Federal Reserve Bank of Dallas

Refining the Workhorse Oil Market Model

Xiaoqing Zhou

Working Paper 1910

Research Department

https://doi.org/10.24149/wp1910

Refining the Workhorse Oil Market Model*

Xiaoqing Zhou†

September 6, 2019

Abstract

The Kilian and Murphy (2014) structural vector autoregressive model has become the workhorse model for the analysis of oil markets. I explore various refinements and extensions of this model, including the effects of (1) correcting an error in the measure of global real economic activity, (2) explicitly incorporating narrative sign restrictions into the estimation, (3) relaxing the upper bound on the impact price elasticity of oil supply, (4) evaluating the implied posterior distribution of the structural models, and (5) extending the sample. I demonstrate that the substantive conclusions of Kilian and Murphy (2014) are largely unaffected by these changes.

Keywords: Oil market; global real activity; structural VAR; narrative sign restrictions; identification; Bayesian inference.

JEL Codes: C32, C52, Q41, Q43.

---

*I thank Jamie Cross, Lutz Kilian and three anonymous referees for helpful comments. An earlier version of this paper was completed while the author was at the Bank of Canada. The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas, the Federal Reserve System, or the Bank of Canada.

†Xiaoqing Zhou, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, xqzhou3@gmail.com.

1 Introduction

Kilian and Murphy (2014) introduced a structural vector autoregressive (VAR) model of the global oil market that for the first time explicitly incorporated storage demand and allowed for forward-looking behavior in oil markets. This model has become the workhorse model for the analysis of oil markets (e.g., Fattouh et al. (2013); Kilian and Lee (2014); Kilian (2017); Herrera and Rangaraju (2019); Cross (2019)). It has been used, for example, to study the role of financial speculation in oil markets, the impact of the U.S. shale oil boom, and the causes of the 2014 oil price decline. While the framework developed by Kilian and Murphy remains as relevant today as when it was introduced, there have been a number of econometric advances as well as new substantive insights in the literature since this paper was written that make it worthwhile to reexamine the substantive conclusions reached by Kilian and Murphy (2014). My paper studies how various refinements, corrections, and extensions motivated by recent advances in the literature affect the results of Kilian and Murphy (2014). Any of these changes, alone or in conjunction, has the potential of substantially changing the estimates of the Kilian-Murphy model. I demonstrate that the substantive conclusions reached by Kilian and Murphy (2014) remain largely unaffected by these changes, lending further credence to their results.

Specifically, I consider five changes. First, I replace the measure of global real activity underlying the analysis by the corrected version of this measure recently introduced by Kilian (2019). Second, Kilian and Murphy (2014) relied on narrative inequality restrictions for selecting the most credible model among the set of model solutions that satisfy the sign restrictions on the impulse responses (see also Kilian and Lee (2014)). I incorporate these narrative restrictions into the estimation of the model rather than imposing them based on the visual inspection of the historical decomposition of the admissible models as in Kilian and Murphy (2014). Third, I relax the upper bound of the impact oil supply elasticity used by Kilian and Murphy (2014) to bring the analysis in line with the recent literature on microeconometric estimates of the oil supply elasticity (see, e.g., Bornstein et al. (2018), Anderson et al. (2018), Newell and Prest (2019)). Fourth, Kilian and Murphy (2014) reported estimates of their sign-identified VAR model conditional on the maximum likelihood estimate (MLE) of the reduced-form VAR model. They did not report credible sets for their preferred estimate because, at the time, no satisfactory solutions existed for the econometric evaluation of the posterior distribution (see Kilian and Lütkepohl (2017)). Since then appropriate econometric methods have been developed (see Inoue and Kilian (2013, 2019)). I use these methods to provide estimates of the most likely structural model estimate as well as joint credible sets for impulse responses and historical decompositions. A similar approach was recently used by Herrera and Rangaraju (2019). The key difference from their work is that my credible sets incorporate, in addition, the narrative restrictions used by Kilian and Murphy (2014), which helps sharpen inference. My analysis also complements the work of Antolin-Diaz and Rubio-Ramirez (2018) by illustrating how the most likely admissible model and joint credible sets may be computed in the presence of narrative restrictions. Finally, I extend the sample to June 2018.

The remainder of the paper is organized as follows. Section 2 briefly reviews the original Kilian and Murphy (2014) model and discusses the baseline specification for my analysis. In section 3, I show that using the corrected index of global real economic activity and relaxing the elasticity bound does not materially change the results obtained conditional on the MLE. Section 4 presents estimates of the most likely model and joint credible sets based on the posterior distribution of the admissible models. I show that the conclusions of Kilian and Murphy (2014) are robust to allowing for estimation uncertainty. Section 5 shows that their key results are also robust to extending the sample. In addition, I explore the effects of replacing the Kilian index in their VAR model by a measure of global industrial production. The concluding remarks are in section 6.

2 The Baseline Structural VAR Model

The Kilian and Murphy (2014) VAR model includes four monthly variables: (1) the percent change in the global production of crude oil ($\triangle q_t$), as reported by the U.S. Energy Information Administration; (2) a measure of cyclical variation in global real economic activity (rea_t) originally proposed by Kilian (2009); (3) the log real price of oil (p_t) obtained by deflating the U.S. refiners' acquisition cost for imported crude oil by the U.S. CPI for all urban consumers; and (4) a proxy for the change in global crude oil inventories ($\triangle inv_t$), as discussed in Kilian and Murphy (2014). I use the same variables, except that throughout this paper I employ the most recent version of the index of global real economic activity, as reported in Kilian (2019).[1] The baseline model is estimated on data from February 1973 to August 2009, as in Kilian and Murphy (2014).

Let $y_t = (\triangle q_t, rea_t, p_t, \triangle inv_t)'$ be generated by the stationary structural VAR(24) process

$$B_0 y_t = B_1 y_{t-1} + ... + B_{24} y_{t-24} + w_t,$$

where the error w_t is mutually uncorrelated i.i.d. white noise and the deterministic terms have been suppressed for expository purposes. Imposing a lag order of 24 allows the model to capture long cycles in the real price of oil and avoids the pitfalls of data-based lag order selection (Kilian and Lütkepohl (2017)). The reduced-form errors can be expressed as $u_t = B_0^{-1} w_t$, where

$$u_t = y_t - A_1 y_{t-1} - ... - A_{24} y_{t-24},$$

$A_l = B_0^{-1} B_l, l = 1, ..., 24$, and the $\{ij\}$th element of B_0^{-1}, denoted b_{ij}^0, represents the impact response of variable i to structural shock j, where $i \in \{1, ..., 4\}$ and $j \in \{1, ..., 4\}$. Given B_0^{-1}, we can recover estimates of the structural impulse responses, variance decompositions and historical decompositions from the reduced-form parameter estimates, as discussed in Kilian and Lütkepohl (2017).

Estimation of B_0^{-1} requires identifying restrictions on the effects of the structural shocks on the model variables. Let $w_t = \left(w_t^{flow\ supply}, w_t^{flow\ demand}, w_t^{storage\ demand}, w_t^{other\ oil\ demand}\right)'$, where $w_t^{flow\ supply}$ denotes a shock to the flow supply of oil, $w_t^{flow\ demand}$ denotes a shock to the flow demand for oil, $w_t^{storage\ demand}$ is a shock to storage demand (sometimes referred to as a speculative demand shock), and $w_t^{other\ oil\ demand}$ is a residual shock designed to capture all other shocks to the demand for oil such as preference shocks, shocks to the storage technology, or politically motivated changes in the Strategic Petroleum Reserve. As in the related literature, my analysis focuses on the first three shocks in this model that have an explicit structural interpretation. All shocks are normalized to represent a shock that

[1]The construction of this index is described in Kilian (2009). The original data and subsequent updates of this series were incorrect because of an unintentional log transformation, as noted by Hamilton (2019). A corrected index was proposed by Kilian (2019), and is available at https://sites.google.com/site/lkilian2019/research/data-sets. The index used for the models considered in sections 3 and 4 has been constructed using only the raw data available to Kilian and Murphy (2014).

increases the real price of oil.

The sign restrictions on the elements of B_0^{-1} are summarized in equation (1):

$$u_t = \begin{pmatrix} u_t^{\triangle q} \\ u_t^{rea} \\ u_t^{p} \\ u_t^{\triangle inv} \end{pmatrix} = B_0^{-1} w_t = \begin{bmatrix} - & + & + & b_{14}^0 \\ - & + & - & b_{24}^0 \\ + & + & + & b_{34}^0 \\ b_{41}^0 & b_{42}^0 & + & b_{44}^0 \end{bmatrix} \begin{pmatrix} w_t^{flow\ supply} \\ w_t^{flow\ demand} \\ w_t^{storage\ demand} \\ w_t^{other\ oil\ demand} \end{pmatrix}. \quad (1)$$

In addition, I follow Kilian and Murphy (2014) in imposing bounds on the one-month price elasticities of oil demand and oil supply, which may be expressed as inequality restrictions on functions of selected impact responses. Let η^{demand} and η^{supply} denote these elasticities, where η^{demand} is defined to incorporate the response of oil production as well as oil inventories to the price shift caused by an exogenous shift in the oil supply curve. In the baseline model, I follow Kilian and Murphy in imposing $-0.8 < \eta^{demand} < 0$ and $0 < \eta^{supply} < 0.0258$. I also follow Kilian and Murphy in imposing dynamic sign restrictions on selected structural impulse response functions. An unexpected oil supply disruption is assumed to raise the real price of oil and to lower global oil production and global real activity for at least twelve months. For details about the estimation of sign-identified VAR models the reader is referred to Kilian and Lütkepohl (2017).

Finally, I impose narrative sign restrictions on the historical decompositions based on extraneous evidence. These restrictions were employed by Kilian and Murphy (2014) to externally validate their preferred model (see also Kilian and Lee (2014)). Whereas Kilian and Murphy (2014) imposed these restrictions by visually inspecting the historical decomposition of the real price of oil for each candidate model, I incorporate the narrative restrictions into the acceptance sampler algorithm much the same way that the dynamic sign restrictions are imposed.[2] This requires me to quantify the informal narrative sign restrictions employed by Kilian and Murphy (2014).

Motivated by the reasoning in Kilian and Murphy (2014, p.460, 469) and Kilian and Lee

[2]For a complementary approach see Antolin-Diaz and Rubio-Ramirez (2018) who impose narrative restrictions on the signs of structural shocks during selected periods and/or on the relative magnitude of the cumulative effect of structural shocks during selected periods.

(2014, p.74), I postulate (1) that storage demand shocks cumulatively raised the log real price of oil by at least 0.2 (or approximately 20%) between May and December 1979, consistent with anecdotal evidence of a dramatic surge in inventory building in the oil market during that time, (2) that storage demand cumulatively lowered the log real price of oil by at least 0.15 between December 1985 and December 1986, after OPEC collapsed, and (3) that storage demand shocks raised the log real price of oil by at least 0.1 cumulatively between June 1990 and October 1990, reflecting market expectations that Iraq would invade its neighbors. Flow supply shocks are assumed to have raised the log real price of oil cumulatively by at least 0.1 between July and October of 1990, reflecting the invasion of Kuwait and the cessation of Iraqi and Kuwaiti oil production in early August. Finally, the cumulative effect of flow demand shocks on the log real price of oil between June and October of 1990 is bounded by 0.1, given that the oil price spike of 1990 was not associated with the global business cycle. These thresholds are chosen to be conservative. For example, few observers would attribute only a 10% increase in the real price of oil in 1990 to storage demand shocks. Likewise, the impact of oil supply shocks on the real price of oil in 1990 is likely to be much higher than 10%. Nevertheless, these weak narrative sign restrictions in conjunction are helpful in narrowing the range of admissible structural models.

3 Inference Based on the Maximum Likelihood Estimator of the Reduced-Form VAR

It is useful to start with the estimates based on the MLE of the reduced-form VAR model, which allows a direct comparison with the estimates plotted in Kilian and Murphy (2014). Figure 1 shows the impulse responses from the set of all admissible structural models conditional on the MLE of the baseline global oil market model described in section 2. All results shown are based on 5 million rotation draws. Because the identifying restrictions are tight and have been further strengthened by the narrative sign restrictions, most estimates are quite precise despite the model being only set identified. The impulse response estimates are very similar to those shown in Kilian and Murphy (2014). For example, the magnitude and pattern of the responses of the real price of oil are virtually unchanged. The only noteworthy difference is the more persistent response of global real activity to flow demand shocks.

Likewise, the historical decomposition for the baseline model in Figure 2 confirms the insights in Kilian and Murphy (2014) about the relative importance of oil supply and oil demand shocks in driving the real price of oil. For example, the surge in the real price of oil in the 2000s is almost entirely explained by the cumulative effects of flow demand shocks. My estimates also support Kilian and Murphy's interpretation of the 1979, 1980, 1986, and 1990 oil price shock episodes. These results show that replacing the originally used index of global real activity by the corrected index for all practical purposes makes no difference for the results.

When Kilian and Murphy derived their oil supply elasticity bound based on historical data for 1990, there were no credible microeconometric estimates of the short-run global price elasticity of oil supply. This question can be revisited now because several recent studies have provided such elasticity estimates, at least at the regional level. The most comprehensive study to date is Newell and Prest (2019), who examine data from all major oil producing regions in the United States, including Texas, Oklahoma, North Dakota, California and Colorado. Newell and Prest's estimate of the one-quarter oil supply elasticity for conventional crude oil is 0.017 with a standard error of 0.006. This estimate is close to the theoretical benchmark of zero derived in Anderson et al. (2018). Based on this evidence, one could make the case for a more conservative upper bound of 0.04 on the one-month price elasticity of oil supply. This bound exceeds Newell and Prest's point estimate by about four standard errors.

Figure 3 illustrates that this change increases the number of admissible models substantially without affecting the substantive conclusions. The impulse responses agree with those for the baseline model in Figure 1. Moreover, even though there is more uncertainty in Figure 4 than in Figure 2 about the relative importance of flow demand and storage demand shocks in the 1980s and during the Asian Crisis of the late 1990s, the central results in Kilian and Murphy (2014) are robust to relaxing the oil supply elasticity bound to match recent extraneous microeconometric evidence. Most importantly, the interpretation of the evolution of the real price of oil in 1990 and during 2003-08 is not affected.

4 Evaluating the Posterior of the Structural Model

None of the results shown thus far, however, account for estimation uncertainty in the reduced-form VAR parameters. In this section, I address this concern by evaluating the posterior distribution of the structural oil market model underlying Figures 3 and 4. My analysis builds on Inoue and Kilian (2013, 2019). Since there is a one-to-one mapping between the joint posterior distribution of the VAR slope parameters, the VAR error covariance matrix, and the rotation matrix, the change-of-variable method may be used to analytically derive the posterior distribution of the set of structural impulse responses associated with each structural model draw from the posterior. This fact allows one to rank these models with the most likely structural model corresponding to the structural model with the highest posterior density value, which provides a unique and economically well-defined estimate. Joint 68% credible sets for impulse responses may be constructed by plotting the impulse responses of the 68% of admissible models with the highest posterior density value. A more detailed discussion of this approach and its advantages compared with alternative approaches can be found in Kilian and Lütkepohl (2017).

In evaluating the posterior, I impose not only the narrative restrictions, but also the added restriction that the dominant autoregressive root is bounded from above by 0.991. This restriction ensures that the effect of a one percent shock at the beginning of the sample on the model data is reduced to nearly zero at the end of the sample. This bound is required for the posterior draws of the historical decomposition to closely resemble the actual historical data for the real price of oil. Without this bound, no meaningful analysis of the cumulative effects of the structural shocks on the real price of oil is possible.

My empirical analysis focuses on the model with the relaxed supply elasticity bound of 0.04. Figure 5 shows that the impulse response functions of the most likely admissible model closely resemble those shown in Figures 1 and 3 based on the MLE. All the qualitative patterns are replicated and the magnitude of the responses of the real price of oil to each of the structural shocks is similar. Even after accounting for estimation uncertainty by constructing joint 68% credible sets, these patterns remain robust.

Figure 6 shows the corresponding estimates of the historical decompositions of the real price of oil. The decomposition based on the most likely admissible model shares many of the key features of the result reported in Kilian and Murphy (2014). For example, flow demand shocks cumulatively accounted for much of the surge in the real price of oil after 2002. This conclusion is robust across virtually all admissible models in the joint credible set. Storage demand shocks, in contrast, were not an important contributor to the surge in the real price of oil between 2003 and mid-2008 nor was there a large contribution from flow supply shocks. The most likely model suggests that flow supply shocks temporarily raised the real price of oil in 1980, following the outbreak of the Iran-Iraq War, but not in late 1979, following the Iranian Revolution. Moreover, both flow supply and storage demand shocks contributed to the oil price spike of 1990. In short, Figures 5 and 6 demonstrate that the conclusions in Kilian and Murphy (2014) are robust both to relaxing the elasticity bound and to allowing for estimation uncertainty.

5 Sensitivity Analysis

An important question is whether the conclusions of Kilian and Murphy (2014) are also robust to extending the estimation period. As Figure 7 shows, the responses implied by the most likely models are similar, when the data are extended to June 2018.[3] Likewise, the variance decomposition in Table 1 is quite robust to extending the estimation period. The only difference is that flow supply shocks play a more important role, reflecting the inclusion of the shale oil boom in recent years. Finally, as Table 2 shows, the historical narrative about the relative importance of oil supply and oil demand shocks remains largely unchanged when the data are extended. For example, the 1979 oil price surge reflects primarily shifts in storage and flow demand. The 1986 drop reflects lower demand rather than higher supply. The 1990 spike was driven by a combination of higher storage demand and lower oil supply, and the surge from late 2002 to mid-2008 was primarily caused by flow demand shocks.

Another interesting question is how robust these results are to replacing the Kilian index in this model by the log-linearly detrended index of global industrial production.[4] As Table 1 shows, even though the modified model in row 3 assigns less importance to flow demand

[3] For more detailed results see Figures A1 and A2 in the not-for-publication appendix.

[4] This index has been discussed in Kilian and Zhou (2018) and Hamilton (2019). The data source is https://econweb.ucsd.edu/ jhamilton/.

shocks than the original specification in row 2, the relative importance of oil demand and oil supply shocks in the variance decomposition is preserved. Moreover, the impulse responses of the respective most likely models in Figure 7 look similar, except for the economically implausible response of oil inventories to a storage demand shock. Nevertheless, the insights of Kilian and Murphy (2014) about the relative importance of oil demand and oil supply shocks between 1979 and 2008 are largely confirmed (see Table 2). There are some differences, of course, since global industrial production is not designed to capture the amplitude or timing of shifts in the flow demand for industrial commodity markets (see Kilian and Zhou (2018)). For example, the modified model assigns more importance to storage demand shocks when there are large flow demand fluctuations in commodity markets that are not fully captured by global industrial production indices (notably in 1979, 1986, 2002-08, and late 2008).

6 Conclusion

This paper established the robustness of the conclusions in Kilian and Murphy (2014) to the use of state-of-the-art methods of evaluating sign-identified structural VAR models. Unlike earlier studies, my analysis explicitly incorporated the narrative restrictions on the historical decomposition of the real price of oil that Kilian and Murphy (2014) informally imposed on the set of admissible structural models. I showed that replacing the original measure of global real economic activity based on Kilian (2009) by the corrected version of this index discussed in Kilian (2019) has no effect on the substantive conclusions. I also demonstrated that the original estimates are robust to relaxing the upper bound on the global price elasticity of oil supply to 0.04, which exceeds the most credible microeconomic estimate in the recent literature by four standard errors. Moreover, the original conclusions of Kilian and Murphy (2014) are supported when using state-of-the-art methods of Bayesian inference for sign-identified VAR models that allow for narrative sign restrictions. Finally, the key conclusions in Kilian and Murphy (2014) are also robust to extending the sample period to June 2018. Replacing the Kilian index by an index of global industrial production does little to change the relative importance of oil demand and oil supply shocks.

References

Anderson, S., R. Kellogg and S. Salant, "Hotelling Under Pressure," *Journal of Political Economy* 126 (2018), 984–1026. https://doi.org/10.1086/697203.

Antolin-Diaz, J. and J. F. Rubio-Ramirez, "Narrative Sign Restrictions for SVARs," *American Economic Review* 108 (2018), 2802–2829. https://doi.org/10.1257/aer.20161852.

Bornstein, G., P. Krusell and S. Rebelo, "Lags, Costs, and Shocks: An Equilibrium Model of the Oil Industry," manuscript, Northwestern University, 2018.

Cross, J., "The Role of Uncertainty in the Market for Crude Oil," manuscript, Norwegian Business School, 2019.

Fattouh, B., L. Kilian and L. Mahadeva, "The Role of Speculation in Oil Markets: What Have We Learned So Far?," *Energy Journal* 34 (2013), 7–33. https://doi.org/10.5547/01956574.34.3.2.

Hamilton, J. D., "Measuring Global Economic Activity," *forthcoming: Journal of Applied Econometrics* (2019), https://doi.org/10.1002/jae.2740.

Herrera, A. M. and S. K. Rangaraju, "The Effect of Oil Supply Shocks on U.S. Economic Activity: What Have We Learned?," *forthcoming: Journal of Applied Econometrics* (2019), https://doi.org/10.1002/jae.2735.

Inoue, A. and L. Kilian, "Inference on Impulse Response Functions in Structural VAR Models," *Journal of Econometrics* 177 (2013), 1–13. https://doi.org/10.1016/j.jeconom.2013.02.009.

———, "Corrigendum to 'Inference on impulse response functions in structural VAR models' [J. of Econometrics 177(2013),1-13]," *Journal of Econometrics* 209 (2019), 139–143. https://doi.org/10.1016/j.jeconom.2017.08.020.

Kilian, L., "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review* 99 (2009), 1053–1069. https://doi.org/10.1257/aer.99.3.1053.

———, "The Impact of the Fracking Boom on Arab Oil Producers," *Energy Journal* 38 (2017), 137–160. https://doi.org/10.5547/01956574.38.6.lkil.

———, "Measuring Global Real Economic Activity: Do Recent Critiques Hold Up to Scrutiny?," *Economics Letters* 178 (2019), 106–110. https://doi.org/10.1016/j.econlet.2019.03.001.

Kilian, L. and T. K. Lee, "Quantifying the speculative component in the real price of oil: The role of global oil inventories," *Journal of International Money and Finance* 42 (2014), 71–87. https://doi.org/10.1016/j.jimonfin.2013.08.005.

Kilian, L. and H. Lütkepohl, "Structural Vector Autoregressive Analysis," Cambridge University Press: New York, NY. https://doi.org/10.1017/9781108164818, 2017.

Kilian, L. and D. P. Murphy, "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics* 29 (2014), 454–478. https://doi.org/10.1002/jae.2322.

Kilian, L. and X. Zhou, "Modeling Fluctuations in the Global Demand for Commodities," *Journal of International Money and Finance* 88 (2018), 54–78. https://doi.org/10.1016/j.jimonfin.2018.07.001.

Newell, R. G. and B. C. Prest, "The Unconventional Oil Supply Boom: Aggregate Price Response from Microdata," *Energy Journal* 40 (2019), 1–3. https://doi.org/10.5547/01956574.40.3.rnew.

Figure 1: Impulse Responses based on MLE ($\eta^{supply} \leq 0.0258$), 1973.2-2009.8

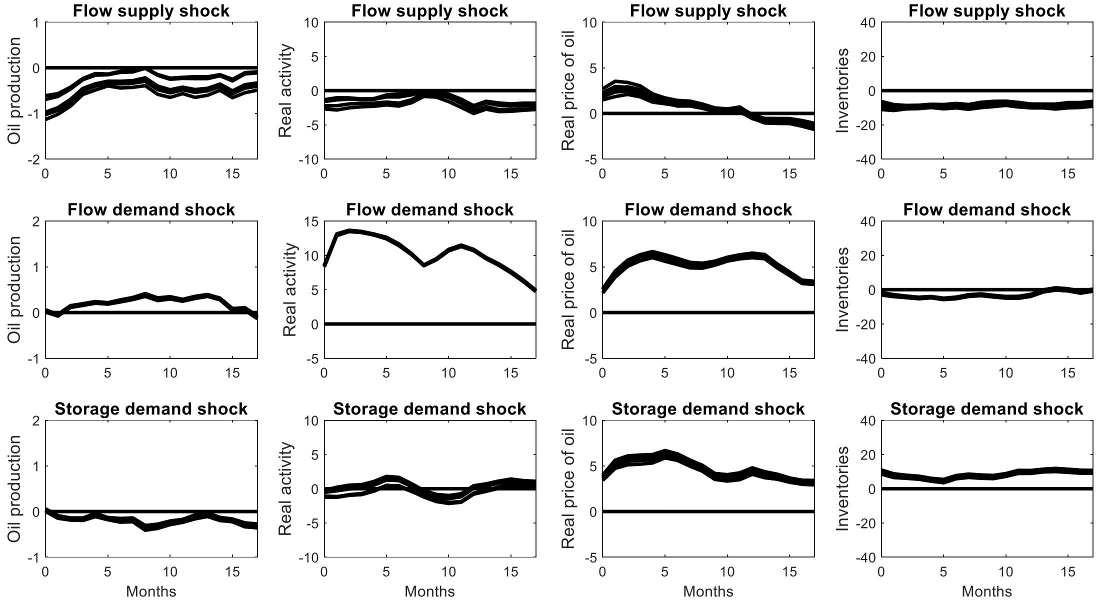

NOTES: The impulse responses shown are obtained from the MLE of the reduced-form VAR model and show the range of model solutions consistent with this estimate based on 5 million draws of the rotation matrix.

Figure 2: Historical Decompositions based on MLE ($\eta^{supply} \leq 0.0258$), 1973.2-2009.8

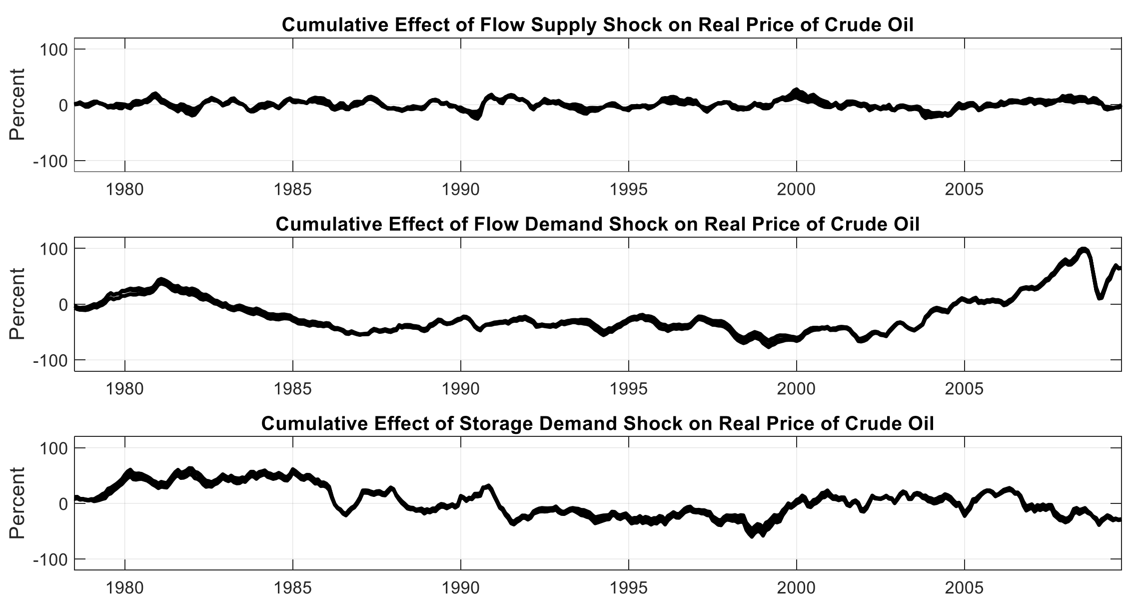

NOTES: The historical decompositions shown are obtained from the MLE of the reduced-form VAR model and show the range of model solutions consistent with this estimate based on 5 million draws of the rotation matrix.

Figure 3: Impulse Responses based on MLE ($\eta^{\text{supply}} \leq 0.04$), 1973.2-2009.8

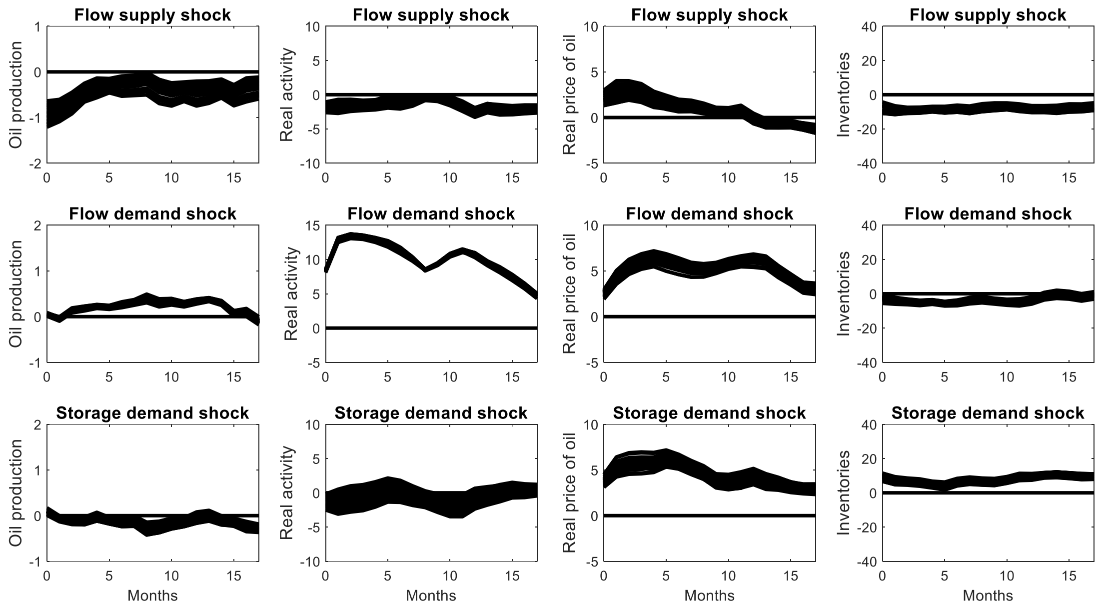

NOTES: The impulse responses shown are obtained from the MLE of the reduced-form VAR model and show the range of model solutions consistent with this estimate based on 5 million draws of the rotation matrix.

Figure 4: Historical Decompositions based on MLE ($\eta^{\text{supply}} \leq 0.04$), 1973.2-2009.8

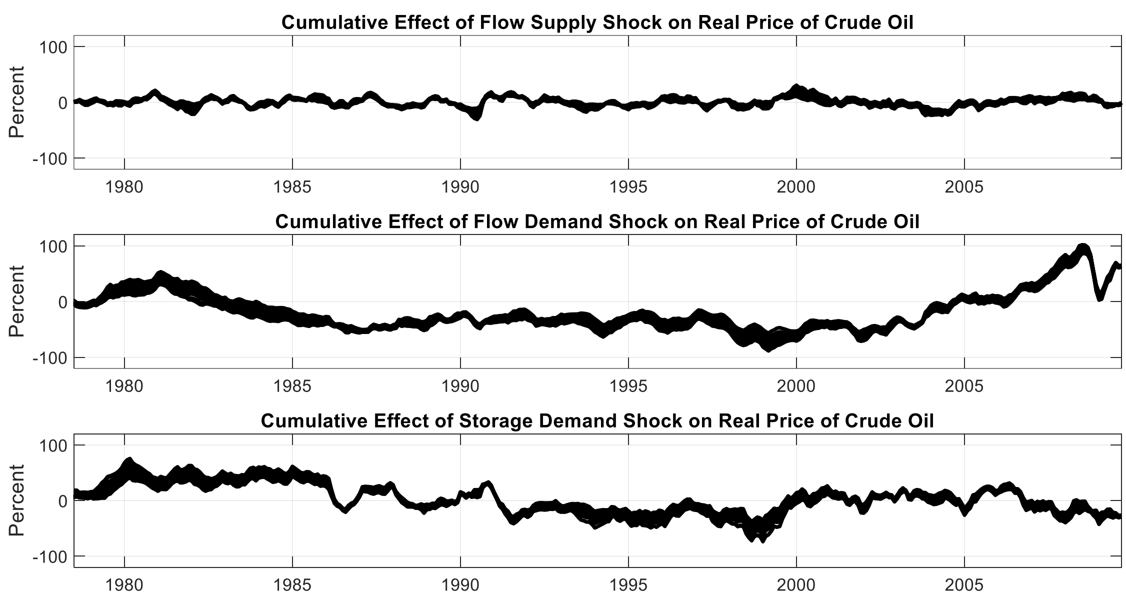

NOTES: The historical decompositions shown are obtained from the MLE of the reduced-form VAR model and show the range of model solutions consistent with this estimate based on 5 million draws of the rotation matrix.

Figure 5: Most Likely Model and 68% Joint Credible Set for Impulse Responses ($\eta^{supply} \leq 0.04$), 1973.2-2009.8

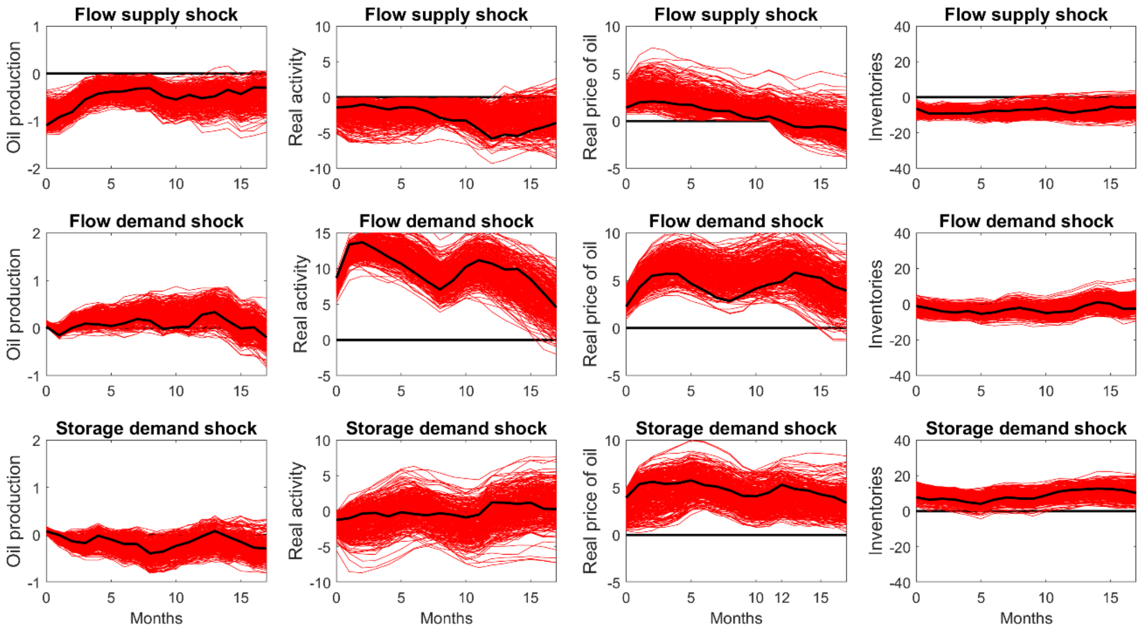

NOTES: The impulse responses shown are from the 68% joint credible set obtained from the posterior distribution of the structural models. The impulse responses implied by the most likely structural model are shown in boldface. Details of the construction of these statistics can be found in Inoue and Kilian (2013, 2019). The results shown are based on 40,000 draws from the reduced-form posterior with 20,000 draws of the rotation matrix each.

Figure 6: Most Likely Model and 68% Joint Credible Set for Historical Decomposition ($\eta^{supply} \leq 0.04$), 1973.2-2009.8

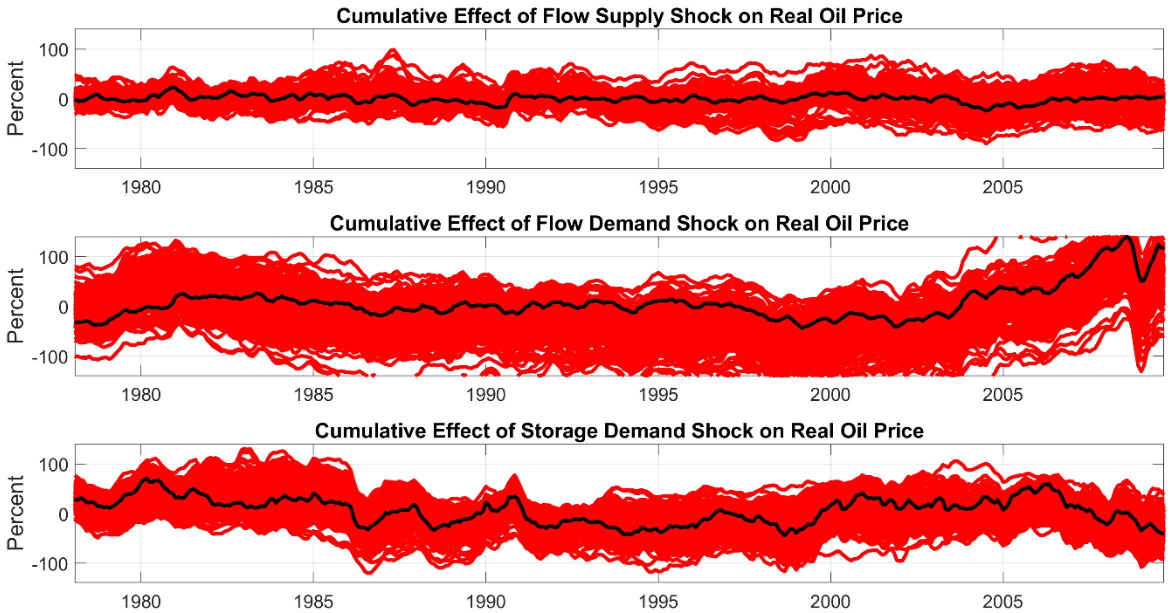

NOTES: The historical decompositions shown are from the 68% joint credible set obtained from the posterior distribution of the structural models, as discussed in Inoue and Kilian (2013, 2019). The cumulative effects implied by the most likely structural model are shown in boldface. The results shown are based on 40,000 draws from the reduced-form posterior with 20,000 draws of the rotation matrix each.

Figure 7: Impulse Response Estimates for the Most Likely Model under Alternative Specifications ($\eta^{supply} \leq 0.04$)

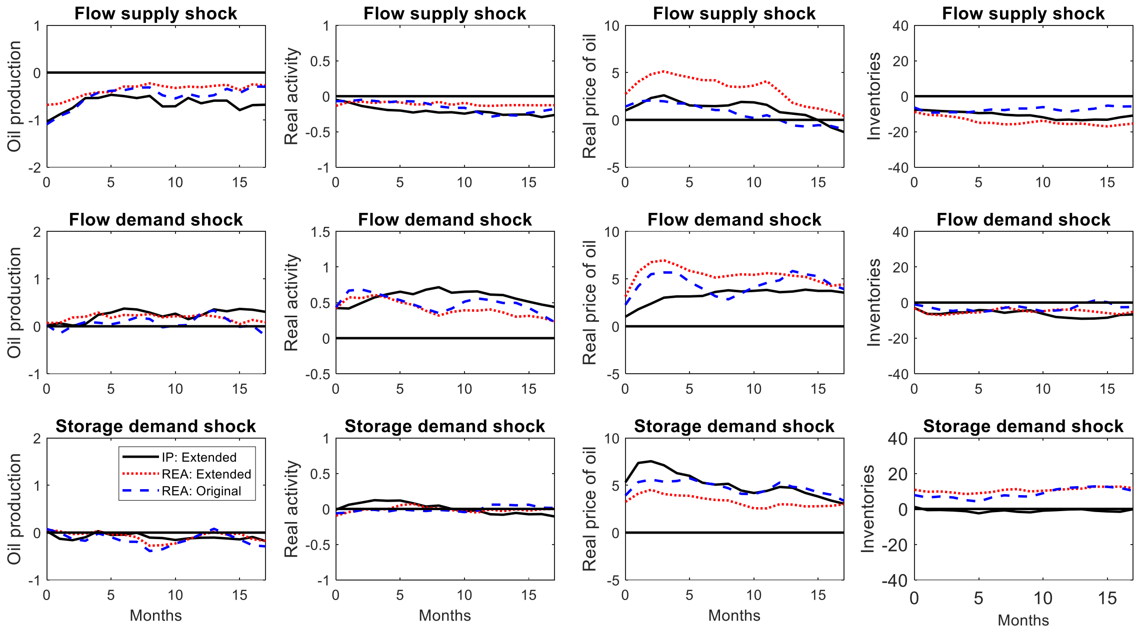

NOTES: The extended sample period is 1973.2-2018.6 The original sample period is 1973.2-2009.8. Details of the construction of these statistics can be found in Inoue and Kilian (2013, 2019). The results shown are based on 40,000 draws from the reduced-form posterior with 20,000 draws of the rotation matrix each. The solid line shows the impulse response estimates for the most likely model when replacing the Kilian index (REA) by a measure of global industrial production (IP). The dotted and dashed lines show the corresponding estimates when using the Kilian index as the proxy for global real economic activity. The responses of the index have been rescaled to match the corresponding response of global industrial production on impact.

Table 1: Variance Decomposition of the Real Price of Oil in the Most Likely Model (%)

		Structural Shock			
Model	Sample	Flow Supply	Flow Demand	Storage Demand	Other Demand
KM with REA	Original	3.6	83.5	11.7	1.2
KM with REA	Extended	23.9	65.0	9.4	1.7
KM with global IP	Extended	26.1	53.5	19.7	0.8

NOTES: The results shown are based on 40,000 draws from the reduced-form posterior with 20,000 draws of the rotation matrix each. All results are based on $\eta^{\text{supply}} \leq 0.04$.

Table 2: Cumulative Effect on the Real Price of Oil in the Most Likely Model (%)

Model	Structural shocks	1979.1-1980.1	1980.9-1980.12	1985.12-1986.12	1990.5-1990.10	2002.7-2008.6	2008.6-2008.12	2014.6-2015.12
KM with global REA (original estimation period)	Flow supply	-2.4	4.2	-1.4	26.9	3.3	0.1	-
	Flow demand	23.7	21.5	-24.7	-0.9	162.6	-85.9	-
	Storage demand	56.5	-17.3	-26.0	21.4	-10.2	-34.9	-
KM with global REA (extended estimation period)	Flow supply	9.8	7.5	-1.8	43.4	11.1	3.7	-64.4
	Flow demand	41.7	8.7	-16.1	7.9	179.9	-112.7	-15.9
	Storage demand	28.2	-8.3	-32.1	11.5	-43.3	-14.9	-8.6
KM with global IP (extended estimation period)	Flow supply	3.7	11.7	-6.2	22.7	25.3	-3.2	-26.3
	Flow demand	2.0	5.5	-2.6	4.1	101.4	-25.7	-27.5
	Storage demand	54.9	-6.7	-56.1	44.4	25.3	-88.0	-44.4

NOTES: The results shown are based on 40,000 draws from the reduced-form posterior with 20,000 draws of the rotation matrix each. All results are based on $\eta^{\text{supply}} \leq 0.04$.

Not-for-Publication Appendix

Figure A1: Updated Estimate of Most Likely Model and 68% Joint Credible Set for Impulse Responses ($\eta^{supply} \leq 0.04$), 1973.2-2018.6

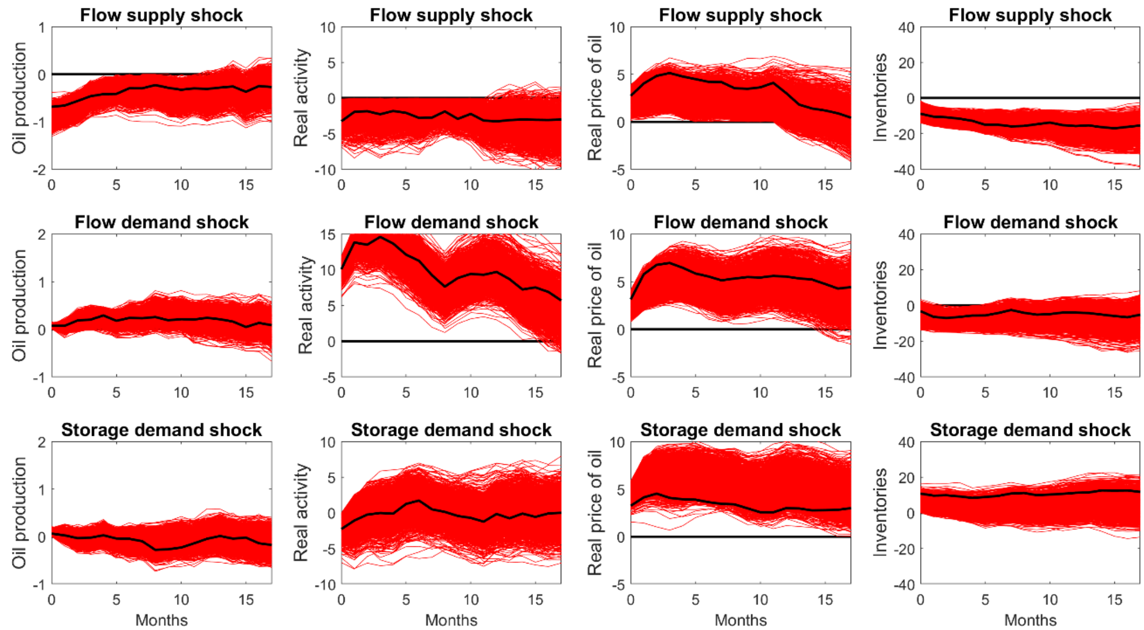

NOTES: The impulse responses shown are from the 68% joint credible set obtained from the posterior distribution of the structural models. The impulse responses implied by the most likely structural model are shown in boldface. Details of the construction of these statistics can be found in Inoue and Kilian (2013, 2019). The results shown are based on 40,000 draws from the reduced-form posterior with 20,000 draws of the rotation matrix each.

Figure A2: Updated Estimate of Most Likely Model and 68% Joint Credible Set for Historical Decomposition ($\eta^{supply} \leq 0.04$), 1973.2-2018.6

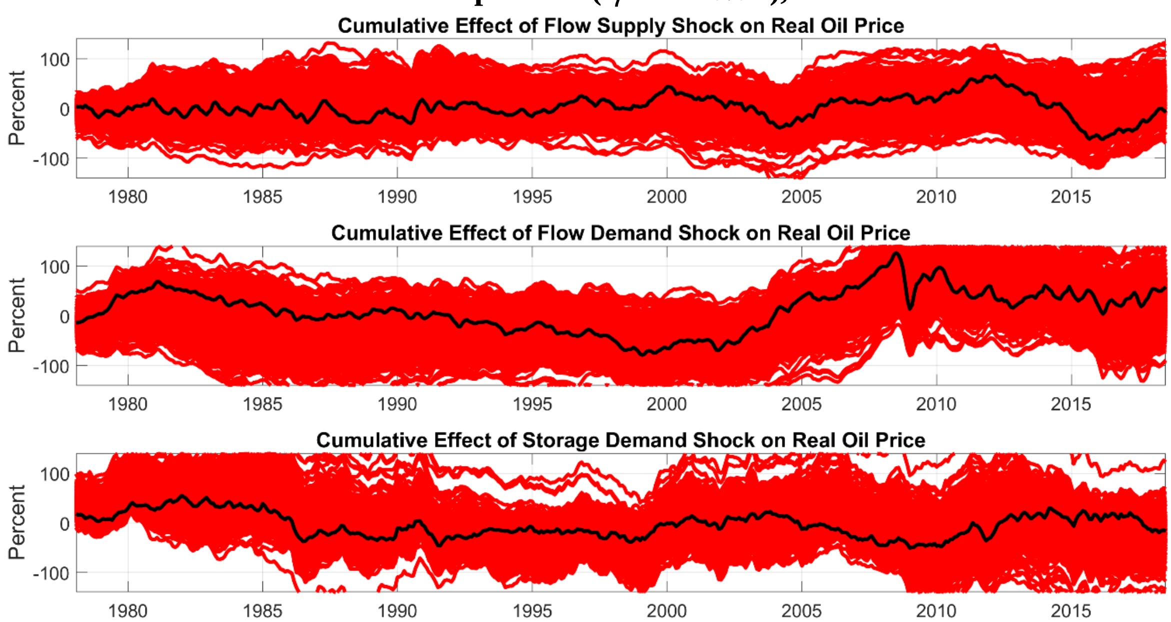

NOTES: The historical decompositions shown are from the 68% joint credible set obtained from the posterior distribution of the structural models. The cumulative effects implied by the most likely structural model are shown in boldface. The results shown are based on 40,000 draws from the reduced-form posterior with 20,000 draws of the rotation matrix each.