Online Appendix for "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 A Time-Varying Parameter VAR: Baumeister and Peersman (2013)

The modeling strategy used by $BP13$ does not allow to identify demand and supply shocks separately. Moreover, focusing only on the oil supply shocks leads to overfitting, which is illustrated by the results reported in Figure A.1. Hence, due to this difference in identification and to the use of time-varying parameters, the results are not directly comparable as they are for the other models.

$BP13$ follow Koop, Pesaran, and Potter (1996) to compute generalized impulse responses, $GIRF$, using Monte Carlo integration. The $GIRF$ is computed as the difference between the expectations, conditional on the information set at time t, ω_t, with and without an exogenous structural shock:

$$GIRF_{t+h} = E[y_{t+h}|\varepsilon_t, \omega_t] - E[y_{t+h}|\omega_t]$$

where y_{t+h} is the forecast of y_t at horizon h, ω_t captures the entire history of y_t up to t, and ε_t denotes the structural innovation at time t .More specifically, one possible state of the economy at time t is drawn from the Gibbs sampler; output, this random draw from the posterior includes time/varying parameters and the elements of the variance-covariance matrix. The transition laws are then used in conjunction with this draw to simulate the future paths of the variance-covariance matrix and the coefficient vector.

Rubio-Ramirez, Waggoner and Zha's (2010) algorithm is then used to compute the time-varying structural impact matrix $B_{0,t}$. Then, the reduced-form residuals are computed as $u_t = B_{0,t}\varepsilon_t$ where the structural shocks ε_t are drawn from a standard normal distribution. The $GIRF$s are computed by comparing the evolution of y_t when the shock is set to $\varepsilon_{i,t} + 1$ and the benchmark where the shock is set to $\varepsilon_{i,t}$. This procedure is repeated until 100 iterations that fulfill the identification restrictions have been found. Then, the responses conditional on the state of the economy are computed as the average over the accepted rotations. To compute the unconditional $GIRFs$, this procedure is repeated by drawing 500 current states of the economy at each point in time. The pointwise median for each variable at each date is reported as the representative impulse response function.

The top left panel of Figure A.1 illustrates the posterior median for the impact response of the real oil price. (Recall that impact on $BP13$ means a quarter.) The oil supply shock is normalized to a 1% decrease in world oil supply. As noted by $BP13$, the oil price increase triggered by a 1% drop in oil production is greater in the 1990s and 2000s than in the 1970s and 1980s. The impact response of oil price increases substantially over time, with minimum of 0.86 % in 1977:IV and a maximum of 28.75% in 2008:IV. For comparison with previous models, we compute the monthly impact response of real oil price by dividing the quarterly estimate by 3, we find that real oil price fluctuates between a minimum of 0.27% in 1977:IV and a maximum of 9.58% in 2008:IV. (This maximum is well below that of $LN12$.) Extending the sample to 2016:IV reveals a steep decline in the response between the onset of the Great Recession and the mid-2010s.

The bottom panel of Figure A.2 depicts the short-run price elasticity of demand implied by the TVP-BVAR model. The decline in the elasticity of demand from the earlier part of the sample relative to the 2000s is consistent with $BP13$'s findings. Extending the sample indicates a slight increase in the elasticity between the onset of the Great Recession and 2013:I that coincides with the reduced responsiveness of real oil prices.

The quarterly median impact response of the real oil price over the 1974:I-2016:IV sample period is 6.75% and the mean impact response is 7.03%. These numbers correspond to a monthly median (2.25%) and mean (2.34%) that fall in the same ballpark as $KM14$ and $BH18$. Yet, they are considerable larger than the modal model estimates in $K09$ or $KM12$.[1] Furthermore, we find a lower degree of precision in the estimates over the last two decades.

2 Computation of Monthly Structural Supply Shocks

Consider the recursive VAR of Kilian (2009), which has the reduced-form representation

$$y_t = c + \sum_{i=1}^{p} A_i y_{t-i} + u_t. \tag{1}$$

Let $A = \begin{bmatrix} c & A_1 & ... & A_p \end{bmatrix}'$ and the structural-form VAR representation be given by

$$B_0 y_t = \sum_{i=1}^{24} B_i y_{t-i} + \varepsilon_t \tag{2}$$

[1]See Figure A.1 in the on-line appendix for the median impact response of the real oil price and associated 68% and 95% credible sets.

so that B_0^{-1} is the structural impact multiplier matrix. After estimating the reduced-form VAR parameters, $\widehat{B}_0^{-1}$ may be obtained via Cholesky decomposition of the estimated variance-covariance matrix $\Omega = E\left[u_t u_t'\right]$ where $u_t = B_0^{-1}\varepsilon_t$. The series of monthly structural supply shocks (ε_{1t}) thus corresponds to the first column of the $(3 \times T - p)$ matrix $B_0 u_t$.

Recovering the series of relevant supply shocks from sign-identified models is less straightforward. To compute the response of U.S. real GDP to the quarterly oil supply measure in $KM12$ and $KM14$, we proceed in the following manner. Let $\overline{P}$ be a potential solution for the impact multiplier matrix so that $\overline{P} = PU$ where P is the lower triangular Cholesky decomposition of Ω (i.e., $\Omega = PP'$) and U is an orthonormal matrix drawn from a uniform prior distribution that is independent from the distribution of the reduced-form VAR parameters. Following Inoue and Kilian (2013), we compute the implied structural impulse responses, $\widetilde{\theta}$, by first taking a random draw (A, Ω) from the posterior of the reduced- form VAR parameters. We then consider N random draws of the rotation U, and compute the implied structural impulse responses, $\widetilde{\theta}$, for each triplet (A, Ω, U). If $\widetilde{\theta}$ satisfies the sign restrictions, we store the value of $\overline{P}$ and the value of the posterior density of $\widetilde{\theta}$. The series of structural supply shocks for any admissible model is given by the first column of the matrix of structural disturbances $\varepsilon_t = \overline{P}^{-1} u_t$. Recovering the series of structural supply shocks for all admissible models allows us to compute not only the response for the modal model but also the $(1 - \alpha)\%$ pointwise HPD credible sets.

3 MIDAS regressions

We compute the impulse responses for U.S. real GDP using MIDAS regressions and local projections (Jordá 2005) in a manner similar to Ferrara and Guérin (2018). The MIDAS regressions are given by

$$\Delta y_{t+h} = \mu_h + \beta_h B\left(L^{1/m}; \theta_h\right) \widehat{\varepsilon}_{st}^{i} + \eta_{t+h} \tag{3}$$

where $B\left(L^{1/m}; \theta_h\right) = \sum_{j=1}^{Q} b\left(j; \theta_h\right) L^{(j-1)/m}$ and $L^{s/m}\widehat{\varepsilon}_{st}^{i(m)} = \widehat{\varepsilon}_{t-s/m}^{i(m)}$ and $i = K09$, $KM12$, $KM14$, $BH18$. Here y_t refers to the quarterly (lower frequency) of the GDP and m refers to the monthly frequency of the structural supply shocks $\widehat{\varepsilon}_{st}^{i(m)}$. μ_h denotes a constant term, and η_{t+h} denotes the regression error. The subscript h in equation (3) indicates that the parameters change with the horizon h as the local projections constitute a series of regressions for each horizon h. The MIDAS exponential almon lag polynomial is given by

$$b\left(j;\theta_h\right)=\frac{\exp\left(\theta_h j\right)}{\sum_{j=1}^{Q}\exp\left(\theta_h j\right)}$$

where Q the lags of the high-frequency variable (the structural supply shock) equals 24 for $K09$, $KM12$, $KM14$ and 12 for $BH18$.[2] θ_h governs the shape of the weight function.

Note that for the sign identified models ($KM12$, $KM14$) we compute the MIDAS regression for each series of structural shocks derived from the monthly SVAR admissible models. We report both the pointwise median response and the response for the modal model. For $BH18$ we compute the MIDAS regressions for each of the 500,000 draws and then compute the pointwise median response. Figure A.5 reports the impulse responses for U.S. real GDP to a 1% decline in world crude oil production. The green lines depict the median pointwise response, whereas the blue lines illustrate the response for the modal model in $KM12$ and $KM14$. For ease of comparison, the figure also plots the responses obtained from projecting the real GDP growth on the quarterly measures of oil supply shocks reported in Figure A.5.

We note that the main conclusions derived in section 5 are robust to using an estimation method that takes into account the mixed-frequency nature of the structural shocks and U.S. real GDP. More specifically, the impulse response estimated via the MIDAS regressions are very similar to those obtained without accounting the mixed-frequency nature of the data for $KM09$ and $KM12$. The only noticeable differences are that the MIDAS regressions suggest a less persistent and smaller impact of oil supply shocks for $KM14$ and a more persistent and pronounced response for $BH18$.

[2] The choice of lag length is consistent with the number of lags included in the original SVAR specifications. Increasing the number of lags to 24 in $BH18$ leads to a more persistence and pronounced response of U.S. real GDP.

Table A.1. Impact sign restrictions

	Oil supply	Aggregate demand	Oil-market specific demand	U.S. demand	RoW demand	U.S. supply	RoW supply
	Panel A: Kilian and Murphy (2012) - $KM12$						
Oil production	-	+	+				
Real price of oil	+	+	+				
Real activity	-	+	-				
	Panel B: Kilian and Murphy (2014) - $KM14$						
Oil production	-	+	+				
Real price of oil	+	+	+				
Real activity	-	+	-				
Inventories			+				
	Panel C: Lippi and Nobili (2012) - $LN12$						
Oil production	-			+	+		
Real price of oil	+			-	+	+	+
US output	-			+	-	+	
RoW output	-			-	+		+
US real output price				+	-	-	+

Table A.2. Priors for A in Baumeister and Hamilton (2018)

	Distribution	Parameter values
Short-run price elasticity of oil supply, α_{qp}	Student $t\left(c_{qp}^{\alpha}, \sigma_{qp}^{\alpha}, \upsilon_{qp}^{\alpha}\right)$	$c_{qp}^{\alpha}=0.1, \sigma_{qp}^{\alpha}=0.2, \upsilon_{qp}^{\alpha}=3$
Short-run price elasticity of oil demand, β_{qp}	Student $t\left(c_{qp}^{\beta}, \sigma_{qp}^{\beta}, \upsilon_{qp}^{\beta}\right)$	$c_{qp}^{\beta}=-0.1, \sigma_{qp}^{\beta}=0.2, \upsilon_{qp}^{\beta}=3$
Contemporaneous effect of oil prices on economic activity, α_{yp}	Student $t\left(c_{yp}^{\alpha}, \sigma_{yp}^{\alpha}, \upsilon_{yp}^{\alpha}\right)$	$c_{yp}^{\alpha}=-0.05, \sigma_{yp}^{\alpha}=0.1, \upsilon_{yp}^{\alpha}=3$
Fraction of world oil inventories held by OECD countries, χ	Beta $\left(\alpha_{\chi}, \beta_{\chi}\right)$	$\alpha_{\chi}=15, \beta_{\chi}=10$

Table A.3. Estimated Short-Run Price Elasticities - Original Sample Period

	Estimates Using Original Sample Period				
	$K09$	$KM12$	$KM14$	$BH18$	$LN12$
Short-run price elasticity of supply	0	0.01	0.01	0.14	0.444
		(0.00, 0.02)	(0.00, 0.02)	(0.05, 0.35)	(0.07, 9.20)
Short-run price elasticity of demand	-4.500	-2.17	-0.35	-0.36	-0.08
in production		(-7.54, -1.12)	(-0.79, -0.11)	(-0.84, -0.16)	(-1.50, -0.00)
Short-run price elasticity of demand			-0.20		
in use			(-0.65, -0.02)		
Impact response of real oil price to	0.024	0.47	2.87	2.86	12.23
supply shock	(-0.02, 0.09)	(0.14, 0.90)	(1.64, 4.62)	(1.70, 3.25)	(0.82, 24.15)
Response of real oil price 1-year	-0.03	0.56	1.14	4.05	4.67
after the shock	(-0.14, 0.12)	(-0.31, 1.93)	(-0.72, 3.65)	(2.15, 5.94)	(-10.02, 18.86)

Notes: This table reports estimates from (or implied by) the original studies. We report the magnitudes implied by $K09$; the posterior median and the 95% HPD for $KM12$ and $KM14$; and the posterior median and the 95% credible sets (in parenthesis) for $LN12$ and $BH18$.

Table A.4. Alternative estimates of B_0^{-1}

Variable	$K09$			$KM12$		
q	1.472	0	0	1.328	0.005	0.080
y	0.024	7.211	0	0.035	5.553	-3.436
p	-0.005	1.050	6.401	-0.771	3.563	4.180

Notes: As is conventional, the first column has signs opposite to the impact response displayed in the figures as the latter depict the response to an unanticipated decline in oil supply.

Table A.5.
Pointwise posterior median estimates of B_0^{-1} for $BH18$

Panel A: Benchmark

	Oil Supply shock	Economic activity shock	Oil consumption demand shock	Oil inventory demand shock
Variable				
q	0.695	0.171	0.243	0.409
y	0.004	0.997	-0.003	-0.006
p	-2.020	1.141	1.570	2.588
Δi	-0.021	-0.064	-0.092	0.845

Panel B: Support for the prior on α_{qp} $(0, 0.0258]$

	Oil Supply shock	Economic activity shock	Oil consumption demand shock	Oil inventory demand shock
Variable				
q	0.980	0.016	0.019	0.032
y	0.001	0.998	-0.001	-0.002
p	-0.915	0.803	0.955	1.559
Δi	-0.042	0.051	0.060	1.099

Panel C: Support for the prior on α_{qp} $(0, 0.04]$

	Oil Supply shock	Economic activity shock	Oil consumption demand shock	Oil inventory demand shock
Variable				
q	0.963	0.030	0.036	0.060
y	0.002	0.998	-0.002	-0.003
p	-1.077	0.927	1.099	1.809
Δi	-0.040	0.463	0.055	1.091

Panel D: Support for the prior on α_{qp} $(0, 0.10]$

	Oil Supply shock	Economic activity shock	Oil consumption demand shock	Oil inventory demand shock
Variable				
q	0.865	0.099	0.122	0.202
y	0.004	0.996	-0.004	-0.006
p	-1.624	1.238	1.486	2.456
Δi	-0.033	-0.008	-0.009	0.983

Table A.6. Prior and Posterior probabilities that the impact of a specific structural shock on the indicated variable is positive for $BH18$

Panel A: Benchmark

	Oil Supply shock		Economic activity shock		Oil consumption demand shock		Oil inventory demand shock	
Variable	Prior	Posterior	Prior	Posterior	Prior	Posterior	Prior	Posterior
q	0.914	1.000	0.970	1.000	0.970	1.000	0.970	1.000
y	0.860	1.000	0.999	1.000	0.029	0.000	0.029	0.000
p	0.139	0.000	0.970	1.000	0.970	1.000	0.970	1.000
Δi	0.699	0.203	0.233	0.159	0.233	0.159	0.970	1.000

Panel B: Support for the prior on α_{qp} $(0, 0.0258]$

	Oil Supply shock		Economic activity shock		Oil consumption demand shock		Oil inventory demand shock	
Variable	Prior	Posterior	Prior	Posterior	Prior	Posterior	Prior	Posterior
q	0.919	1.000	0.971	1.000	0.971	1.000	0.971	1.000
y	0.860	1.000	0.999	1.000	0.028	0.000	0.028	0.000
p	0.139	0.000	0.971	1.000	0.971	1.000	0.971	1.000
Δi	0.701	0.098	0.230	0.902	0.230	0.902	0.971	1.000

Panel C: Support for the prior on α_{qp} $(0, 0.04]$

	Oil Supply shock		Economic activity shock		Oil consumption demand shock		Oil inventory demand shock	
Variable	Prior	Posterior	Prior	Posterior	Prior	Posterior	Prior	Posterior
q	0.918	1.000	0.971	1.000	0.971	1.000	0.971	1.000
y	0.861	1.000	0.999	1.000	0.028	0.000	0.028	0.000
p	0.138	0.000	0.971	1.000	0.971	1.000	0.971	1.000
Δi	0.705	0.101	0.228	0.847	0.228	0.847	0.971	1.000

Panel D: Support for the prior on α_{qp} $(0, 0.10]$

	Oil Supply shock		Economic activity shock		Oil consumption demand shock		Oil inventory demand shock	
Variable	Prior	Posterior	Prior	Posterior	Prior	Posterior	Prior	Posterior
q	0.919	1.000	0.972	1.000	0.972	1.000	0.972	1.000
y	0.861	1.000	1.000	1.000	0.028	0.000	0.028	0.000
p	0.139	0.000	0.971	1.000	0.972	1.000	0.972	1.000
Δi	0.702	0.131	0.230	0.437	0.230	0.437	0.972	1.000

Table A.7. Short-run Run Price Elasticities - $BP13$ Model		
	Original Sample (1974:I -2011:I)	Sample Period (1973:I-2016:IV)
Short-run price elasticity of supply		
Short-run price elasticity of demand in production	-0.141(**-0.291**)	-0.148 (**-0.285**)
Short-run price elasticity of demand in use		
Impact response of real oil price to supply shock	2.352 (**2.518**)	2.250 (**2.343**)
	(1.651, 3.639)	(1.142, 4.484)
Response of US real GDP 1-year after the shock	-0.125 (**-0.154**)	-0.155 (**-0.174**)
	(-0.268, -0.016)	(-1.036, 0.513)

Notes: The table reports the posterior median responses for the sample period and mean estimate in bold for Baumeister and Peersman (2013), $BP13$, model. We report 95% posterior credible sets. We also report the quarterly median and mean estimate (in bold) for short-run price elasticity of demand in production.

Figure A.1: Quarterly Response of Real Price of Oil and Real GDP to an Unexpected Decline in World Oil Supply-BP(2013)

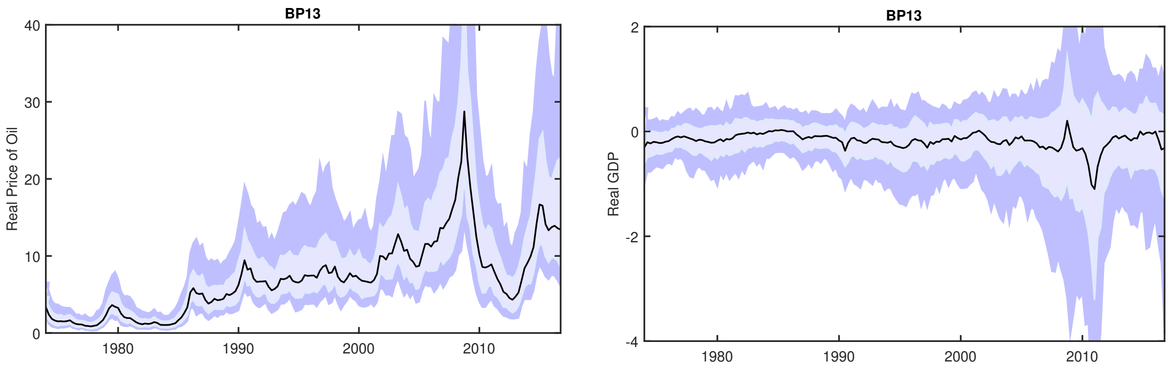

Notes: The figure plots the median impact responses (black line) of real price of crude oil and real GDP to oil supply shock normalized to a 1 percent decrease in world oil production. The light and dark shaded areas indicate respectively 68% and 95% posterior credible sets.

Figure A.2: Quarterly Responses of Real Price of Oil and Real GDP to an Unexpected Decline in World Oil Supply and Oil Demand Elasiticity-BP(2013)

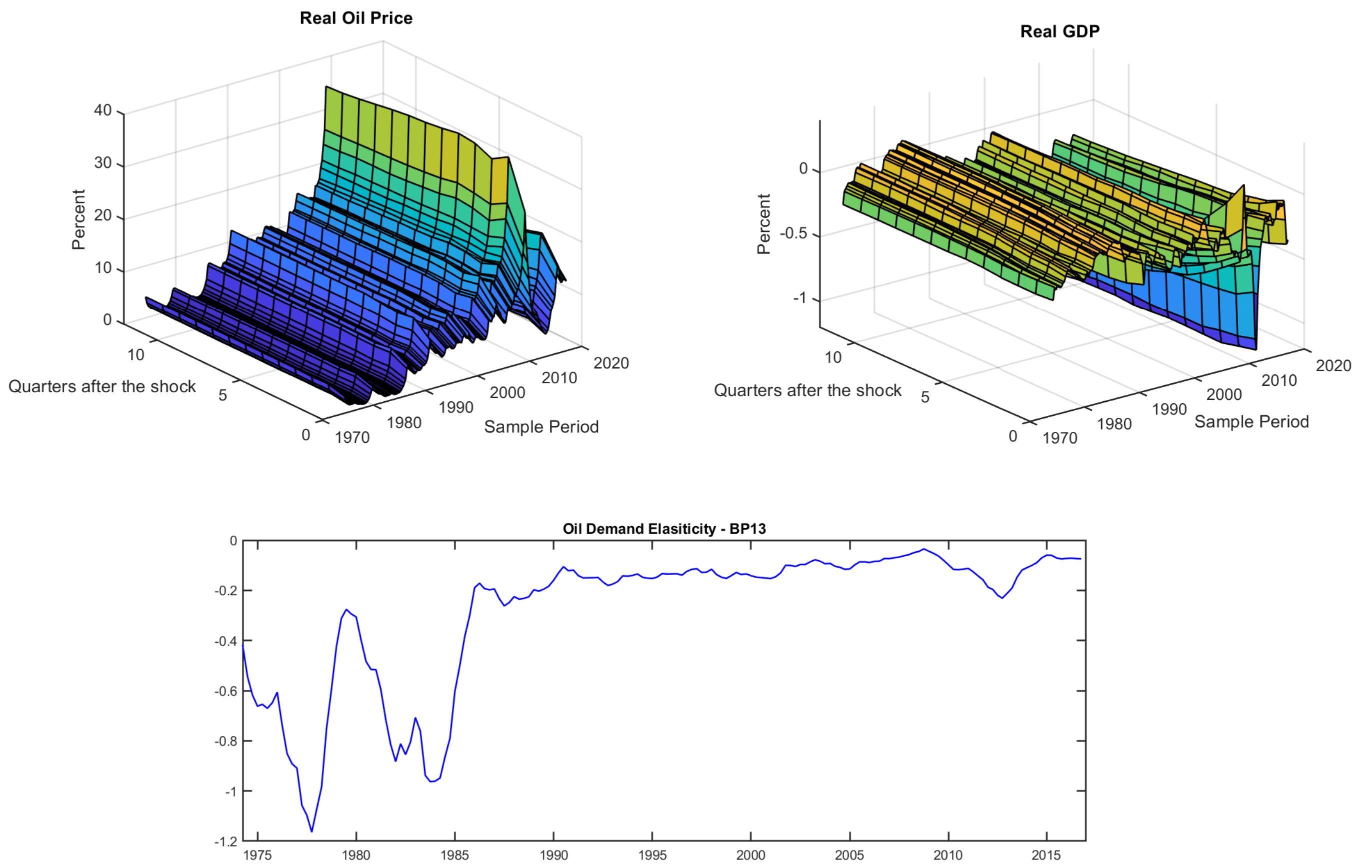

Figure A.3: Monthly Responses of Oil Production, Real Price of Oil and Quarterly Response of Real GDP to an 1% Unexpected Decline in World Oil Supply

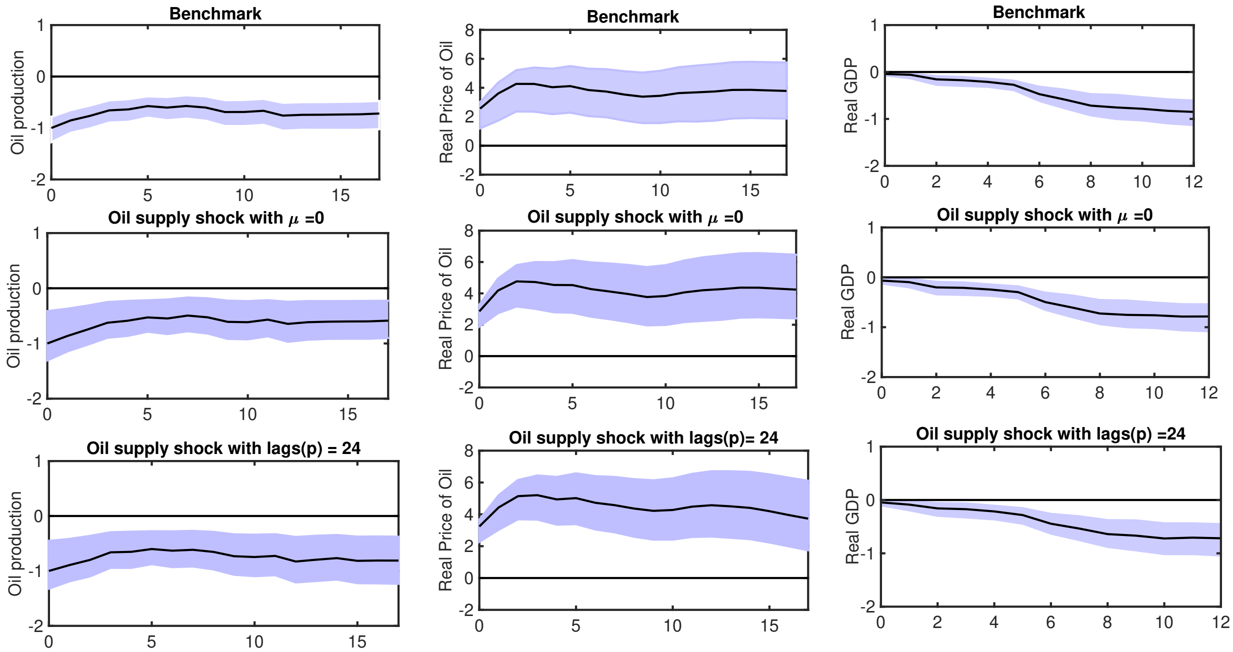

Notes: The figure plots response to an oil supply disruption normalized to represent the effect of an unexpected 1% monthly decline in world oil production for *BH18* with different values of μ. *BH18* benchmark analysis sets μ=0.5, which regards observations in the first sample i.e., data prior to 1975 as only half as informative as those in the second sample i.e., data in post 1975. We run the benchmark specification with $\mu = 0$, which completly discards the first sample. The black line denotes Bayesian posterior median responses and shaded regions denote 95% posterior credible sets. We plot the quarterly response of Real GDP to quarterly 1% supply shock.

Figure A.4: Monthly Responses of Oil Production, Real Price of Oil and Quarterly Response of Real GDP using alternative measures of economic activity

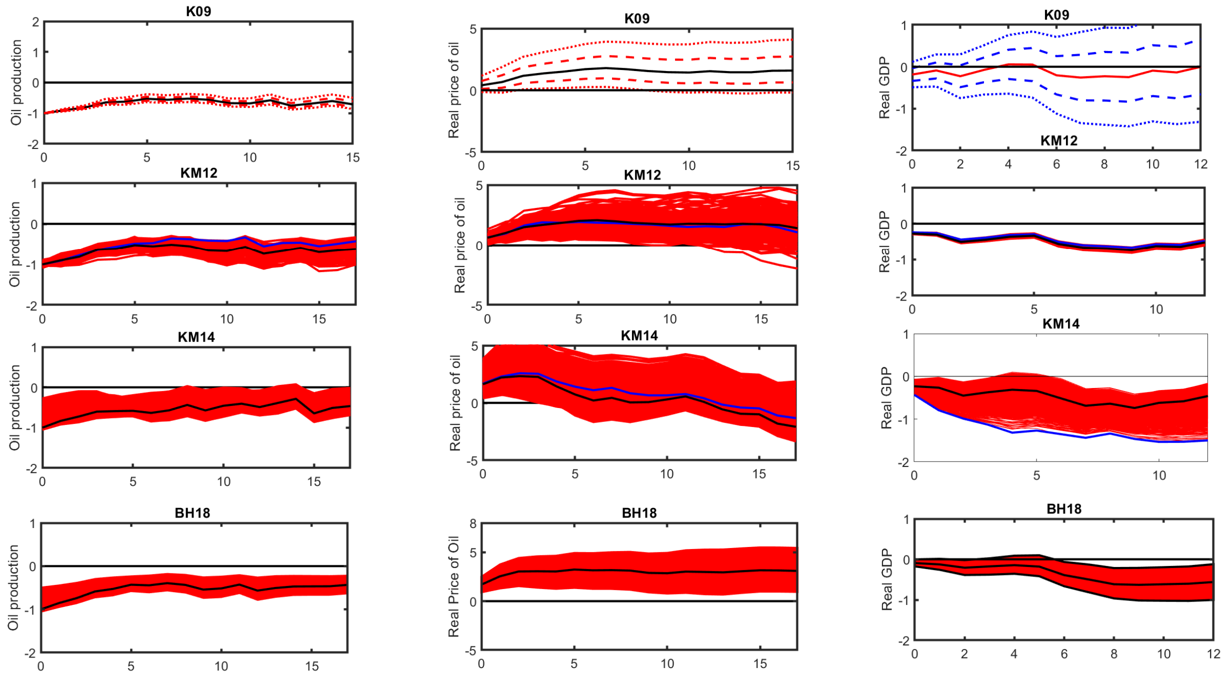

Notes: The figure plots response to an oil supply disruption normalized to represent the effect of an unexpected 1% monthly decline in world oil production for each of the SVAR specifications. We use linearly detrended World IP as a measure of economic acitivity in *K09*, *KM12*, *KM14* and Kilian index of real economic activity in *BH18* model. We report monthly response of real price of oil. We plot the quarterly response of Real GDP to quarterly 1% supply shock. 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*. The black line denotes Bayesian posterior median responses and shaded regions denote 95% posterior credible sets for *BH18*.

Figure A.5: Response of Real GDP using MIDAS model

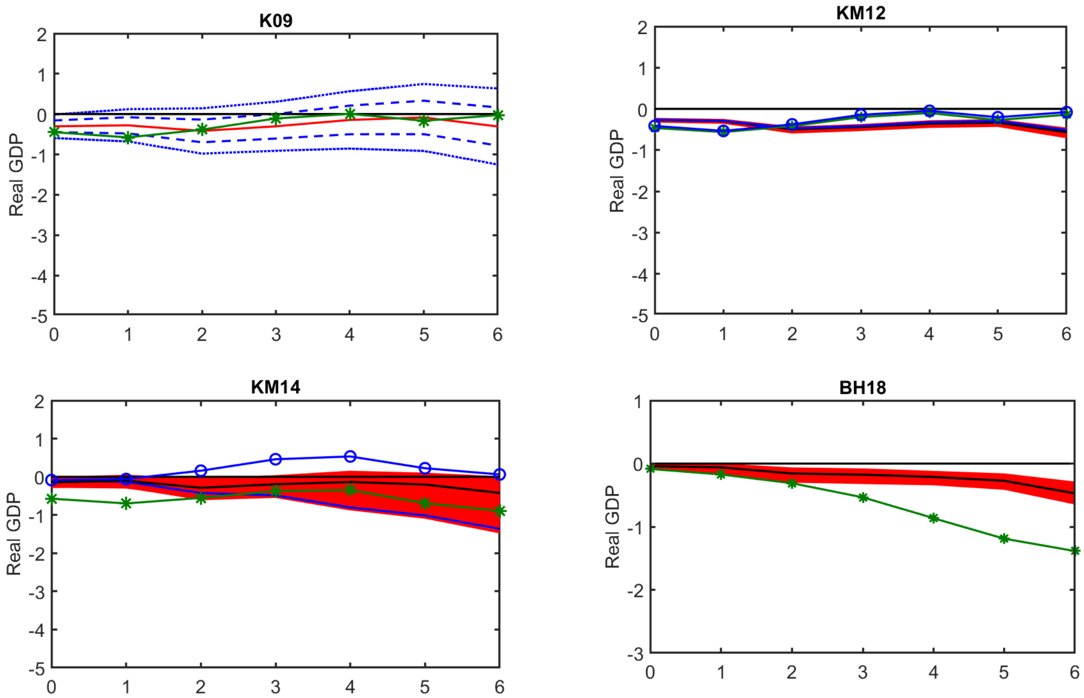

Notes: We report impulse response of Real GDP to an oil supply shock: MIDAS model. The green line with asterisks are the estimates for *K09* using MIDAS model. The baseline point estimates for *K09* are reported with one and two standard error bands that are constructed using a recursive-design wild bootstrap. The green line with asterisks and blue line with circles are the posterior median and modal model response for*KM12*, *KM14* using MIDAS model. The black and blue line denotes the baseline posterior median and modal model response for *KM12*, *KM14*. The green line with asterisks are the estimates for *BH18* model using MIDAS model. The black line denotes Bayesian posterior median responses and shaded regions denote 95% posterior credible sets for *BH18*.