-design wild bootstrap. World oil production remains below its trend for more than twelve months after the shock, but the increase in the real oil price

¹¹We refer the reader to Inoue and Kilian (2013, 2018) for a step-by-step description of the computation procedure. We use $M = 5000$ draws from the joint posterior distribution of the reduced-form VAR parameters with $N = 20,000$ rotations each.

¹²Baumeister and Hamilton (2018) justify reporting the median response function because it corresponds to the solution to a loss function involving the sum of the absolute loss of the vector of structural responses. As noted in Kilian and Lütkepohl (2017), this loss function assumes the user does not care about dependence of the impulse response across horizons and variables and, hence, is not likely to reflect the preferences of typical VAR users.

is small and short-lived. For instance, six months after the shock the real oil price has increased by only 0.05%. These results are very similar to the original estimates obtained by Kilian (2009) using data spanning the period between January 1973 and December 2007.

Figure 1a illustrates the responses of oil production and the real price of oil obtained from *KM12*'s sign-identified SVAR. The black and blue lines represent the pointwise posterior median and the structural modal model, respectively. The red area corresponds to the 95% HPD credible set. The response of oil production is persistently negative and very similar to that implied by *K09*. As for the real oil price, the response function for the modal model reveals a significant increase for the first twelve months. The increase in the real oil price is considerably larger than in *K09*. For instance, six months after the shock the real oil price is estimated to rise 1.4% instead of the less than 0.05% increase in *K09*. Although imprecisely estimated, the response in the modal model peaks at 2.1% sixteen months after the shock.

Figure 1a displays the responses for *KM14*. Modeling storage demand leads to a somewhat less persistent response of oil production but larger increase in real oil prices than in *KM12*. The real price of oil rises for two months, reaches a peak increase of 2.38% and then starts to decline. Eight months after the shock, the real oil price has returned to its initial level.

The responses for *BH18* are plotted in Figure 1a. The black line represents the pointwise posterior median and the red area depicts the 95% posterior credible sets. The dynamic response of oil production is similar to the previous specifications. However, the response of the real oil price is quite distinct. For instance, an unexpected 1% decline in world oil production raises the real price of oil by 4.27% at the peak (2 months after the shock). Contrast this with the 2.38% increase (two months after the shock) in *KM14*'s modal model, the alternative specification with the largest effect and the only other model that includes inventories.

How do differences in the response of the real oil price relate to the underlying SVAR specification? To answer this question Panel A of Table 1 reports the short-run (impact) elasticities of supply and demand implied by the SVARs.¹³ Recall that the short-run elasticity of oil supply is assumed to be zero in *K09*, and bounded above by 0.0258 in *KM12* and *KM14*. The estimated supply elasticity in the modal model for *KM12* and *KM14* is not very different from zero (0.019 and 0.017, respectively) and very close to the state-of-the-art microeconomic estimate of the U.S. oil supply elasticity in Newell and Prest (2018). In contrast, for *BH18* where a Student- t prior distribution is assumed for α_{qp} , which puts substantial probability weight on larger elasticity

¹³For reference, estimates obtained using the original samples are reported in the online appendix.

values than permitted in *KM12* and *KM14*, the posterior median (0.144) is considerably larger than current microeconomic elasticity estimates.

The traditional approach to computing the short-run price elasticity of oil demand has been to divide the impact response of oil production to an oil supply shock by the impact response of the real oil price to the same shock. As Table 1 shows, the smaller the price elasticity of oil supply, the larger the corresponding traditional price elasticity of oil demand. Under the recursive identification scheme of *K09*, we obtain an estimate of -3.131, and of -1.723 in the sign-identified model of *KM12*. However, because storage demand is not modeled in *K09*, or *KM12*, the traditional measure of the short-run price elasticity of oil demand is invalid (see *KM14*). In fact, the price elasticity of oil demand is not well-defined in models excluding oil inventories (For an in-depth discussion see Kilian and Zhou, 2018a).

KM14 emphasize the importance of distinguishing between the price elasticity of oil demand in production and in use in a world where crude oil inventories change over time. The latter measures the change in the quantity of oil brought about by an oil supply shock as the change in the flow of oil production plus the change in crude oil inventories. *KM14* show that traditional estimates of the price elasticity of oil demand can be quite misleading. Indeed, the short-run demand elasticity in production for the modal model in *KM14* equals -0.546, whereas the corresponding demand elasticity in use is -0.260. In addition, extending the sample used by *KM14* leaves the estimate of the demand elasticity in use virtually unchanged (see Table A.3 of the online appendix for estimates obtained with the original sample).

While *BH18's* setup does not allow for the computation of the oil demand elasticity in use (see Kilian and Zhou, 2018b) –the relevant elasticity when accounting for storage demand– we are able to compare the demand elasticity in production. Note that the posterior median for the demand elasticity in production in *BH18* equals -0.356 and the 95% posterior credible set covers (-0.803 -0.152). The posterior median for *BH18* is considerably smaller than the demand elasticity in production estimated for *K09* and *KM12*.

These results confirm that modeling storage demand is key for pinning down the relevant elasticity of demand and is linked to a larger and more persistent response of oil prices to supply disruptions. Hence, practitioners interested in estimating the effect of oil supply (and demand) shocks on the real oil price should give serious consideration to the role of storage demand in the world oil market as indicated by *KM14*.

Finally, differences in the impulse response functions can be traced to differences in the estimated

impact multiplier, B_0^{-1} . Panel C of Table 1 reports the estimates of B_0^{-1} for the structural modal model in *KM14* and the posterior median for *BH18*, the two models with the largest divergence in response functions and similar definitions of y_t .¹⁴ Three discrepancies emerge. First, whereas *KM14* specification offers little evidence against the null of a vertical oil supply curve –all but the first element of the first row are very close to zero–, *BH18*’s specification would not support an inelastic supply. Second, the first column indicates very different impact responses to oil supply disturbances. *KM14* model implies a large impact on real economic activity and a muted impact response of real oil prices –see (2,1) and (3,1) elements of B_0^{-1} , respectively– whereas *BH18* entails a small impact of world IP and a large initial response of real oil prices. Third, a large discrepancy is evident in the impact of oil supply disturbances on inventories in *KM14* and *BH18*. The estimate of 5.088 in *KM14*’s modal model indicates a large impact on inventories whether the median estimate of -0.022 in *BH18* is very close to zero. All in all, while both identification schemes imply the same posterior probabilities that the impact of a structural shock is positive (see Table A.6 in the online appendix), they do entail different estimates of B_0^{-1} and, thus, discrepancies in the impulse response functions.

4.3 Disentangling the effects of U.S. and rest-of-the-world shocks?

We now turn to *LN12* who seek to disentangle the effects of U.S. and rest-of-the world driven shocks. The Figure 1a and the fifth column of Table 1 show that the impact effect of an oil supply disturbance is larger for this specification than all other models. An unexpected 1% decline in world oil supply leads to an impact increase in monthly oil prices of 12.07%. This increase is persistent (a year after the shock the price is estimated to have risen 7.15%), but imprecisely estimated. The posterior median for the short-run price elasticity of supply is very large (0.447) while the estimated short-run elasticity of demand is close to zero, reflecting the absence of supply elasticity bounds, which were only introduced by *KM12* after *LN12* was written.

In brief, adding additional variables to the sign-identified VAR might shed some light into the role of U.S. versus rest-of-the-world demand shocks. Yet, without additional restrictions, the estimated elasticities fall way off the magnitudes found in the literature. Moreover, the number of unrestricted parameters that can be estimated with precision appears to be limited.

¹⁴Estimates of B_0^{-1} for the other models can be found in Table A.4 of the online appendix.

5 Oil Supply Shocks and Aggregate Economic Activity

This section evaluates how alternative SVAR models for the global oil market affect our inference regarding the effect of oil supply shocks on U.S. real GDP. The main hurdle in tackling this question is that data on GDP are only available at the quarterly frequency. The practitioner might thus consider constructing an analogous VAR for the world oil market at the quarterly frequency. Yet the use of quarterly data would invalidate identification schemes that impose restrictions on the impact multiplier matrix. Instead, we follow a two-step method similar to *K09* whereby we first extract a measure of structural oil supply shocks from the SVAR models, use these series to compute a quarterly measure of oil supply shocks and then project the real GDP growth on the shocks.

5.1 Empirical Strategy

The first step in estimating the effect of oil supply shocks on U.S. real GDP is to extract the monthly series of structural supply shocks from the SVAR models.¹⁵ For the recursive VAR of Kilian (2009), an estimate of the structural impact multiplier matrix, \hat{B}_0^{-1} , is obtained via Cholesky decomposition of the estimated variance-covariance matrix $\Omega = E[u_t u_t']$ where $u_t = B_0^{-1} \varepsilon_t$. The series of monthly structural supply shocks (ε_{1t}) corresponds to the first column of the $(3 \times T - p)$ matrix $B_0 u_t$. For *KM12* and *KM14* we retain the series of structural supply shocks for all admissible models. This allows us to compute the response of the modal model and the $(1 - \alpha)\%$ pointwise HPD credible sets. For *BH18* and *LN12* we retain the posterior distribution for the structural VAR model parameters and the corresponding posterior distribution of monthly structural supply shocks ε_{1t}^* in (4). This allows us to compute the pointwise posterior median as well as credible regions for the GDP response.

We construct quarterly measures of oil supply shocks by averaging the monthly structural innovations derived from the SVARs for each quarter (see Kilian 2009). Let $\tilde{\varepsilon}_{stj}^i$ denote the estimated supply shock in month j of the t^{th} quarter implied by each of the $i = K09, KM2012, KM2014, BH18, LN12$ SVAR specifications.¹⁶ Then, for each specification i , the quarterly measure of oil

¹⁵See online appendix for a detailed description of the methodology.

¹⁶The structural supply shocks series for *K09*, *KM12*, *KM14*, and *BH18* exhibit very similar patterns. In fact, the correlations range between 0.96 for *K09/KM12* and 0.77 for *K09/BH18*. In contrast, the correlation between *LN12* and the other models ranges from 0.26 for *K09* to 0.44 for *BH18*. Very similar correlations are obtained when the quarterly measures are used.

supply shocks is given by

$$\widehat{\xi}_{st}^i = \frac{1}{3} \sum_{j=1}^3 \widehat{\varepsilon}_{stj}^i, \quad i = K09, KM2012, KM2014, BH18, LN12. \quad (5)$$

Under the identifying assumption that the quarterly measure of oil supply shocks, $\widehat{\xi}_{st}^i$, does not respond to changes in U.S. GDP (Δy_t) within the quarter, we may treat $\widehat{\xi}_{st}^i$ as predetermined with respect to Δy_t . As Kilian (2009) notes the assumption that the series of quarterly shocks are predetermined with respect to U.S. real GDP growth is not testable. Nevertheless, a correlation with the innovations that is close to zero would suggest that unanticipated changes in crude oil supply are not associated with contemporaneous changes in real GDP growth. Examination of the empirical correlation between the autoregressive residuals of GDP growth and the different measures of oil supply shocks reveals that the correlations are quite small,¹⁷ thus leading credence to the assumption that the supply shock is predetermined with respect to US real GDP growth. It is important to note that $\widehat{\xi}_{st}^i$ does not denote the actual decrease in production in a quarter. Instead, it is an average of the j estimated monthly structural supply shocks in a quarter t .

The third step involves estimating the effect of oil supply shocks on U.S. real GDP, y_t , and real oil prices, op_t , via OLS according to the following equations

$$\Delta y_t = \alpha_{1,k} + \sum_{j=0}^{12} \beta_{1,j} \widehat{\xi}_{st-j}^i + v_{1,t}. \quad (6)$$

$$\Delta op_t = \alpha_{2,k} + \sum_{j=0}^{12} \beta_{2,j} \widehat{\xi}_{st-j}^i + v_{2,t}. \quad (7)$$

where Δy_t and Δop_t represent quarter-to-quarter rates of growth.¹⁸ Impulse response coefficients at horizon h are given by the β_{ij} 's and, therefore, the number of lags ($p = 12$) is determined by the maximum horizon of the impulse response function.¹⁹ Standard errors for the recursive VAR are computed using a block bootstrap (Kilian, 2009) to control for serial correlation in the error term.

To conduct inference on the impulse response functions for *KM12* and *KM14* we retain the vector of structural supply innovations for all admissible models. For each of these models, we

¹⁷The correlation coefficients for all the models are statistically insignificant.

¹⁸The use of quarter-to-quarter rates of growth differs from Kilian (2009) who uses an annualized rate of growth for real GDP.

¹⁹This estimation approach was originally applied by Kilian (2009). To ensure that lagged values of Δy_t have no explanatory power we include up to four lags Δy_t and find no individual or joint significance.

compute the quarterly measure of oil supply shocks and estimate the effect of the supply shocks on U.S. real GDP. This allows us to report the response for the modal model as well as 95% joint HPD credible sets that are consistent with the admissible SVAR models. For ease of comparison, we also report the pointwise posterior median.

For *BH18* and *LN12*, we retain the series of estimated monthly structural supply shocks for all the draws or admissible models, respectively. Then, for each replication we compute the quarterly measure as in (5) and estimate the OLS model in (6). We report the pointwise median response function as well as the 95% pointwise credible sets.

5.2 The Effect of Oil Supply Shocks on U.S. GDP

Figure 1b summarizes the cumulative responses of the real oil price and the U.S. real GDP to a one unit increase in $\tilde{\xi}_{st}^i$ for each time-invariant model. The top panels report the responses for the *K09* specification, as well as one and two-standard error bands. The second and third panels illustrate the posterior median response in black, the response for the modal model in blue, and the responses for the 95% HPD credible sets for *KM12* and *KM14*, respectively. For *BH18* and *LN12* we report the posterior median response and the 95% credible sets.

Figure 1b shows the response of real GDP to an unexpected decrease in world oil production is, as expected, negative for all specifications. Differences in the SVAR methodology lead to dissimilar conclusions regarding the response of real GDP to an oil supply shock. First, the magnitude and shape of the response functions is very similar for the two models that exclude inventories. However, whereas *K09* and *KM12* suggest a moderate impact on GDP, the response for *KM12* implies a long-lasting contraction. Second, modeling storage demand in the sign-identified setup of *KM14* would lead the researcher to downplay the importance of oil supply shocks.²⁰ The response of real GDP in *KM14* is considerably muted in the short-run relative to *KM12* and estimated with a lower degree of precision. Third, the median impact (-0.454) and one-year (-0.515) responses of U.S. GDP is largest for *LN12* (see Panel B in Table 1).

Nevertheless, the response of the real oil price to a one unit increase in $\tilde{\xi}_{st}^i$ also differs considerably across models (see Figure 1b). Hence, to evaluate the response of U.S. real GDP to a one percent increase in real oil prices, we calculate the dynamic GDP-oil price multiplier. This multiplier is computed as the ratio of the cumulative impulse response of these two variables to the

²⁰This result is consistent with the findings in Kilian and Murphy (2014).

oil supply shock $\widehat{\xi}_{st}^i$ at horizons h :²¹

$$\Phi_t^i(h) = \frac{\sum_{k=0}^h \frac{\partial \Delta y_{t+k}}{\partial \widehat{\xi}_{st}^i}}{\sum_{k=0}^h \frac{\partial \Delta op_{t+k}}{\partial \widehat{\xi}_{st}^i}} \quad (8)$$

Panel B of Table 1 reports the dynamic multiplier on impact and one year after the shock conditional on an oil supply shock, as well as the corresponding median (modal in bold) responses for real GDP.²² Three results stand out. First, as expected, the impact multipliers for all specifications are negative, which is consistent with oil supply shocks moving U.S. real GDP and real oil prices in opposite directions. Second, the impact multiplier is somewhat larger for *K09* than *KM14* or *BH18*. In *K09* a negative oil supply shock that drives the oil price down by one percentage point simultaneously reduces (quarterly) real GDP by about 0.1 percentage points whereas in *KM14* and *BH18* GDP declines 0.02 percentage points (0.04 in the modal model) and 0.01, respectively.²³ Last but not least, note that the one-year dynamic multiplier suggests that the effect of an oil supply shock that leads to a 10% increase in real oil prices is of similar magnitude across specifications: the estimated contraction of U.S. real GDP four quarters after the shock ranges between 0.1% and 0.6% for *KM14* and *LN12*, respectively.

We note here that estimates from a time-varying model (*BP13*) suggest the cumulative loss in U.S. GDP growth induced by a 1% decline in world oil production is considerably larger in the 2000s than in the earlier part of the sample (see top right panel of Figure A.2 in the online appendix). The largest losses occur in the early 2010s when real oil prices became more responsive and the short-run price elasticity of demand increased. Yet, the degree of precision is considerably lower in the later part of the sample. In fact, estimates of the four-quarter GDP loss for the TVP-BVAR model could be positive or negative for most of the sample.²⁴

Summarizing, we find that SVAR models that place a larger probability on values of the short-run (one-month) price elasticity of supply away from zero imply a somewhat greater response of U.S. real GDP to oil supply disturbances a year after the shock. However, once we account

²¹See Galí and Gambetti (2019) and Ramey and Zubairy (2018) for the computation of dynamic multipliers in alternative setups.

²²Note that the definition of the dynamic multiplier here differs from the multiplier in macroeconomic textbooks where all else is held equal. Here, except for *K09*, an unexpected change in oil production does not leave all variables unchanged.

²³Note that the responses for the modal model are very similar in magnitude to the pointwise median.

²⁴See Figure A.1 in the appendix for the median impact response of the real GDP associated 68% and 95% credible sets.

for differences in the responsiveness of real oil prices to the quarterly supply shocks, we obtain similar estimates of the dynamic GDP-oil price multiplier. Including more variables in the SVAR to disentangle U.S. versus ROW driven shocks results in larger but rather imprecise estimates of the GDP loss. Finally, while allowing for time-varying parameters suggest some variation in the response of GDP over time, the estimates lose precision in the last two decades.

We conclude this section by noting that the main insights derived in this section are robust to using an estimation method that takes into account the mixed-frequency nature of the structural shocks and U.S. real GDP (see Figure A.5 of the online appendix). If we use a series of mixed data sampling (MIDAS) regressions to estimate the GDP responses using local projections (see Jordá 2005 and Ferrara and Guerin 2018), the GDP responses for *K09* and *KM12* are very similar to those obtained without accounting the mixed-frequency nature of the data. The only noticeable differences are that the MIDAS regressions suggest a less persistent and smaller response of oil supply shocks for *KM14* and a more persistent and pronounced response for *BH18*.

6 The role of the prior on the short-run supply elasticity

Estimation results presented in the previous sections suggest that the researcher’s prior beliefs regarding the short-run price elasticity of supply are key to estimating the impact of oil supply shocks on real oil prices and U.S. GDP. To dig deeper into this question, we re-estimate the baseline model proposed by *BH18* imposing a modified prior on α_{qp} . We employ a Student $t(c_{qp}^\alpha, \sigma_{qp}^\alpha, \nu_{qp}^\alpha)$ with mode at $c_{qp}^\alpha = 0.1$, scale parameter $\sigma_{qp}^\alpha = 0.2$ and truncated to be positive as in *BH18*; however, we truncate the prior to impose a tighter upper bound. First, we truncate our prior at 0.0258 to impose the same upper bound as in *KM14*. This moves probability mass away from the right tail of the prior and towards the support $\alpha_{qp} \in (0, 0.0258]$. Note that in *BH18*’s benchmark specification the prior probability that the elasticity exceeds the 0.0258 bound imposed in *KM14* equals 94%.²⁵

Table 2 reports the posterior median and the 68% credible regions (in parenthesis) for the short-run price elasticities of oil supply, α_{qp} , and demand, β_{qp} for *BH18*’s benchmark and the doubly truncated Student- t prior. Imposing this upper bound on the support results in posterior median estimates of 0.019 and -1.386 for the supply and demand elasticities, respectively. That is, the posterior median for the short-run elasticity of supply is an order of magnitude smaller and

²⁵The reader should note that we only change the prior on α_{qp} and do not explore interactions between this prior and other priors (e.g., the priors on the lagged coefficients of the structural equation).

the short-run elasticity of demand in production is an order of magnitude larger than in *BH18*'s benchmark ($\alpha_{qp} = 0.144$, $\beta_{qp} = -0.356$), highlighting other important differences between the *KM14* and *BH18* specifications besides this bound.

Figure 2a illustrates the responses of oil production and real oil prices for the benchmark and the alternative prior. The only noticeable difference in the response of oil production is a narrower credible region for the impact response, consistent with the smaller support for the prior. In contrast, the increase in real oil prices is quite smaller when the prior is bounded above by 0.0258. The median response of the real oil price twelve months after the shock equals 1.421% instead of 3.685%.

KM12 and *KM14* motivate the bound of 0.0258 by appealing to extraneous information. On the one hand, they cite Kellogg(2011), who uses monthly well-level oil production data from Texas and finds that oil production does not respond to changes in spot or futures oil prices. On the other hand, they argue that –even in the presence of spare capacity due to the existence of adjustment costs–, oil production would barely respond to changes in oil prices within a month. Recent empirical studies would seem to suggest that the price elasticity for conventional oil production is indeed low. Anderson, Kellogg and Salant (2018) find that oil production from existing wells in Texas does not respond to changes in oil production. Using data for North Dakota, Bjørnland, Nordvik and Rohrer (2017) estimate an elasticity of 0.035 for conventional wells and 0.041 for shale oil wells.²⁶

Motivated by the estimates in Bjørnland, Nordvik and Rohrer (2017), we evaluate the sensitivity of the results a priori for α_{qp} where we impose an upper bound of 0.04 on the supply elasticity ($\alpha_{qp} \in (0, 0.04]$). This bound is justified as follows. Given that shale oil production amounted to 4% of world oil production in recent years, Bjørnland, Nordvik and Rohrer (2017) short-run price elasticity estimates (0.076 and 0.035 for shale and conventional wells, respectively, in the model with fixed effects) suggest a short-supply elasticity of $0.04 \times 0.076 + 0.96 \times 0.035 = 0.04$ for the world oil production. Moreover, if we assume that shale oil accounts for 10% of the world oil production –i.e., the Energy Information Agency's (2018) estimate of the percentage of world's crude oil technically recoverable from shale oil resources in the United States and 137 shale formations in 41 other countries–, then we obtain the same the world short-run supply elasticity ($0.9 \times 0.035 + 0.1 \times 0.072 = 0.04$). This bound is still an order of magnitude smaller than the benchmark estimate in *BH18*. In fact, for most of our sample the production of unconventional wells was negligible; hence, this

²⁶Note that a higher estimate of 0.072 is obtained for the short-run elasticity of shale oil wells when year and well fixed effects (FE) are excluded. However, neither estimate (with or without FE) is statistically different from zero.

bound of 0.04 is rather generous. Indeed, using drilling and production data from wells in Texas, Oklahoma, North Dakota, California and Colorado, Newell and Prest (2018) estimate the elasticity of oil supply for conventional and unconventional wells to equal -0.02 and 0.12, respectively. These estimates would suggest a world elasticity of supply of 0.01, which is even smaller than the 0.0258 bound of *KM12* and *KM14*.

In related work, Caldara, Cavallo and Iacoviello (2018) infer higher oil supply elasticities than *KM12* and *KM14*. They first identify episodes of large country-specific drops in oil production. Narrative records are then used to classify these episodes as endogenous or exogenous. Finally, for each country and month, all the exogenous declines in other countries are used to construct an instrument. Estimates of the short-run elasticity of supply vary from 0.054 for narrow set of IV to 0.081 for a broader IV set. A larger supply elasticity of 0.10 is implied by their baseline VAR when a pair of short-run supply and demand (in production) elasticities is selected by minimizing the Euclidean distance between the VAR admissible elasticities and the target elasticities ($\alpha_{qp} = 0.081$ and $\beta_{qp} = -0.080$). Thus, for our last thought experiment we bound the prior above by 0.10, $\alpha_{qp} \in (0, 0.10]$.²⁷ Of course, unlike the micro estimates discussed earlier, this estimate cannot be considered extraneous evidence since it was derived from a model and data set similar to that of *BH18* in several dimensions, nor is it based on micro-level evidence. Thus, our analysis of this oil supply elasticity value is only exploratory.

Table 2 reports the results obtained with these two alternative priors. The estimation results confirm that when the upper bound on the support of α_{qp} increases, the responsiveness of real oil prices to oil supply shocks rises while that of the world industrial production decreases. As Figure 2a illustrates, the median response of oil production varies little while, not surprisingly, a tighter bound on the prior results in a tighter 95% credible set for the impact response. Regarding the real oil price, a tighter upper bound on the prior leads to a smaller response on impact and overtime. Indeed, significant differences are observed across alternative priors in the response of the real oil price twelve months after the shock (see Figure 2a). As for real GDP, imposing a prior that attaches a larger probability mass to values of α_{qp} away from zero, would lead to infer that the recessionary effect of oil supply shocks is somewhat smaller (see Figure 2b). This result is consistent with the decreasing responsiveness of the world industrial production index as the prior is relaxed (see last column of Table 2). In brief, the prior on α_{qp} plays a key role in the estimated responsiveness of

²⁷Because the baseline VAR of Caldara, Cavallo and Iacoviello (2018) does not include inventories, the target for β_{qp} is akin to a demand elasticity in production and not in use. Hence, the pair of admissible elasticities is not directly comparable to Kilian and Murphy (2014).

real oil prices and U.S. real economic activity to oil supply shocks.

We conclude this section with a caveat. We have restricted ourselves to investigate alternative priors on the short-run supply elasticity by maintaining the prior Student t distribution with location parameter, $c_{qp}^\alpha = 0.1$, scale parameter $\sigma_{qp}^\alpha = 0.2$, and $v_{qp}^\alpha = 3$ degrees of freedom, but truncating it above at different values suggested by the existing literature.²⁸ Exploring the role of alternative prior distributions is a research question that we leave for future work. For example, imposing an exponential prior would have been more economically appealing than the Student- t prior. However, this section shows that for any prior that is consistent with state-of-the-art microeconomic estimates, the impulse response estimates for *BH18* imply effects on real oil prices and U.S. real GDP that are very similar to the earlier studies of *K09*, *K12*, and *K14*.

7 Other differences in specification

Additional specification differences between *BH18* and their predecessors comprise the use of earlier data, the VAR lag length, the loss function –median response function vs. modal structural model–, and the measure of real economic activity. We discuss the role of these differences here, but relegate the figures to the online appendix (see Figure A.3).

Before 1973 the price of oil was regulated in the U.S. by the Texas Railroad Commission (see e.g., Hamilton 1983, 1985). A new regime in the global oil market seems to have begun in the early 1970s as the U.S. dependence on oil imports increased and oil prices became more responsive to supply and demand movements. It has been noted that there was a major break in the process governing the price of oil in late 1973 as well as a major break in the correlation between growth in real GDP and the real price of oil (see, e.g., Alquist, Kilian and Vigfusson 2013). These changes in the global oil market motivated *K09*, *KM12*, and *KM14*'s choice to restrict the sample to the post-1973 period. In contrast, *BH18* use earlier data to construct their prior but assumed the pre-1973 data to be 1/2 or 1/4 less informative ($\mu = 0.5, 0.25$) than the later data. We further investigate the sensitivity of *BH18* and our results by setting $\mu = 0$, which amounts to using the post-1973 data only to form the priors. Table 2 reveals only minimal changes in the estimated short-run elasticities of demand and supply when $\mu = 0$. Yet, the researcher would infer that oil supply shocks have a somewhat smaller effect on the world's industrial production. We find that real oil prices are slightly more responsive to oil supply shocks and inference regarding the U.S.

²⁸Estimation results using a location parameter of 0.01 result only in a slight reduction in the posterior median for the short-run elasticity of supply.

real GDP differs only slightly from the baseline. In brief, the impact of reducing μ to 0 is small.

BH18's SVAR model includes 12 lags. To explore the sensitivity of the results to the larger number of lags included in previous studies, we re-estimate *BH18* including 24 lags but maintaining all other features of *BH18's* model.²⁹ Increasing the number of lags results in a slight decrease in the short-run price elasticity of demand, and a very small increase in the effect of oil supply shocks on world industrial production 12 months after the shock (see Table 2). The researcher would infer a slightly larger increase in real oil prices. Increasing the number of lags suggest oil supply shocks have no effect on U.S. real GDP in the short run and have a similar effect in the long run. Summarizing, the only noticeable change of increasing the number of lags is the smaller impact response of U.S. real GDP to oil supply shocks.

As for the loss function, comparing the median pointwise responses and the response for the modal model in *KM14* reveals a slightly larger and more persistent increase in oil prices and a somewhat bigger reduction in U.S. real GDP in the long run in the modal model (see Figure 1b). However, the divergences between the modal model response function and the pointwise median for *KM12*, and *KM14*, are small. That is, the results do not appear to be very sensitive to using the pointwise median or the structural modal model to conduct inference.

Finally, to investigate how alternative measures of real economic activity affect the inferences derived from alternative models we re-estimate *KM09*, *KM12* and *KM14* replacing Kilian's real economic activity measure with the world IP used by *BH18*.³⁰ More specifically, we use a linearly detrended world IP series in the *K09*, *KM12*, and *KM14* specifications. In addition, we investigate the role of the real activity measure in *BH18* by replacing the world IP with Kilian's Index of real economic activity and re-estimating the model. Because the real activity index starts in the 1970s we set up $\mu = 0$; as we showed earlier, down-weighting the earlier data has virtually no effect on the responses.

On the one hand, estimation results reported in Table 3 indicate that using the world IP results in a larger increase in real oil prices a year after the shock as well as a bigger loss in U.S. real GDP in *K09*, *KM12*, and *KM14* (see Figure A.4 in the online appendix). Yet, estimates of the supply and demand elasticities in these models are rather robust to the change in the measure of real economic activity. On the other hand, when we use Kilian's real economic activity index in *BH18*, we find that the pointwise posterior median for the short-run elasticity of supply is about

²⁹In particular, we use the same informative prior on the lags used by *BH18*.

³⁰We use $M = 5000$ draws from the joint posterior distribution of the reduced-form VAR parameters with $N = 2000$ rotations each for *KM12* and *KM14* models.

half of that in the benchmark model (0.071 instead of 0.144). In turn, the elasticity of demand in production is almost twice as large (-0.604 instead of -0.356). As a result, the impact response of the real oil price to an unexpected 1% decline in world oil supply is considerably smaller than in the benchmark model (1.692% instead of 2.561%), as is the one-year loss in U.S. real GDP (-0.140% instead of -0.212%). In brief, the use of alternative measures of real economic activity plays an important role in explaining the differences in the estimated response of real oil prices and U.S. real GDP. If we condition on the same measure of global economic activity and we consider the dynamic GDP-real oil multiplier, the differences in the responses are considerably smaller than they initially appear.

8 Conclusions

Estimates of the dynamic effect of oil supply disruptions on real oil prices differ greatly across studies. The source for these dissimilarities can be traced to differences in methodology, especially alternative identification schemes and model specification. This paper reviewed recent SVAR models used to study the world oil market and evaluated the effect of the implied structural supply shocks on U.S. real GDP.

We found that models that impose a prior that attaches a large probability mass to large values of the short-run price elasticity of oil supply such as Baumeister and Hamilton (2018) result in larger estimates of the short-run price elasticity of oil supply and smaller estimates of the short-run price elasticity of demand in production. They also imply a larger response of real oil prices to oil supply disruptions and a larger and longer-lived contraction in U.S. real GDP. We also found that *LN12*, a sign-identified model that separates U.S. and rest-of-the-world demand and supply shocks, leads to estimates that fall outside the ballpark magnitudes obtained by the above listed studies. The credible bands attached to these estimates are large, which suggests sign restrictions are not enough to pin down the quantities of interest.

The GDP-oil price dynamic multiplier conditional on an oil supply shock is estimated to fall in a range between -0.01 and -0.06. Thus, the effect of an oil supply shock that leads to a 10% increase in real oil prices is estimated to result in a contraction of U.S. real GDP four quarters after the shock ranges between 0.1% and 0.6% for *KM14* and *LN18*, respectively. In addition, estimation results reported in the appendix show that extending the sample period for the TVP-BVAR model of *BP13* implies a larger impact response of oil prices and a greater one-year GDP loss following the Great Recession than in earlier years. Nevertheless, the estimates obtained for the later samples

are rather imprecise.

To further investigate the role of the prior on the short-run price elasticity of oil supply, α_{qp} , we re-estimated *BH18* after truncating the prior on α_{qp} in a manner that shifted probability mass closer to zero. Specifically, we first truncated the prior at the 0.0258 bound imposed in *K12* and *KM14*; we then allowed for larger bounds motivated by recent microeconomic studies (e.g., Bjørnland, Nordvik and Rohrer 2017). We found that imposing a tighter prior on α_{qp} , which rules out the possibility that the elasticity could be very large, considerably reduces the differences between the estimated short-run price elasticity of supply obtained in *BH18* and *KM14*, yet it considerably increases the estimate of the elasticity of demand in production for *BH18*. Moreover, imposing a tighter prior on α_{qp} in *BH18* reduces the responsiveness of real oil prices to oil supply disruptions, but *has* little impact on the response of U.S. real GDP.

We also found that the responsiveness of real oil prices and U.S. real GDP is somewhat affected by the SVAR lag length and the measure of world economic activity. Namely, a researcher would infer a smaller contraction in *BH18*, if she doubled the lag length in line with earlier studies or, especially, if she replaced the world IP with Kilian’s real economic activity index. In contrast, the researcher would infer a larger contraction if she replaced Kilian’s real economic activity index with the world IP measure used by *BH18* in *K09*, *KM12*, or *KM14*.

What have we learned? First, future investigations into the role of oil supply (and demand) shocks on real oil prices should heed the advice of Kilian and Murphy (2014) model storage demand in SVAR models for the crude oil market. Doing so is key in modeling the short-run comovement between oil production and prices. Second, even if inventories are accounted for and the short-run price elasticity of demand has been pinned down, the importance of oil supply shocks in generating recessions hinges heavily on the value of the short-run price elasticity of oil supply.

The question is whether to consider priors that assign a high probability mass to values of the short-run price elasticity of supply far away from zero. For example, the analysis in *BH18* allowed for a positive probability mass on elasticity values approaching infinity and assigned 94% prior probability mass to elasticity values larger than the bound of imposed by *KM14*. This approach is difficult to support based on extraneous evidence. Both the economic theory in Anderson et al. (2018) and extraneous microeconomic estimates argue against priors with large probability mass on higher values of the price elasticity of oil supply. Estimates based only on data from North Dakota by Bjørnland, Nordvik and Rohrer (2017) imply a slightly higher global price elasticity of 0.04 whereas those of Newell and Prest (2018) imply an elasticity of 0.01. The highest plausible

bound is distinctly smaller than the posterior elasticity estimate of 0.15 reported by Baumeister and Hamilton (2018). We conclude that the apparent disagreements in the literature regarding the effect of oil supply shocks on real oil prices are not nearly as large as they seem if we are willing to condition on a range of supply elasticity values that is supported by external evidence.

References

- [1] Alquist, R., Kilian, L., and Vigfusson, R. J. (2013). "Forecasting the price of oil," In *Handbook of economic forecasting* (Vol. 2, pp. 427-507). Elsevier.
- [2] Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant (2018). "Hotelling under pressure," *Journal of Political Economy* 126(3), 984-1026.
- [3] Backus, D. K., and Crucini, M. J. (2000). "Oil prices and the terms of trade," *Journal of international Economics*, 50(1), 185-213.
- [4] Barsky, R. B., and Kilian, L. (2001). "Do we really know that oil caused the great stagflation? A monetary alternative," NBER Macroeconomics annual, 16, 137-183.
- [5] Barsky, R. B., and Kilian, L. (2004). Oil and the Macroeconomy since the 1970s," *Journal of Economic Perspectives*, 18(4), 115-134.
- [6] Baumeister, C. and Hamilton, J. D. (2018). "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks," *American Economic Review*, forthcoming.
- [7] Baumeister, C., and Kilian, L. (2014). "Real-time analysis of oil price risks using forecast scenarios," *IMF Economic Review*, 62(1), 119-145.
- [8] Baumeister, C., and Kilian, L. (2016). "Forty years of oil price fluctuations: Why the price of oil may still surprise us," *Journal of Economic Perspectives*, 30(1), 139-60.
- [9] Baumeister, and Peersman, G. (2013a). "The role of time-varying price elasticities in accounting for volatility changes in the crude oil market," *Journal of Applied Econometrics* 28 (7), 1087-1109.
- [10] Baumeister, C., and Peersman, G. (2013b). "Time-varying effects of oil supply shocks on the US economy," *American Economic Journal: Macroeconomics*, 5(4), 1-28.

- [11] Bjørnland, H. C., Frode, M. N., and Rohrer, M. (2017). "Supply Flexibility in the Shale Patch: Evidence from North Dakota," Manuscript.
- [12] Blanchard, Olivier J. and Jordi Galí (2010). "The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so Different from the 1970s?," in J. Galí and M. Gertler (eds). *International Dimensions of Monetary Policy*, University of Chicago Press, Chicago, IL, (2010): 373-428.
- [13] Bornstein, G., Krusell, P., and Rebelo, S. (2017). "Lags, costs, and shocks: An equilibrium model of the oil industry," (No. w23423). National Bureau of Economic Research.
- [14] Caldara, D., Cavallo, M., and Iacoviello, M. (2018). "Oil price elasticities and oil price fluctuations," *Journal of Monetary Economics*.
- [15] Dvir, E., and Rogoff, K. S. (2009). "*Three epochs of oil*," (No. w14927). National Bureau of Economic Research.
- [16] Edelstein, P., and Kilian, L. (2007). "The response of business fixed investment to changes in energy prices: A test of some hypotheses about the transmission of energy price shocks," *The BE Journal of Macroeconomics*, 7(1).
- [17] Edelstein, P., and Kilian, L. (2009). "*How sensitive are consumer expenditures to retail energy prices?*," *Journal of Monetary Economics*, 56(6), 766-779.
- [18] Ferrara, L., and Guérin, P. (2018). "What are the macroeconomic effects of high-frequency uncertainty shocks?," *Journal of Applied Econometrics*, 33(5), 662-679.
- [19] Fry, R., and Pagan, A. (2011). "Sign restrictions in structural vector autoregressions: A critical review," *Journal of Economic Literature*, 49(4), 938-60.
- [20] Hamilton, J. D. (1983). "Oil and the macroeconomy since World War II," *Journal of political economy*, 91(2), 228-248.
- [21] Hamilton, J. D. (1985). "Historical causes of postwar oil shocks and recessions," *The Energy Journal* 6(1), 97-116.
- [22] Hamilton, J. D. (2009a). "Causes and consequences of the oil shock of 2007–08," *Brookings Papers on Economic Activity* 1, 215–261.
- [23] Hamilton, J. D. (2009b). "Understanding crude oil prices," *Energy Journal* 30, 179–206.

- [24] Hamilton, J. D., and Herrera, A.M. (2004). "Comment: oil shocks and aggregate macroeconomic behavior: the role of monetary policy," *Journal of Money, Credit and Banking*, 265-286.
- [25] Hausman, J. A., and Newey, W. K. (1995). "Nonparametric estimation of exact consumers surplus and deadweight loss," *Econometrica*, 1445-1476.
- [26] Herrera, A. M. and Pesavento, E. (2009). "Oil Price Shocks, Systematic Monetary Policy and the 'Great Moderation'," *Macroeconomic Dynamics* 13(1), 107-137.
- [27] Inoue, A. and Kilian, L. (2013). "Inference on impulse response functions in structural VAR models," *Journal of Econometrics* 177(1), 1-13.
- [28] Inoue, A. and Kilian, L. (forthcoming). "Corrigendum to "Inference on impulse response functions in structural VAR models," [Journal of Econometrics 177 (2013), 1-13]. *Journal of Econometrics*.
- [29] Jordà, Ò. (2005). "Estimation and inference of impulse responses by local projections," *American Economic Review*, 95(1), 161-182.
- [30] Kellogg, R. (2011). "Short-run responses of oil production to spot and future oil prices," mimeo, University of Michigan.
- [31] Kilian, L. (2008). "The Economic Effects of Energy Price Shocks," *Journal of Economic Literature*, 871-909.
- [32] Kilian, L. (2009). "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *The American Economic Review*, 99(3), 1053-1069.
- [33] Kilian, L. (2016). "The impact of the shale oil revolution on US oil and gasoline prices," *Review of Environmental Economics and Policy*, 10(2), 185-205.
- [34] Kilian, L. (2017). "The Impact of the Fracking Boom on Arab Oil Producers," *Energy Journal*, 38(6), 137-160.
- [35] Kilian, L. and Lee, T. K. (2014). "Quantifying the speculative component in the real price of oil: The role of global oil inventories," *Journal of International Money and Finance* 42, 71-87.
- [36] Kilian, L. and Lütkepohl, H. (2017). "Structural vector autoregressive analysis," Cambridge University Press.

- [37] Kilian, L. and 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.
- [38] Kilian, L. and Murphy, D. P. (2014). "The role of inventories and speculative trading in the global market for crude oil," *Journal of Applied Econometrics* 29(3), 454-478.
- [39] Kilian, L. and Zhou, X. (2018a). "Modeling Fluctuations in the Global Demand for Commodities," *Journal of International Money and Finance*, 88, 54-78.
- [40] Kilian, L. and Zhou, X. (2018b). "Structural Interpretation of Vector Autoregressions with Incomplete Information: Revisiting the Role of Oil Supply and Demand Shocks: Comment," available at http://www-personal.umich.edu/~lkilian/kz0_071518.pdf.
- [41] Lippi, F. and Nobili, A. (2012). "Oil and the macroeconomy: a quantitative structural analysis," *Journal of the European Economic Association* 10(5), 1059-1083.
- [42] Newell, R. G., and Prest, B. C. (2018). "The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata," *The Energy Journal*, forthcoming.
- [43] Yatchew, A. and No, J. A. (2001). "Household gasoline demand in Canada," *Econometrica* 69(6), 1697-1709.
- [44] Zhou, X. (2019), "Refining the Workhorse Oil Market Model," manuscript, Bank of Canada.

Table 1. Summary of Main Estimates (1973:1-2016:12)

Panel A: Monthly SVAR Estimates					
	<i>K09</i>	<i>KM12</i>	<i>KM14</i>	<i>BH18</i>	<i>LN12</i>
Short-run price elasticity of supply	0	0.015 (0.019) (0.000, 0.024)	0.013 (0.017) (0.0007, 0.025)	0.144 (0.052, 0.367)	0.447 (0.071, 10.551)
Short-run price elasticity of demand in production	-3.131	-1.743 (-1.723) (-6.072, -1.038)	-0.511 (-0.546) (-1.103, -0.145)	-0.356 (-0.803, -0.152)	-0.087 (-1.638 -0.005)
Short-run price elasticity of demand in use			-0.275 (-0.260) (-0.760, -0.016)		
Impact response of real oil price to supply shock	0.322	0.580 (1.170, 0.992)	1.675 (0.870, 3.344)	2.561 (1.195, 3.008)	12.073 (0.691, 23.049)
Response of real oil price 12-months after the shock	0.734	1.171 (-0.053, 3.050)	-0.735 (-1.590, 0.584)	3.743 (1.195, 3.008)	7.147 (-6.675, 19.937)
Panel B: Quarterly Estimates					
	<i>K09</i>	<i>KM12</i>	<i>KM14</i>	<i>BH18</i>	<i>LN12</i>
Impact response of US real GDP	-0.302	-0.269 (-0.260) (-0.287 -0.253)	-0.146 (-0.169) (-0.224 -0.040)	-0.042 (-0.119 0.009)	-0.454 (-0.689 -0.678)
Response of US real GDP 4-quarter after the shock	-0.144	-0.357 (-0.323) (-0.393 -0.320)	-0.135 (-0.150) (-0.474 0.108)	-0.212 (-0.361 -0.097)	-0.515 (-1.068 0.050)
Dynamic multiplier ($h = 0$)	-0.098	-0.061 (-0.072) (-0.076 -0.052)	-0.016 (-0.040) (-0.038 -0.004)	-0.006 (-0.018 0.001)	-0.044 (-0.605 0.556)
Dynamic multiplier ($h = 4$)	-0.023	-0.038 (-0.040) (-0.040 -0.035)	-0.012 (-0.038) (-0.069 0.0121)	-0.017 (-0.032 -0.007)	-0.055 (-1.218 1.240)
Panel C: Alternative estimates of B_0^{-1}					
	<i>KM14</i>	<i>BH18</i>			
0.011	0.000	0.007	0.695	0.172	0.409
1.985	4.337	-3.784	0.005	0.997	-0.006
-0.020	0.045	0.018	-2.021	1.142	2.588
5.088	-2.501	-8.430	-0.022	-0.065	0.845

Notes: Panel A reports the magnitudes implied by *K09*; the posterior median, the modal model (in bold) and the 95% HPD for *KM12* and *KM14*; and the posterior median and the 95% credible sets (in parenthesis) for *LN12* and *BH18*. Panel B reports quarterly estimates for real GDP derived from equation (6). The dynamic multipliers are defined as in (8). Panel C reports alternative estimates of B_0^{-1} . The estimates correspond to the modal model for *KM14* and the posterior median for *BH18*.

Table 2. Sensitivity Analysis for *BH18*

	Short-run price elasticity of supply	Short-run price elasticity of demand in production	Effect of supply shock that raises real oil price by 10% on economic activity 12 months later
Benchmark	0.144 (0.094, 0.222)	-0.356 (-0.508, -0.242)	-0.503 (-0.908, -0.168)
Alternative support for prior on α_{qp}			
$\alpha_{qp} \in (0, 0.0258]$	0.019 (0.011, 0.024)	-1.386 (-1.779, -1.115)	-1.601 (-2.692, -0.665)
$\alpha_{qp} \in (0, 0.04]$	0.033 (0.022, 0.038)	-1.014 (-1.406, -0.815)	-1.239 (-2.097, -0.516)
$\alpha_{qp} \in (0, 0.10]$	0.082 (0.062, 0.094)	-0.533 (-0.460, -0.654)	-0.758 (-1.244, -0.313)
Down-weighting the earlier sample			
$\mu = 0$	0.140 (0.088, 0.223)	-0.350 (-0.499, -0.238)	-0.163 (-0.524, 0.156)
Increasing the number of lags			
$p = 24$	0.141 (0.094, 0.214)	-0.311 (-0.440, -0.213)	-0.520 (-0.908, -0.192)

Notes: This table reports the posterior median and the 68% credible sets (in parenthesis) implied by the modified *BH18* specification.

Table 3. Alternative measures of economic activity

Panel A: Monthly SVAR Estimates												
	<i>K09</i>			<i>KM12</i>			<i>KM14</i>			<i>BH18</i>		
Measure of economic activity	REA index	World IP	World IP	REA index	World IP	World IP	REA index	World IP	World IP	REA index	World IP	
Short-run price elasticity of supply	0	0	0.01	0.02	(0.00, 0.02)	(0.00, 0.02)	0.01	0.01	0.01	0.07	0.14	
Short-run price elasticity of demand in production	-3.13	-2.37	-1.57	-1.74	(-6.07, -1.04)	(-4.35, -1.01)	-0.51	-0.61	-0.61	(0.02, 0.29)	(0.05, 0.37)	
Short-run price elasticity of demand in use							(-1.10, -0.15)	(-1.20, -0.19)	(-1.20, -0.19)	-0.60	-0.36	
Impact response of real oil price to supply shock	0.32	0.40	0.63	0.58	(1.17, 0.99)	(0.23, 0.97)	1.68	1.62	1.62	1.69	2.56	
Response of real oil price 12-months after the shock	0.73	1.45	1.77	1.17	(-0.05, 3.05)	(0.23, 4.04)	(0.87, 3.34)	0.09	0.09	(0.91, 2.47)	(1.20, 3.01)	
	(-0.85, 2.96)	(-0.24, 3.83)					(-1.59, 0.58)	(-0.98, 2.85)	(-0.98, 2.85)	2.98	3.74	
										(0.78, 4.93)	(1.20, 3.01)	

Panel B: Quarterly Estimates												
	<i>K09</i>			<i>KM12</i>			<i>KM14</i>			<i>BH18</i>		
Measure of economic activity	REA index	World IP	World IP	REA index	World IP	World IP	REA index	World IP	World IP	REA index	World IP	
Impact response of US real GDP	-0.30	-0.29	-0.27	-0.27	(-0.29 -0.25)	(-0.29 -0.24)	-0.15	-0.23	-0.23	-0.09	-0.04	
Response of US real GDP 4-quarters after the shock	-0.14	-0.35	-0.36	-0.36	(-0.39 -0.32)	(-0.42 -0.30)	(-0.22 -0.04)	(-0.38 -0.11)	(-0.38 -0.11)	(-0.17 -0.00)	(-0.12 0.00)	
Dynamic multiplier ($h = 0$)	-0.10	-0.08	-0.06	-0.06	(-0.07 -0.05)	(-0.08 -0.05)	-0.13	-0.31	-0.31	-0.14	-0.21	
Dynamic multiplier ($h = 4$)	-0.02	-0.04	-0.04	-0.03	(-0.04 -0.03)	(-0.05 -0.03)	(-0.47 0.11)	(-0.87 0.03)	(-0.87 0.03)	(-0.36 0.00)	(-0.36 -0.10)	
							-0.02	-0.04	-0.04	-0.02	-0.01	
							(-0.04 -0.00)	(-0.09 -0.01)	(-0.09 -0.01)	(-0.05 0.00)	(-0.02 0.00)	
							-0.01	-0.05	-0.05	-0.02	-0.02	
							(-0.07 0.01)	(-0.66 0.02)	(-0.66 0.02)	(-0.06 0.01)	(-0.03 -0.01)	

Notes: This table reports the magnitudes implied by *K09*; the posterior median and the 95% HPD (in parenthesis) for *KM12* and *KM14*; the posterior median and the 95% credible sets (in parenthesis) for *BH18*. Panel B reports the quarterly estimates for real GDP derived from equation (6). The dynamic multiplier estimates are derived from equation (8).

Figure 1a: Monthly Response of Oil Production, and Real Price of Oil to an Unexpected Decline in World Oil Supply

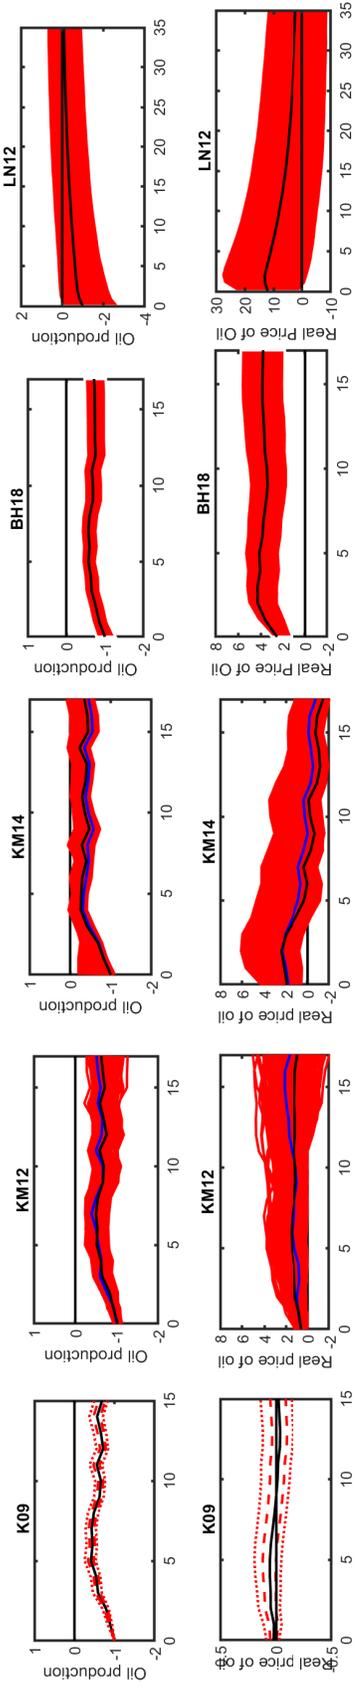
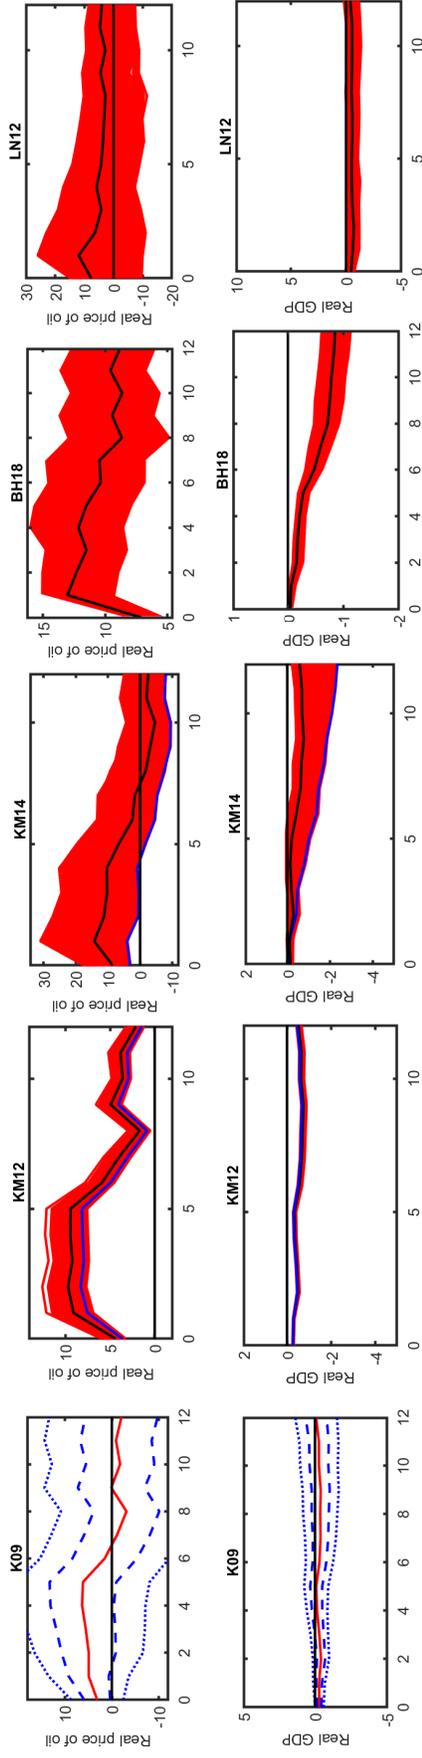


Figure 1b: Quarterly Response of Real Price of Oil and Real GDP to an Unexpected Decline in World Oil Supply



Notes: Figure 1a plots response to an oil supply disruption normalized to represent the effect of an unexpected 1% monthly decline in world oil production for each of the SVAR specifications. We report monthly response of real price of oil. The point estimates for *K09* are reported with one and two standard error bands that are constructed using a recursive-design wild bootstrap. The responses for the modal model (blue line) and the 95% joint regions of high posterior density (HPD) are reported for *KM12*, *KM14*. The black line denotes the posterior median response of real oil price for *KM12*, *KM14*. We report posterior median response and 95% posterior credible sets for *LN12* model. The black line denotes Bayesian posterior median responses and shaded regions denote 95% posterior credible sets for *BH18*. Figure 1b plots quarterly response of Real GDP and real price of oil. These estimates are derived from equation (6) and (7).

Figure 2a: Monthly Response of Oil Production, and Real Price of Oil to Unexpected Decline in World Oil Supply under Alternative Priors for α_{qp}

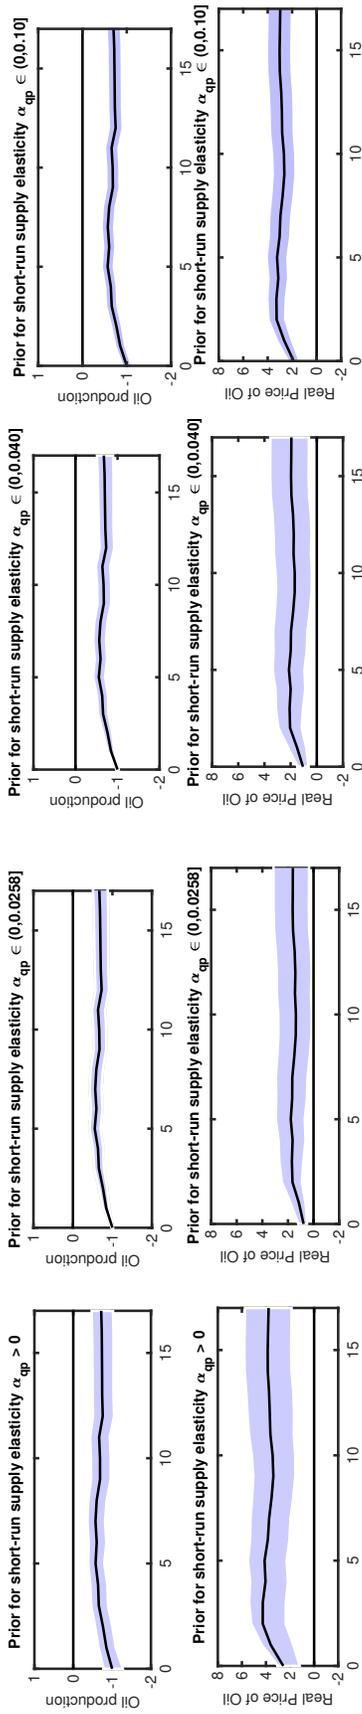
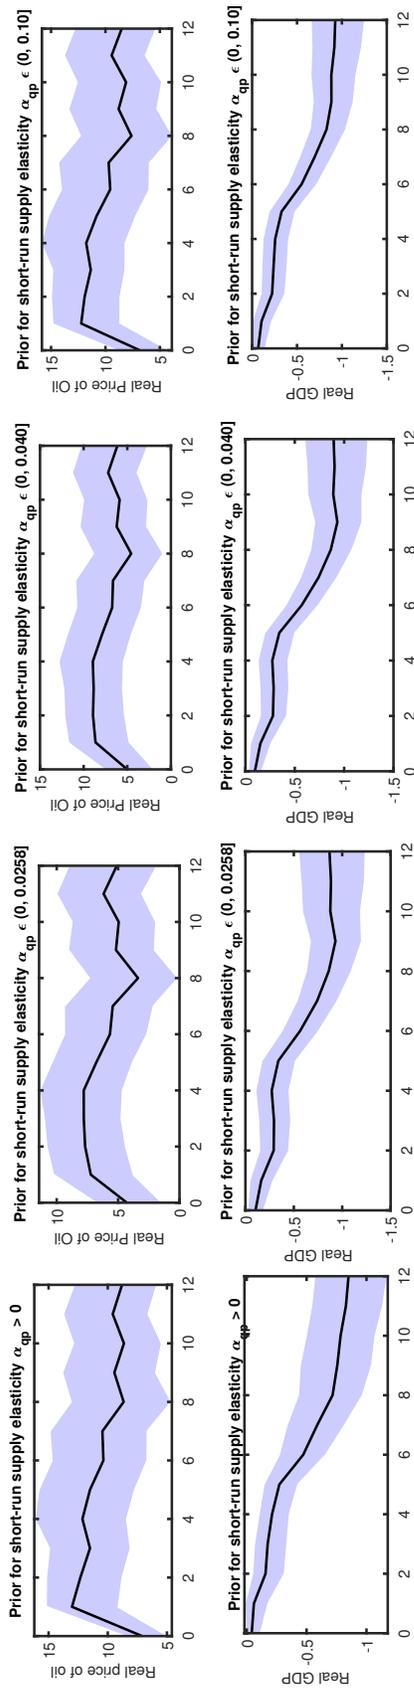


Figure 2b: Quarterly Response of Real Price of Oil, and Real GDP to Structural Oil Supply Shocks under Alternative Priors for α_{qp}



Notes: Figure 2a plots response to an oil supply disruption normalized to represent the effect of an unexpected 1% monthly decline in world oil production for *BH18* imposing a modified prior on α_{qp} . We report monthly response real price of oil. The black line denotes Bayesian posterior median responses and shaded of regions denote 95% posterior credible sets. The Figure 2b plots the response of real price of oil and Real GDP for *BH18* imposing a modified prior on α_{qp} . Estimates are derived from equation (6) and (7). The black line denotes Bayesian posterior median responses and shaded regions denote 95% posterior credible sets.