What Drives Oil Price Fluctuations? Revisiting the Role of Oil Supply and Demand Shocks

Christiane Baumeister, University of Notre Dame
James D. Hamilton, UCSD

8th Atlantic Workshop on Energy and Environmental Economics
June 21-22, 2018

What explains the recent decline in oil prices?

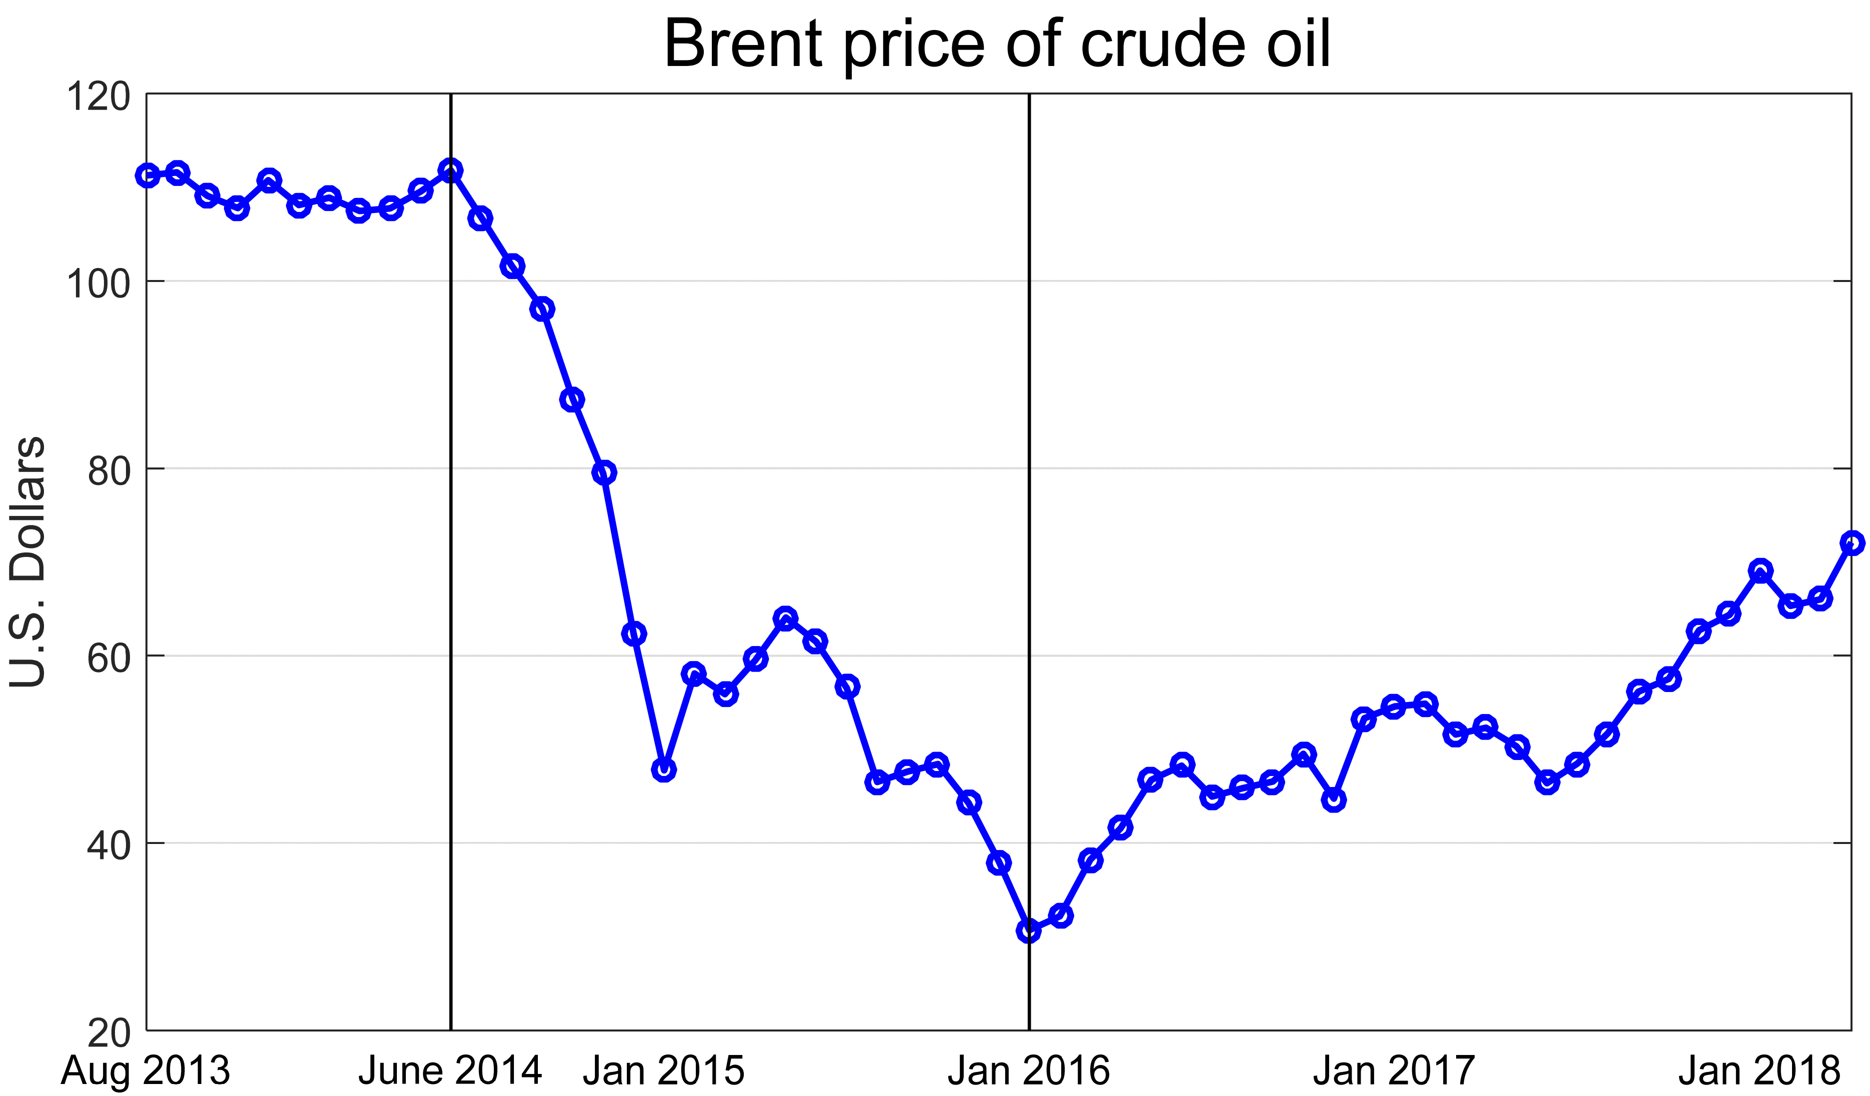

How should we go about it?

- Causal interpretation of correlations requires information about economic structure

- Express what we know as a Bayesian prior

- Formal Bayesian methods permit less dogmatic and more flexible structural inference

Structural model of interest:

$$\underset{(n\times n)}{\mathbf{A}}\ \underset{(n\times 1)}{\mathbf{y}_t} = \boldsymbol{\lambda} + \mathbf{B}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{B}_m \mathbf{y}_{t-m} + \mathbf{u}_t$$

$\mathbf{u}_t \sim$ i.i.d. $N(\mathbf{0}, \mathbf{D})$

$\mathbf{D}$ diagonal

Example: oil supply and demand

$$q_t = k^s + \alpha^s p_t + b_{11}^s p_{t-1} + b_{12}^s q_{t-1} + b_{21}^s p_{t-2} + b_{22}^s q_{t-2} + \cdots + b_{m1}^s p_{t-m} + b_{m2}^s q_{t-m} + u_t^s$$

$$q_t = k^d + \beta^d p_t + b_{11}^d p_{t-1} + b_{12}^d q_{t-1} + b_{21}^d p_{t-2} + b_{22}^d q_{t-2} + \cdots + b_{m1}^d p_{t-m} + b_{m2}^d q_{t-m} + u_t^d$$

$$\mathbf{A} = \begin{bmatrix} 1 & -\alpha^s \\ 1 & -\beta^d \end{bmatrix}$$

4 structural parameters:

$\alpha^s, \beta^d, \text{Var}(u_t^d), \text{Var}(u_t^s)$

3 VAR parameters:

$\text{Var}(\varepsilon_t^p), \text{Var}(\varepsilon_t^q), \text{Cov}(\varepsilon_t^p, \varepsilon_t^q)$

Model is unidentified if all we observe is quantity and price.

Our approach:

represent imperfect information about (α^s, β^d) in the form of a Bayesian prior distribution $p(\alpha^s, \beta^d)$

How does this relate to other approaches?

(1) Traditional identified structural VAR is a special case of a Bayesian prior that is dogmatic.

For example: Cholesky identification
We know with absolute certainty that short-run supply elasticity α^s is 0
and have no information at all about short-run demand elasticity β^d.

The global oil market

Kilian AER (2009)

q_t = world oil production

y_t = real global economic activity

p_t = real price of oil

Structural model of the global oil market

oil supply:

$$q_t = \alpha_{qy} y_t + \alpha_{qp} p_t + \mathbf{b}_1' \mathbf{x}_{t-1} + u_{1t}$$

economic activity:

$$y_t = \alpha_{yq} q_t + \alpha_{yp} p_t + \mathbf{b}_2' \mathbf{x}_{t-1} + u_{2t}$$

inverse of oil demand curve:

$$p_t = \alpha_{pq} q_t + \alpha_{py} y_t + \mathbf{b}_3' \mathbf{x}_{t-1} + u_{3t}$$

Note: $\alpha_{pq} =$ inverse of short-run price-elasticity of oil demand

What does Cholesky identification imply?

$$\alpha_{qy} = \alpha_{qp} = \alpha_{yp} = 0$$

oil supply:

$$q_t = \alpha_{qy} y_t + \alpha_{qp} p_t + \mathbf{b}_1' \mathbf{x}_{t-1} + u_{1t}$$

economic activity:

$$y_t = \alpha_{yq} q_t + \alpha_{yp} p_t + \mathbf{b}_2' \mathbf{x}_{t-1} + u_{2t}$$

inverse of oil demand curve:

$$p_t = \alpha_{pq} q_t + \alpha_{py} y_t + \mathbf{b}_3' \mathbf{x}_{t-1} + u_{3t}$$

Special case of Bayesian prior beliefs

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 \\ -\alpha_{yq} & 1 & 0 \\ -\alpha_{pq} & -\alpha_{py} & 1 \end{bmatrix}$$

$$p(\mathbf{A}) = p(\alpha_{yq})p(\alpha_{pq})p(\alpha_{py})$$

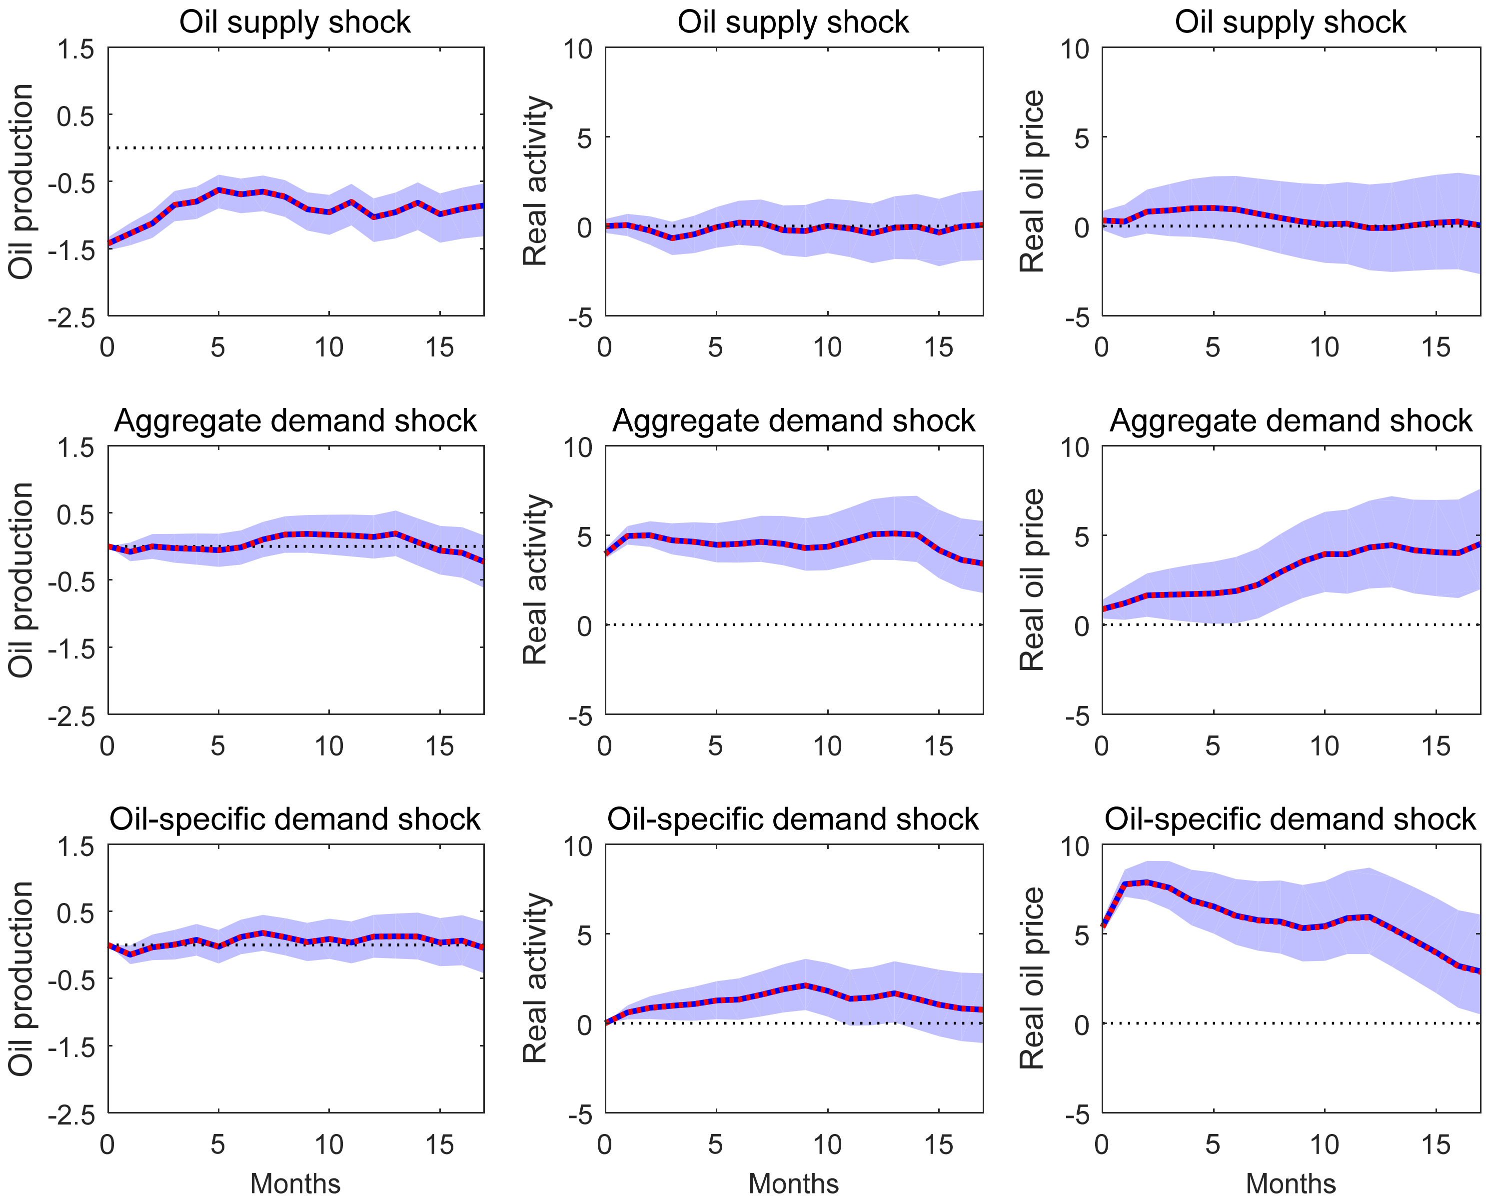

Blue: posterior median IRF as calculated using Baumeister-Hamilton algorithm for dogmatic prior.

Red: IRF calculated using Kilian's (AER 2009) Cholesky analysis.

Prior (red) and posterior (blue) distributions for unknown elements of **A**

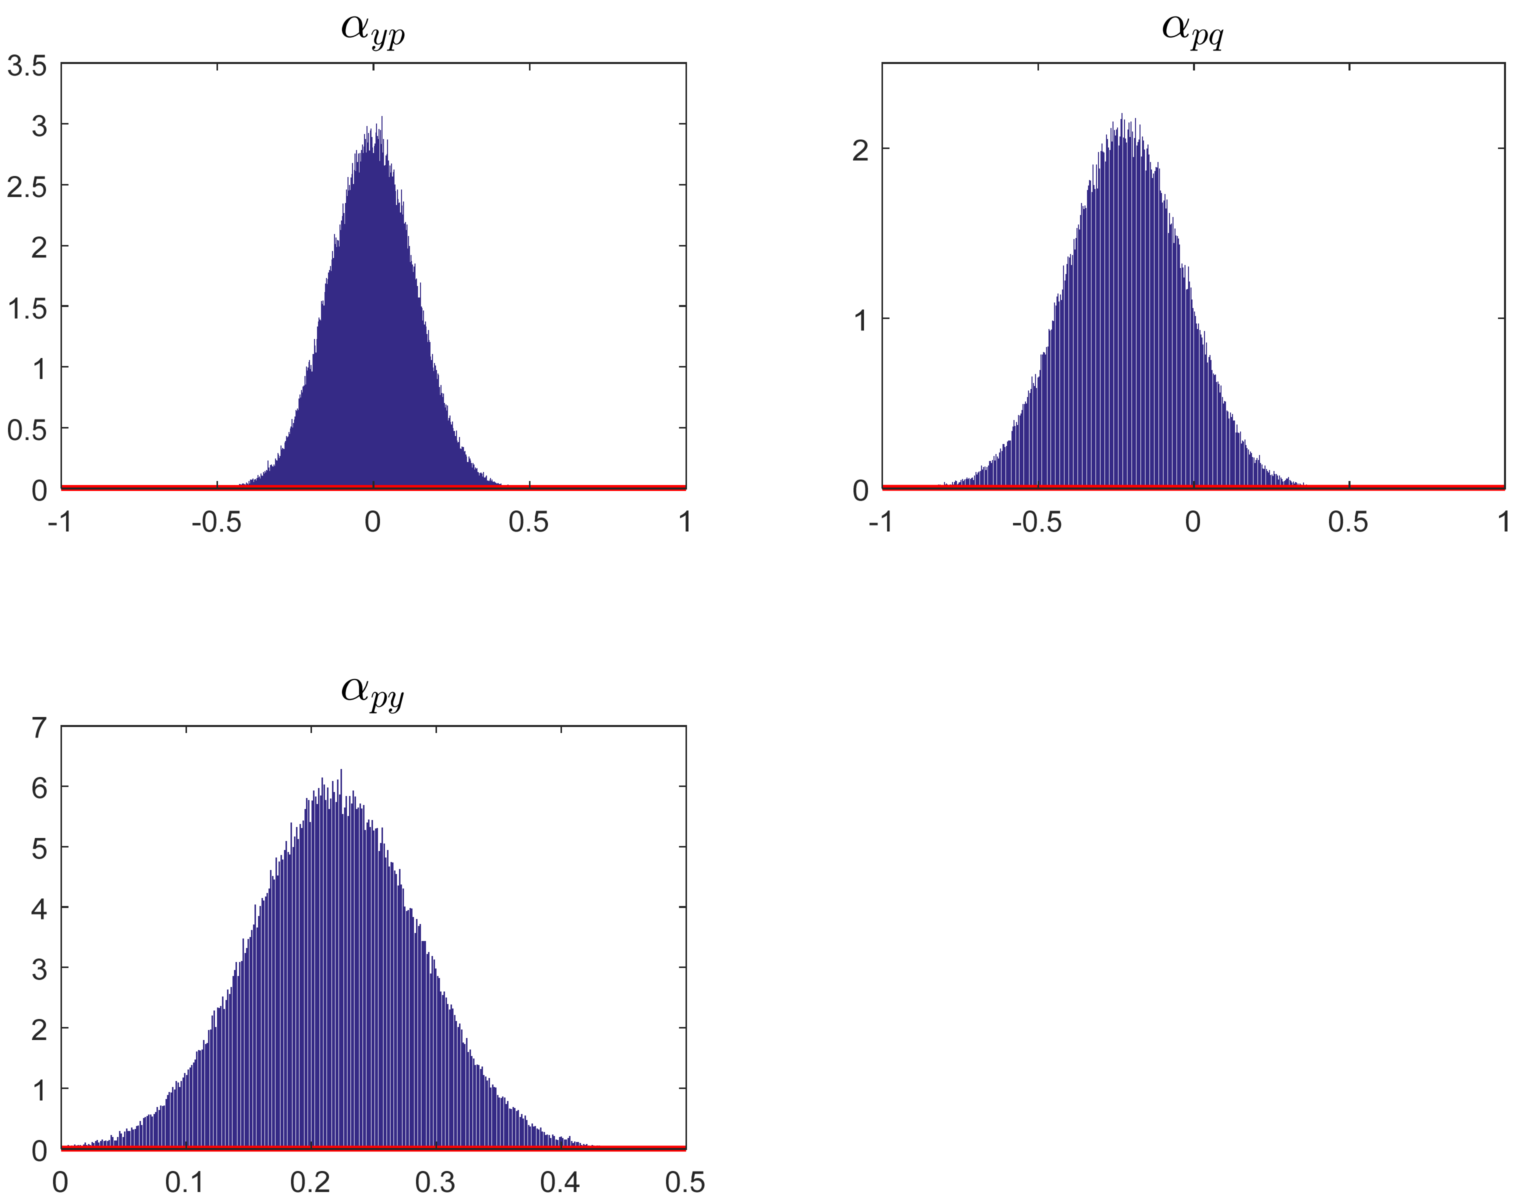

Prior (red) and posterior (blue) distributions for unknown elements of **A**

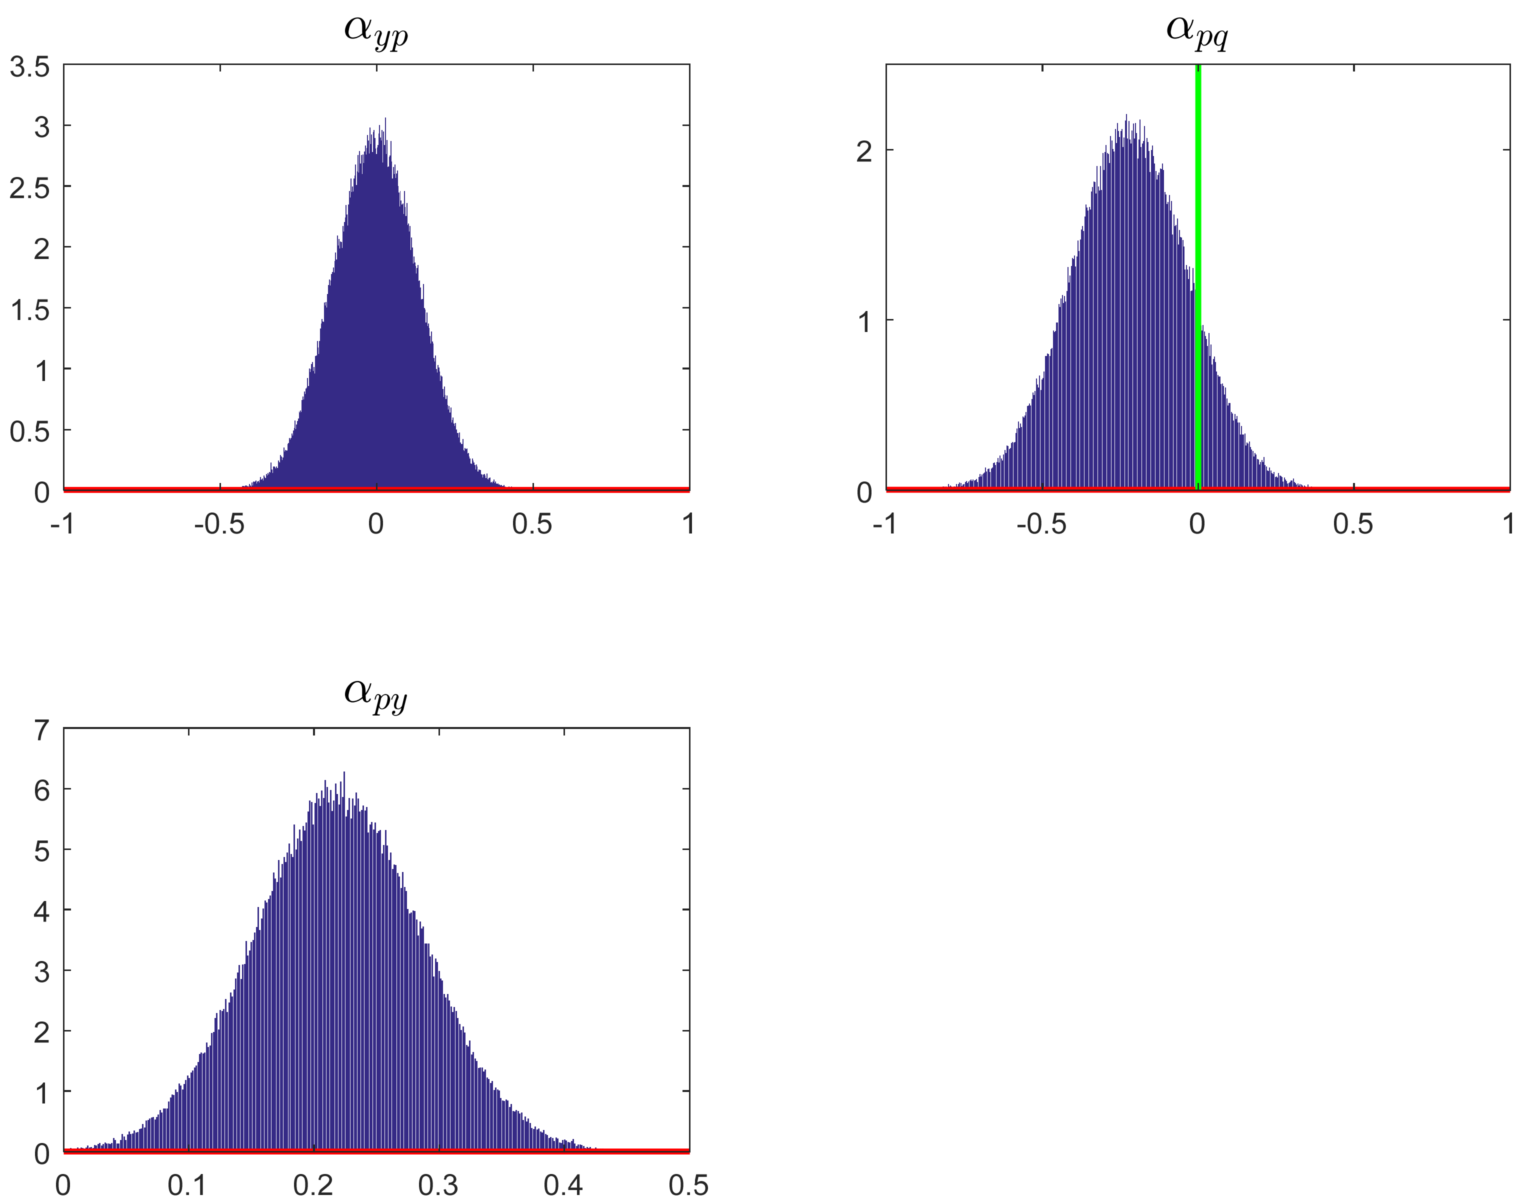

Posterior density of short-run oil demand elasticity

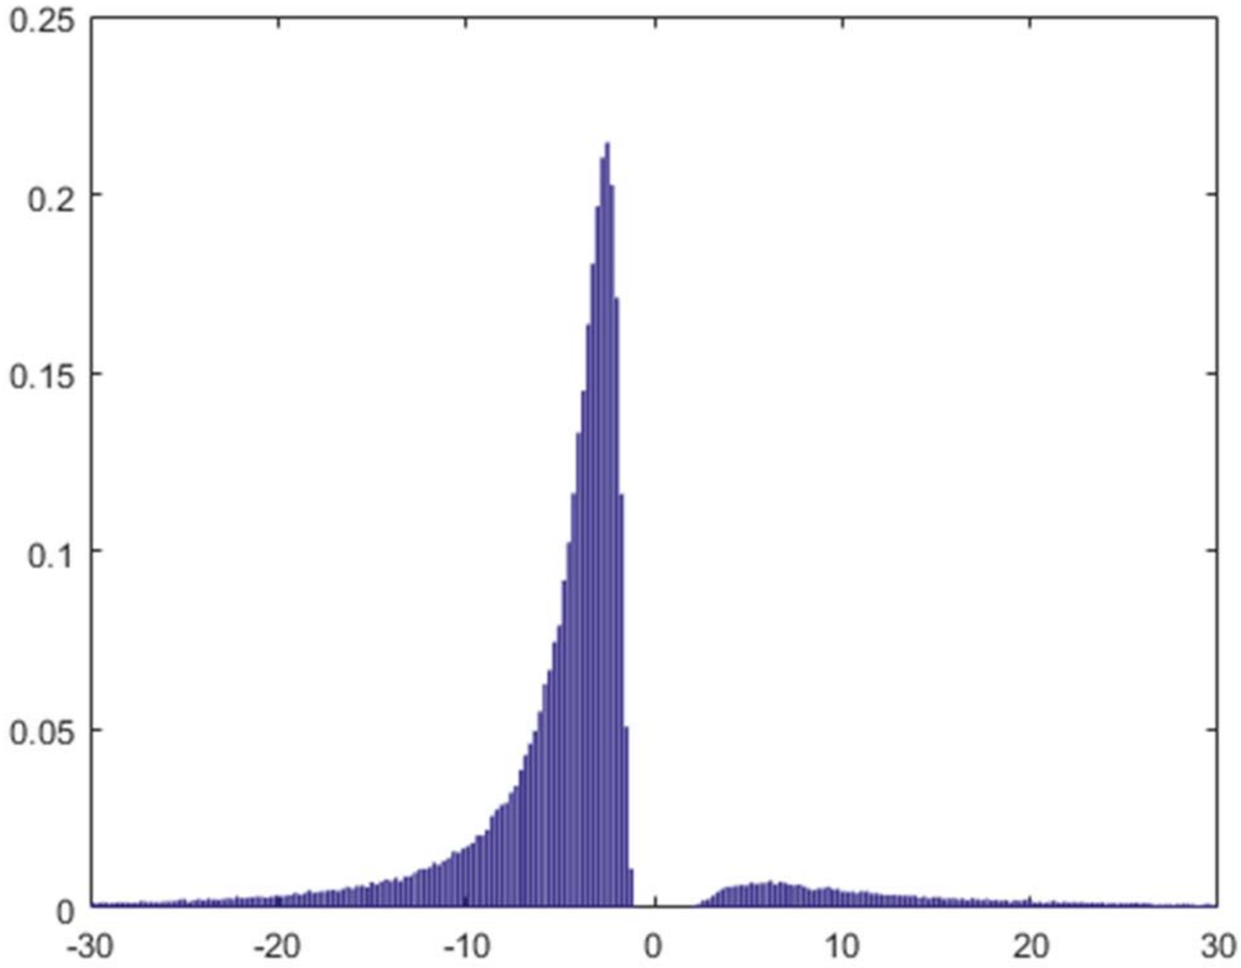

12% posterior probability that demand elasticity > 0
94% posterior probability that abs(elasticity) > 2

How does this relate to other approaches?

(2) Sign restrictions:

- Idea: $\alpha^s \geq 0,\ \beta^d \leq 0$

How Does This Relate to Other Approaches?

(2) Sign restrictions:

- Idea: $\alpha^s \geq 0,\ \beta^d \leq 0$

- Question:
  What would happen if we ran the standard "sign-restricted structural VAR" *without* imposing any restrictions?

Dynamic effect of 1% increase in oil price on oil production (median and 68% credible set)

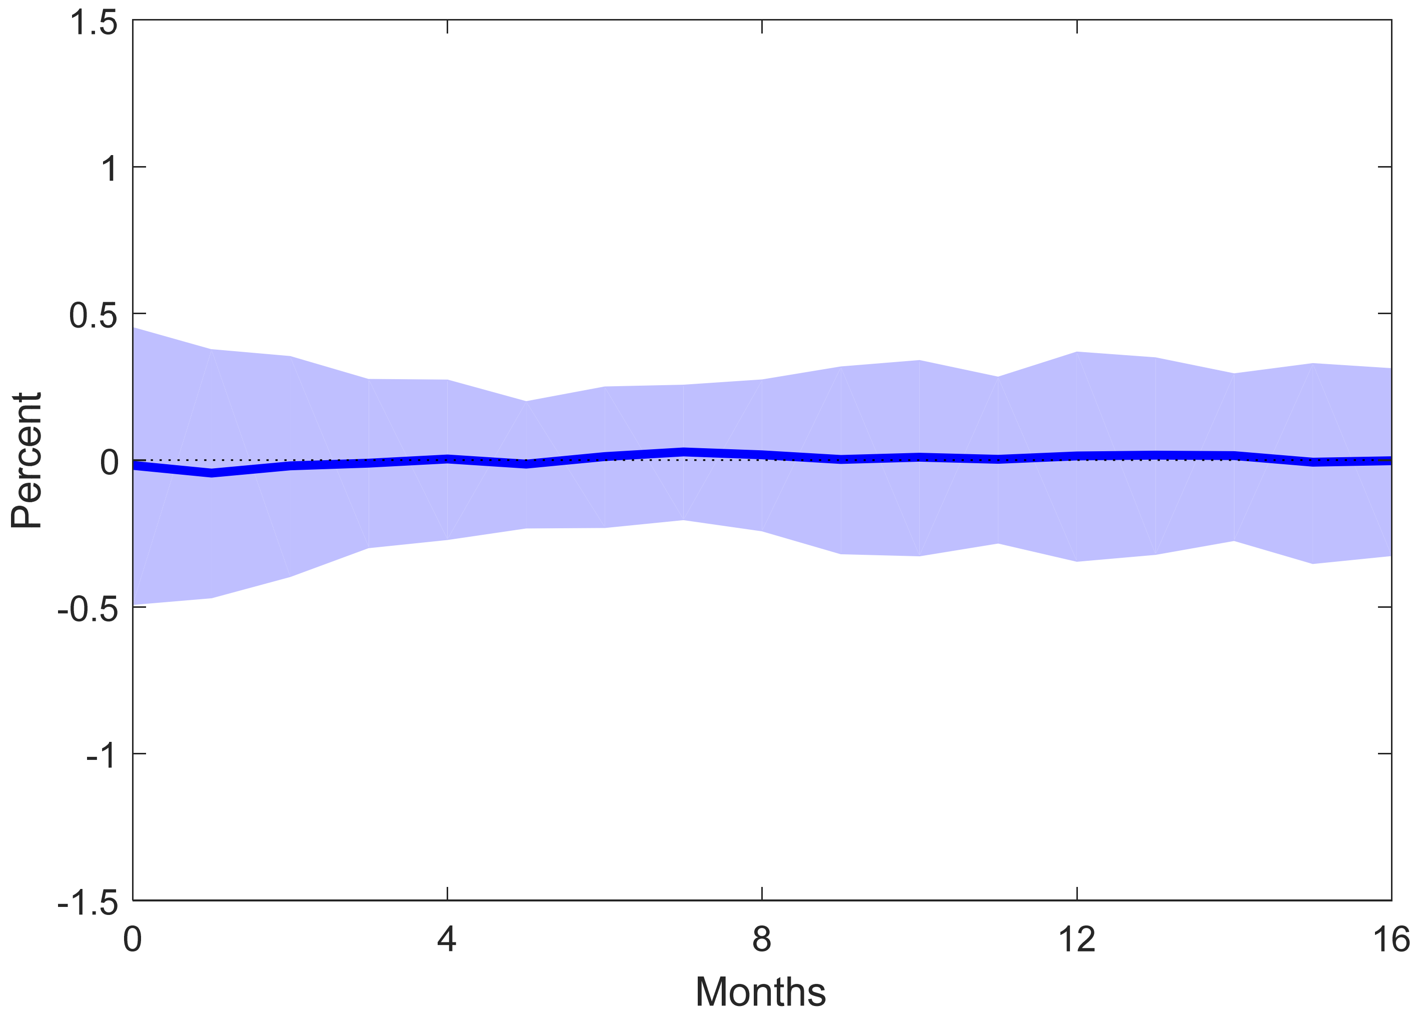

Distribution for horizon *h* = *0*

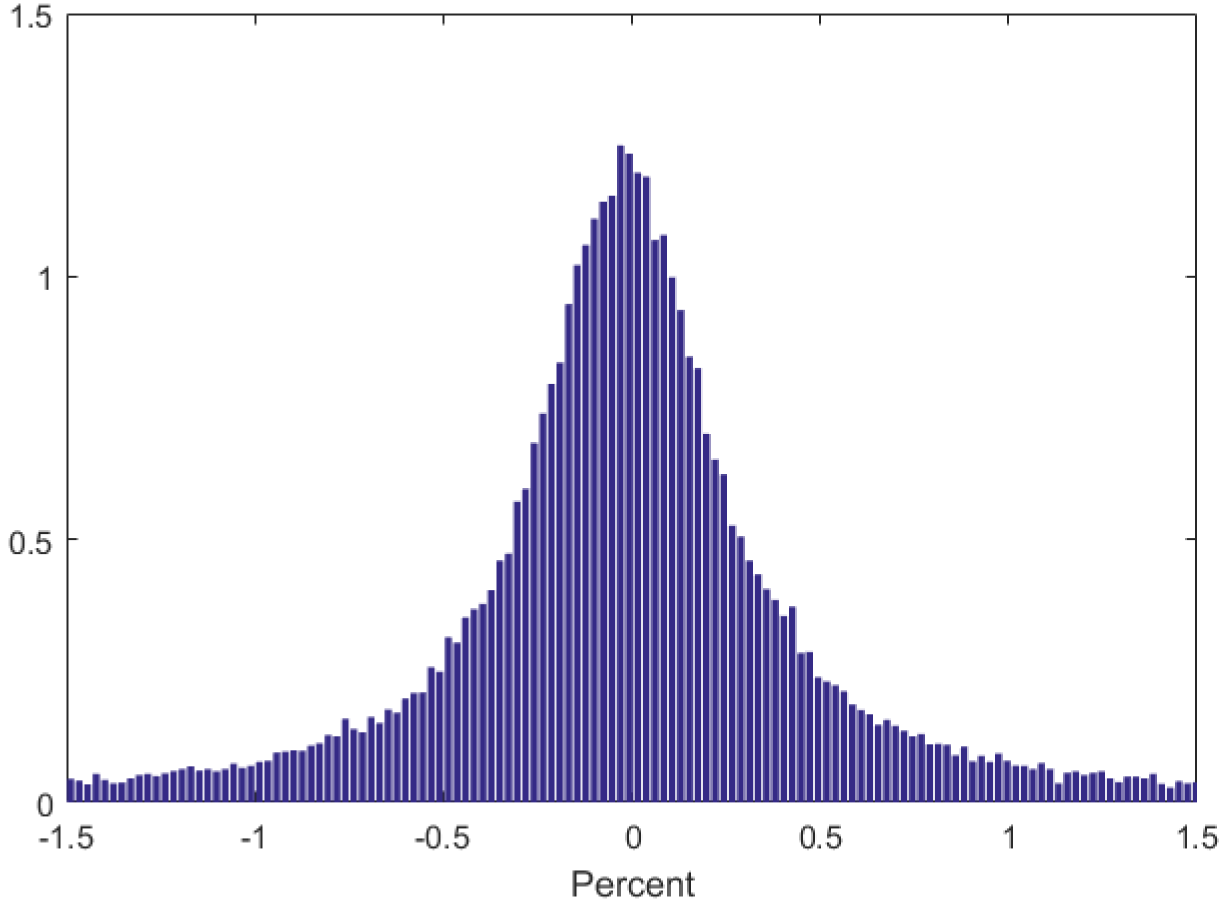

Distribution for horizon *h = 0*

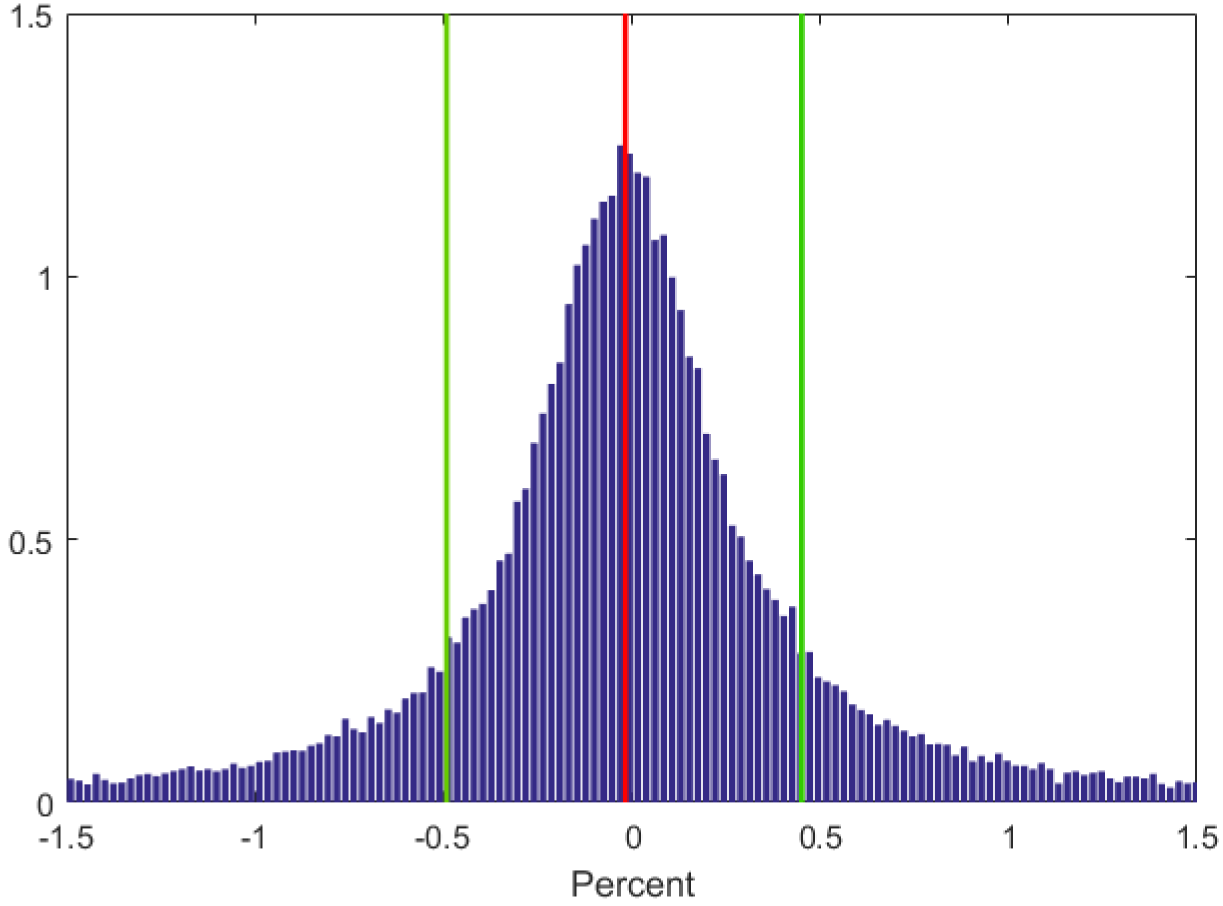

Red line: median

Green lines: 16^{th} and 84^{th} percentiles

Implied elasticity of oil supply

(with restriction that $\alpha^s \geq 0$)

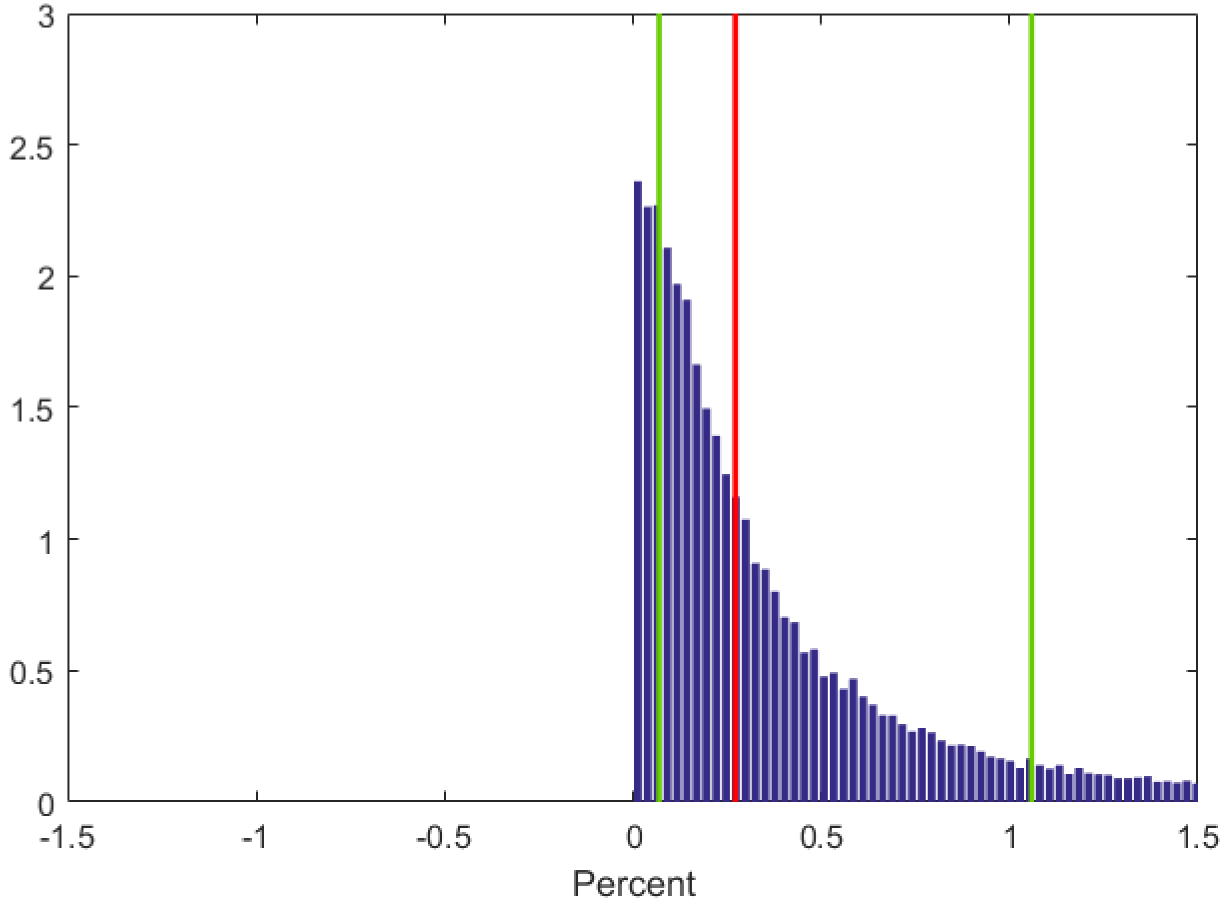

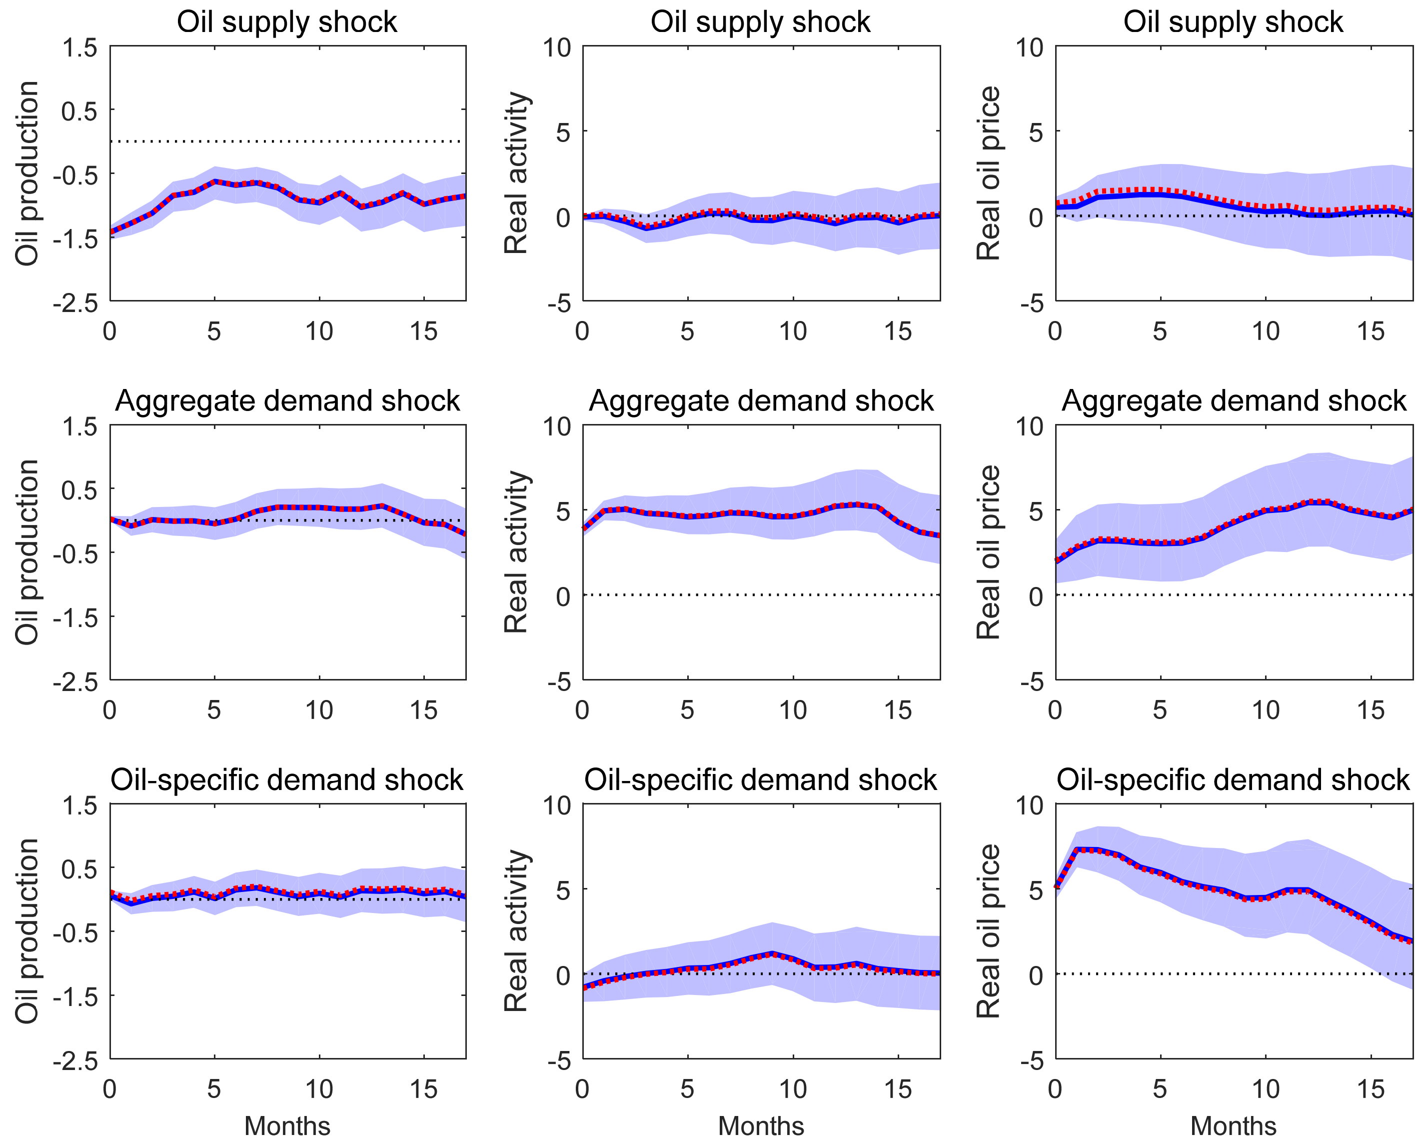

Blue: posterior median IRFs calculated using Baumeister-Hamilton algorithm

Red: IRFs with Kilian-Murphy's (JEEA 2012) sign restrictions

Posterior density of short-run demand elasticity

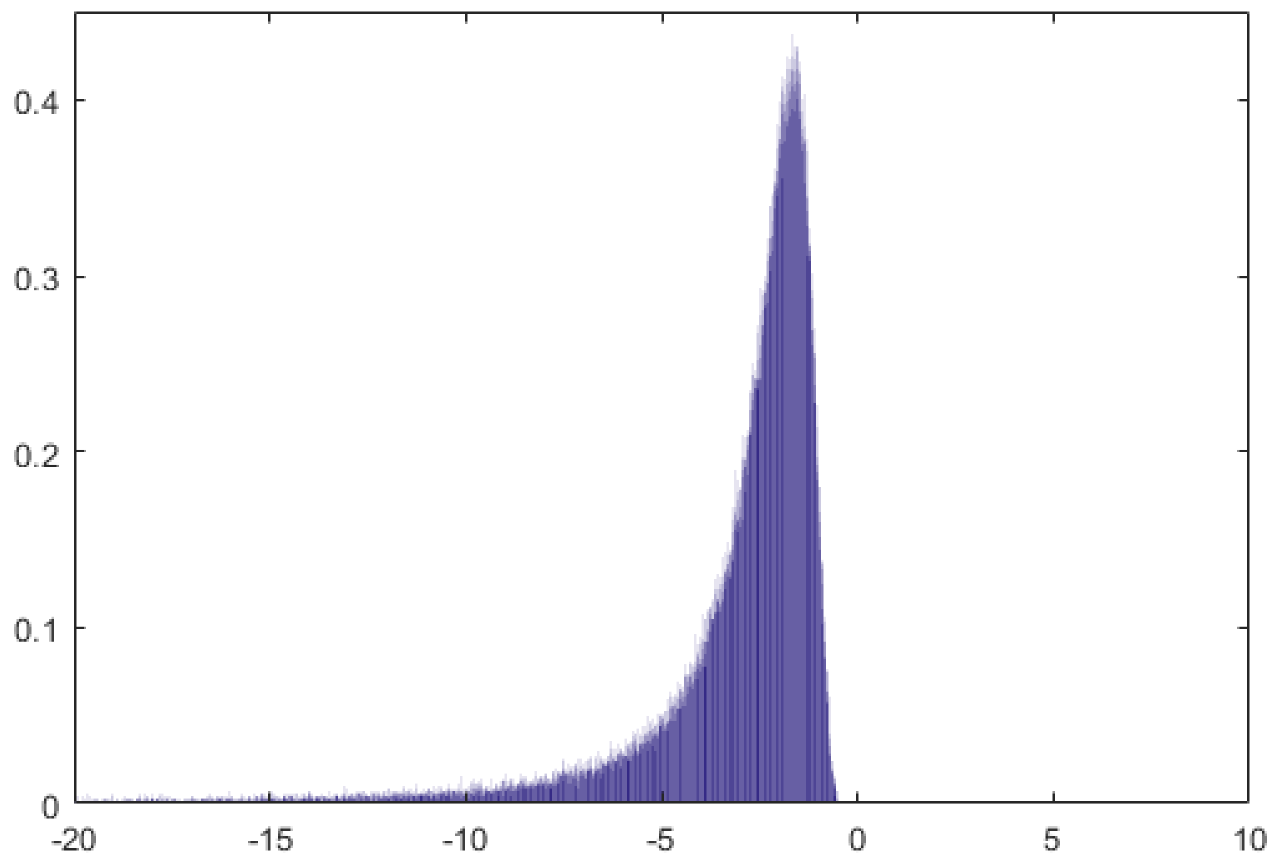

96% posterior probability that demand elasticity < -1

Our proposal:

- Better to use nondogmatic priors and use all available information
- Uncertainty about identifying assumptions should be incorporated in statements about any feature of the model
- Bayesian posterior distribution reflects both randomness in a finite data set and incomplete confidence in identifying assumptions

Structural model of the global oil market

$$q_t = \alpha_{qp} p_t + \mathbf{b}_1' \mathbf{x}_{t-1} + u_{1t}^*$$ (supply)

$$y_t = \alpha_{yp} p_t + \mathbf{b}_2' \mathbf{x}_{t-1} + u_{2t}^*$$ (economic activity)

$$q_t = \beta_{qy} y_t + \beta_{qp} p_t + \Delta i_t^* + \mathbf{b}_3' \mathbf{x}_{t-1} + u_{3t}^*$$ (demand)

$$\Delta i_t^* = \psi_1^* q_t + \psi_2^* y_t + \psi_3^* p_t + \mathbf{b}_4^{*\prime} \mathbf{x}_{t-1} + u_{4t}^*$$

(inventory demand)

$$\Delta i_t = \chi \Delta i_t^* + e_t$$ (inventory measurement error)

What do we know about the price elasticity of demand?

- Cross-country regression of log of petroleum use per dollar of GDP on real price of gasoline for 23 OECD countries in 2004

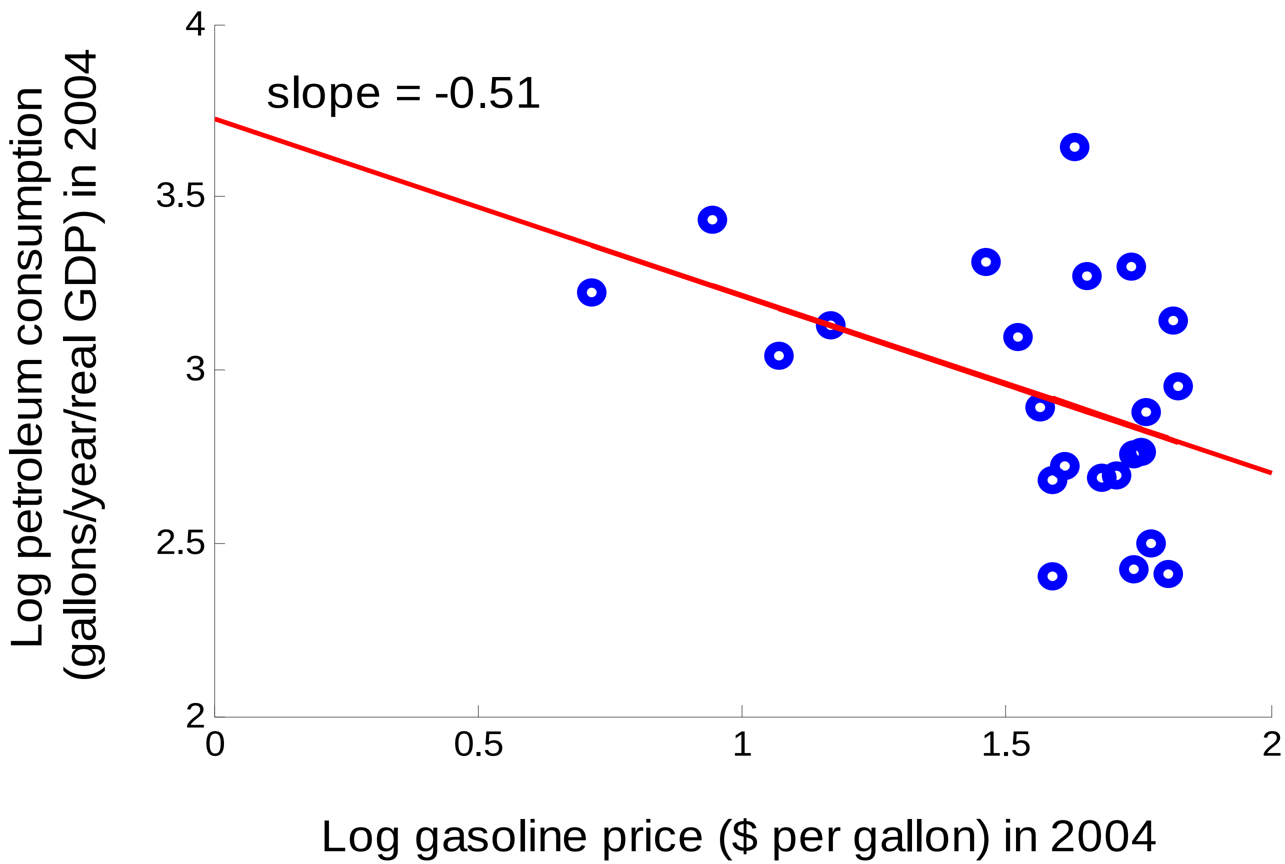

What do we know about the price elasticity of demand?

- Cross-sectional evidence based on household surveys
  - Newey and Hausman (1995): -0.81
  - Yatchew and No (2001): -0.9

- Estimates from previous literature surveys:
  - Dahl and Sterner (1991): -0.86
  - Graham and Glaister (2004): -0.77
  - Brons et al. (2008): -0.84

⇒ short-run elasticity < long-run elasticity

β_{qp}: short-run price elasticity of demand

Student t (location -0.1, scale $0.2 \rightarrow 1.0$, d.f. 3)

truncated ≤ 0

Student *t* prior for price elasticity of demand

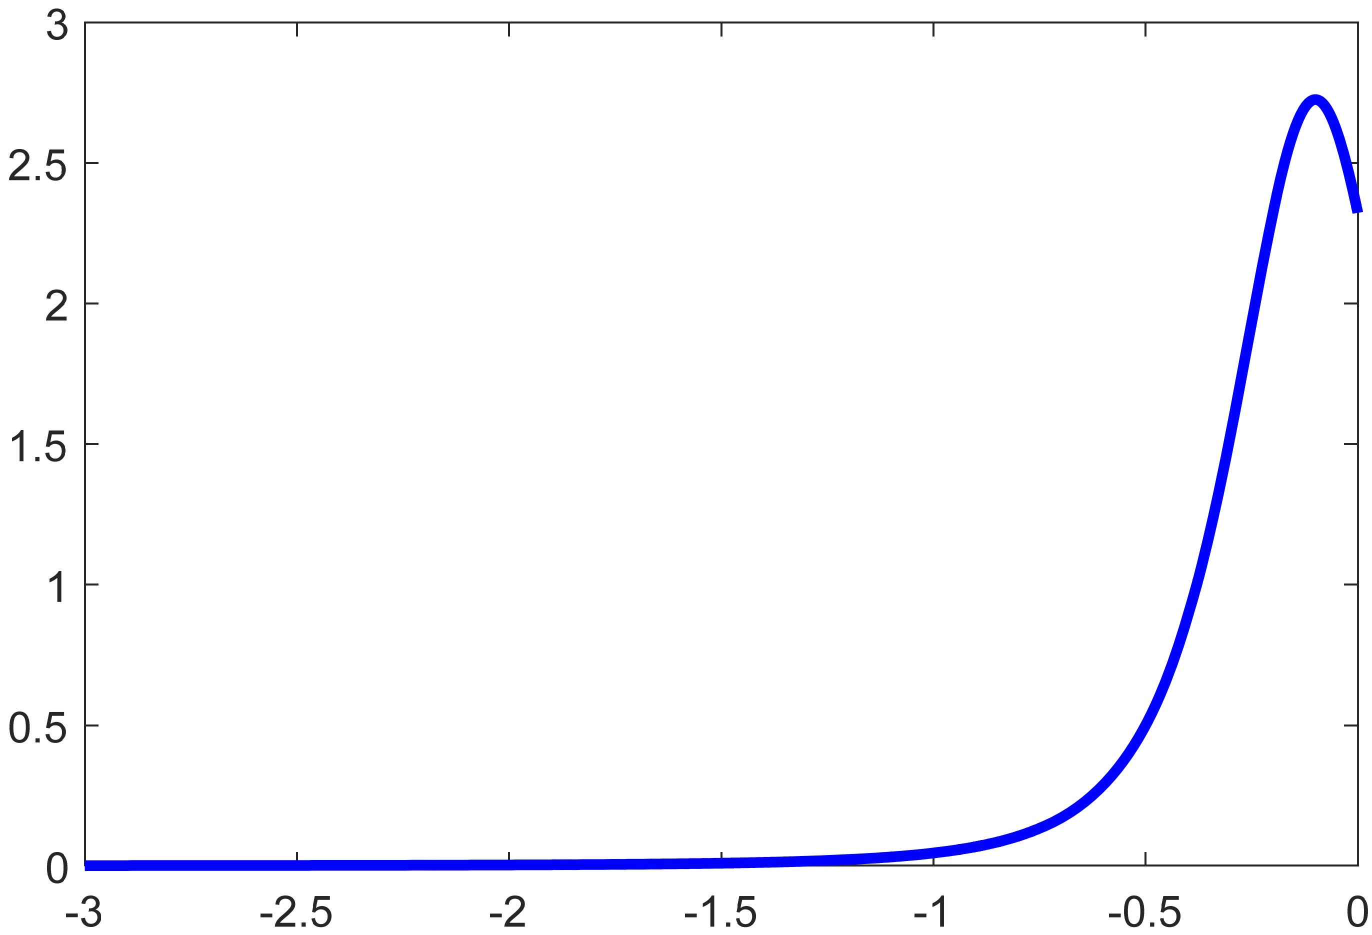

Student *t* prior for price elasticity of demand

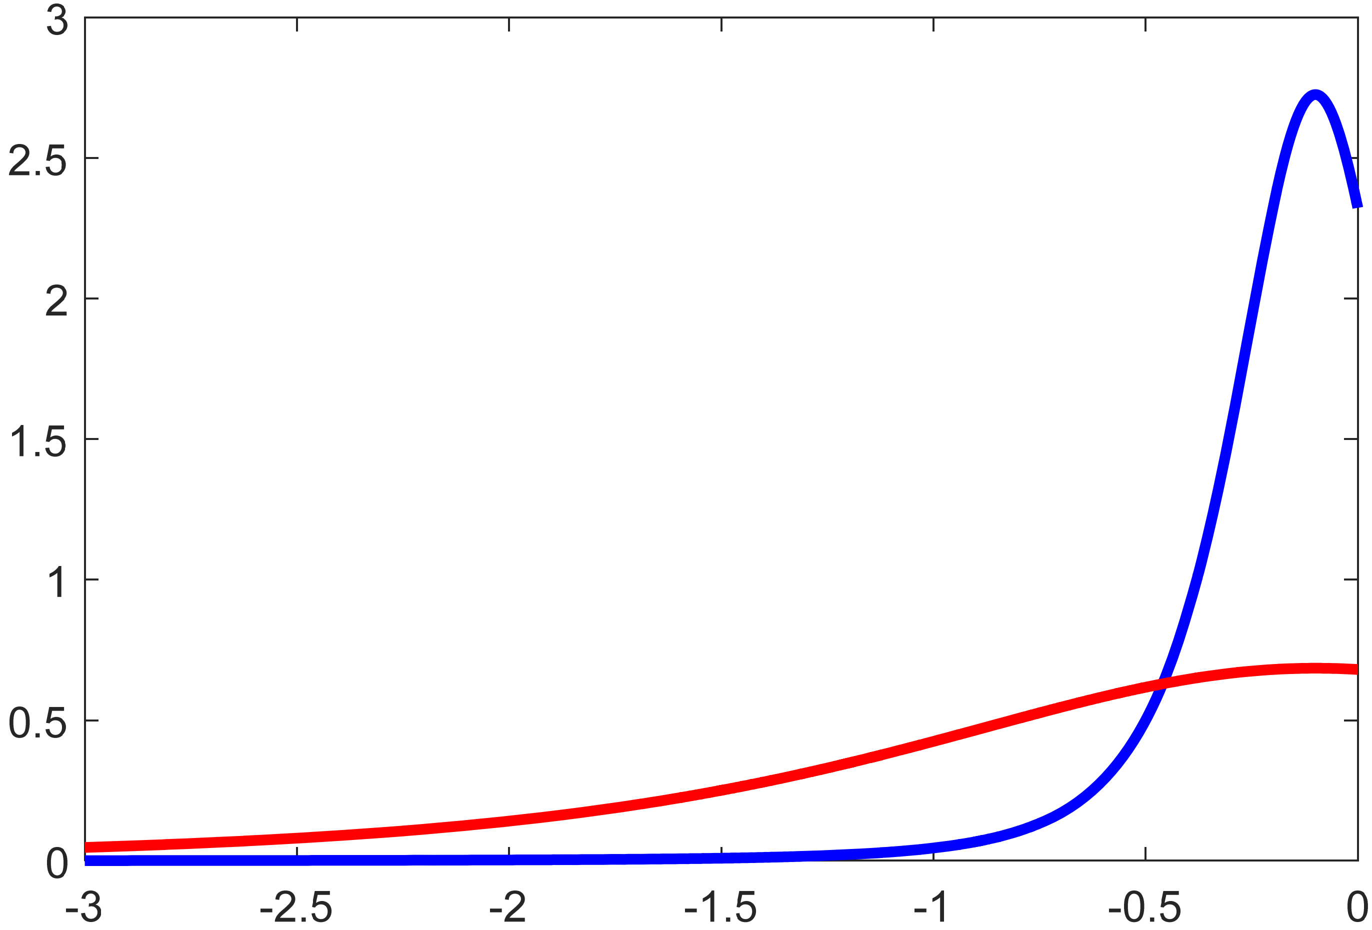

What about the price elasticity of supply?

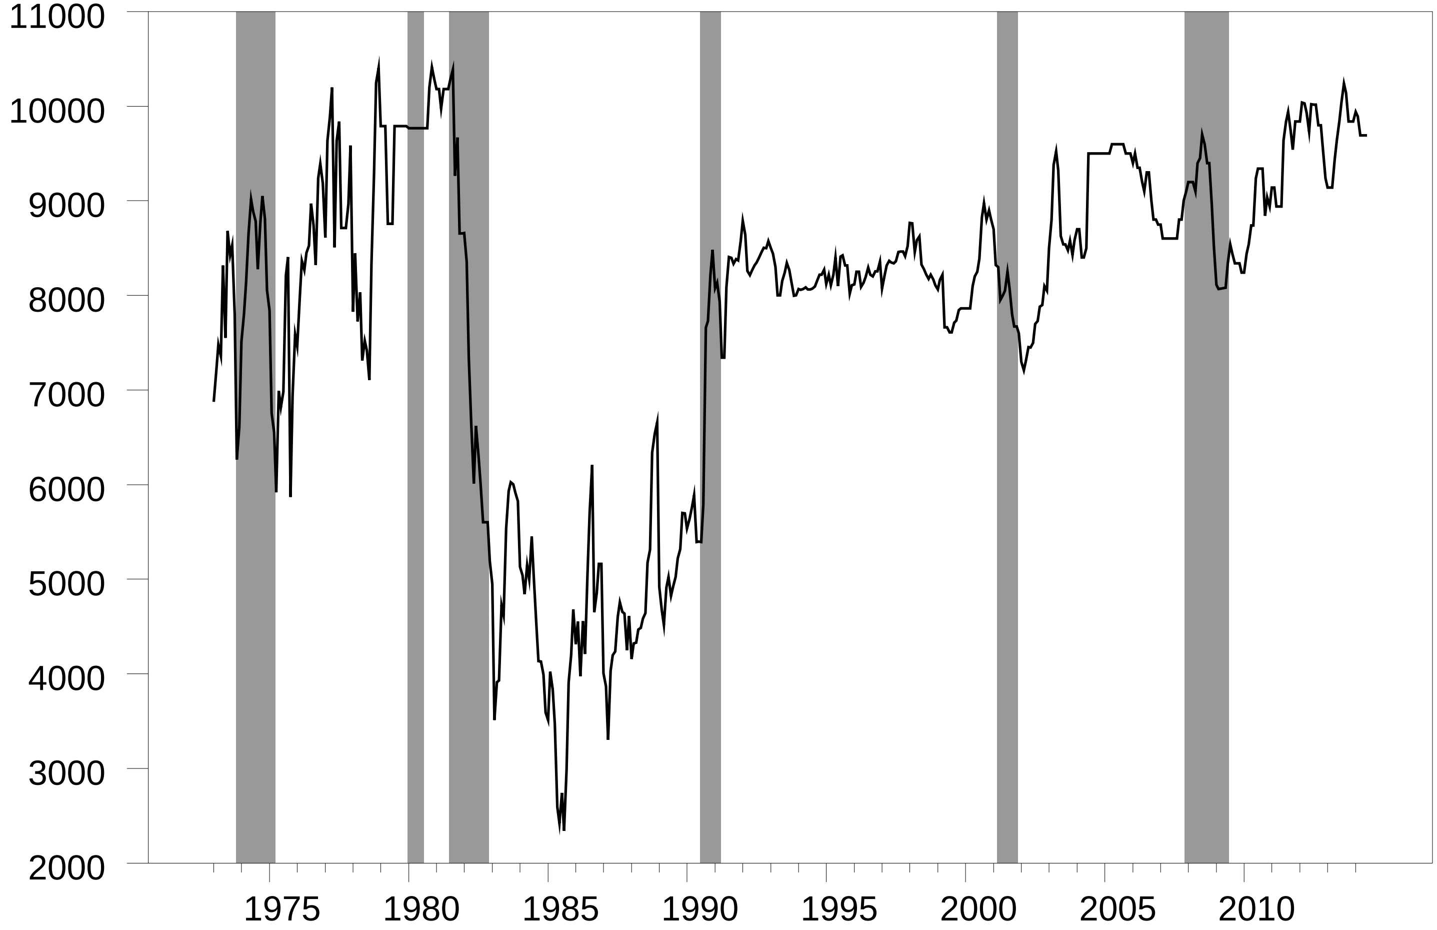

Saudi Arabian monthly production and U.S. recessions: significant monthly response of supply is not implausible

What do we know about the price elasticity of supply?

- Estimates from multiple historical episodes
  - Caldara, Cavallo and Iacoviello (2017): 0.077
- Estimates from individual oil wells in North Dakota
  - Bjornland, Nordvik and Rohrer (2017): 0.2

α_{qp}: short-run price elasticity of supply

Student t (location 0.1, scale $0.2 \to 1.0$, d.f. 3)

truncated ≥ 0

Student *t* prior for price elasticity of supply

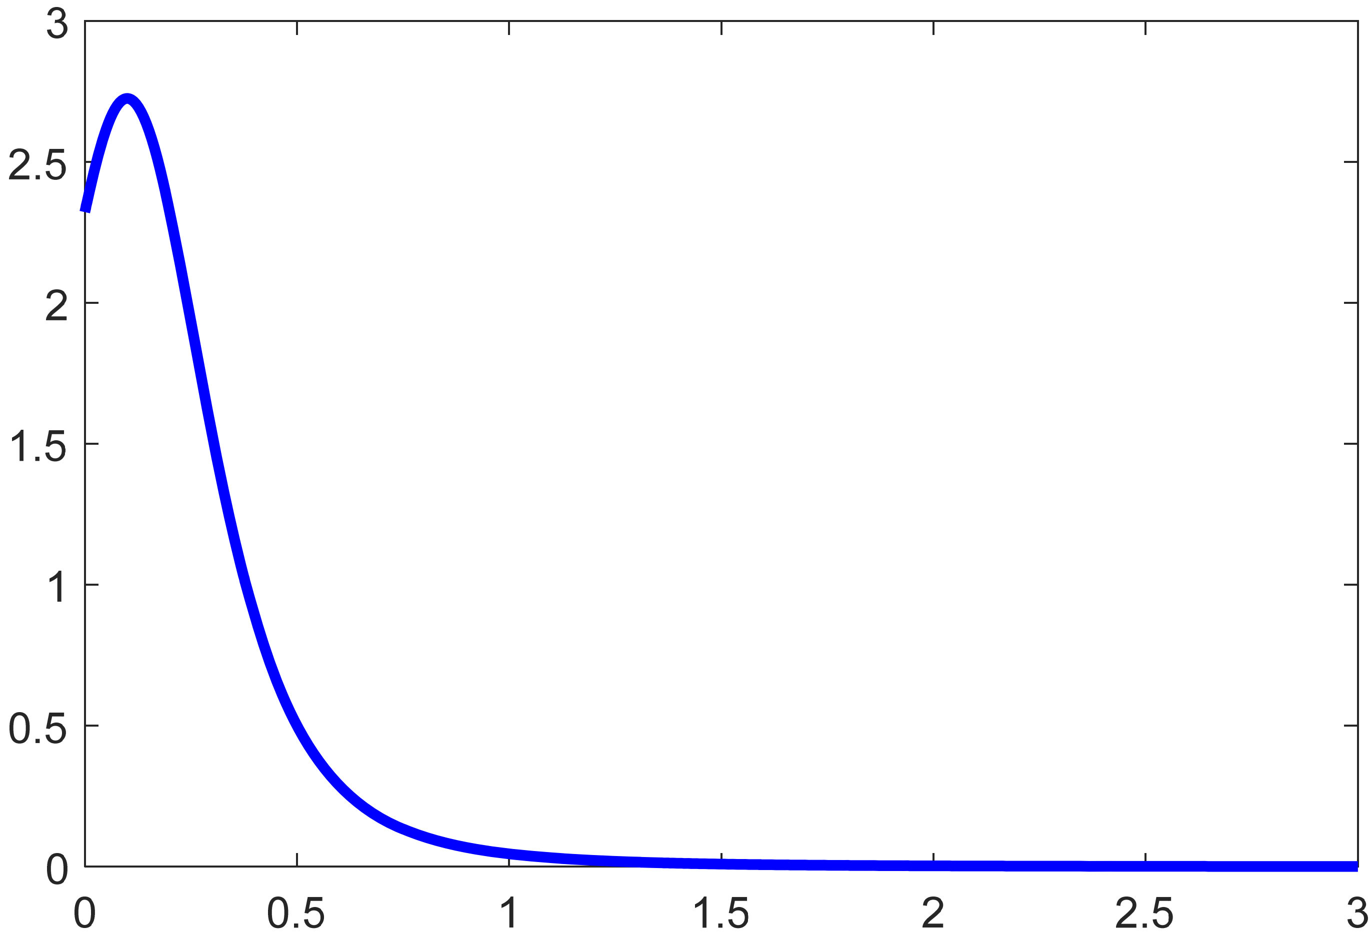

Extensions

- Additional sources of prior information
  - Impact effects of shocks: sign and magnitude restrictions
- Longer time span of data
  - Allow for changes in the information content of observations: downweight more distant ones
- Accommodation of larger systems
  - Informative priors only on subset of structural parameters
  - Identify only subset of shocks

Prior (red) and posterior (blue) distributions for unknown elements of **A**

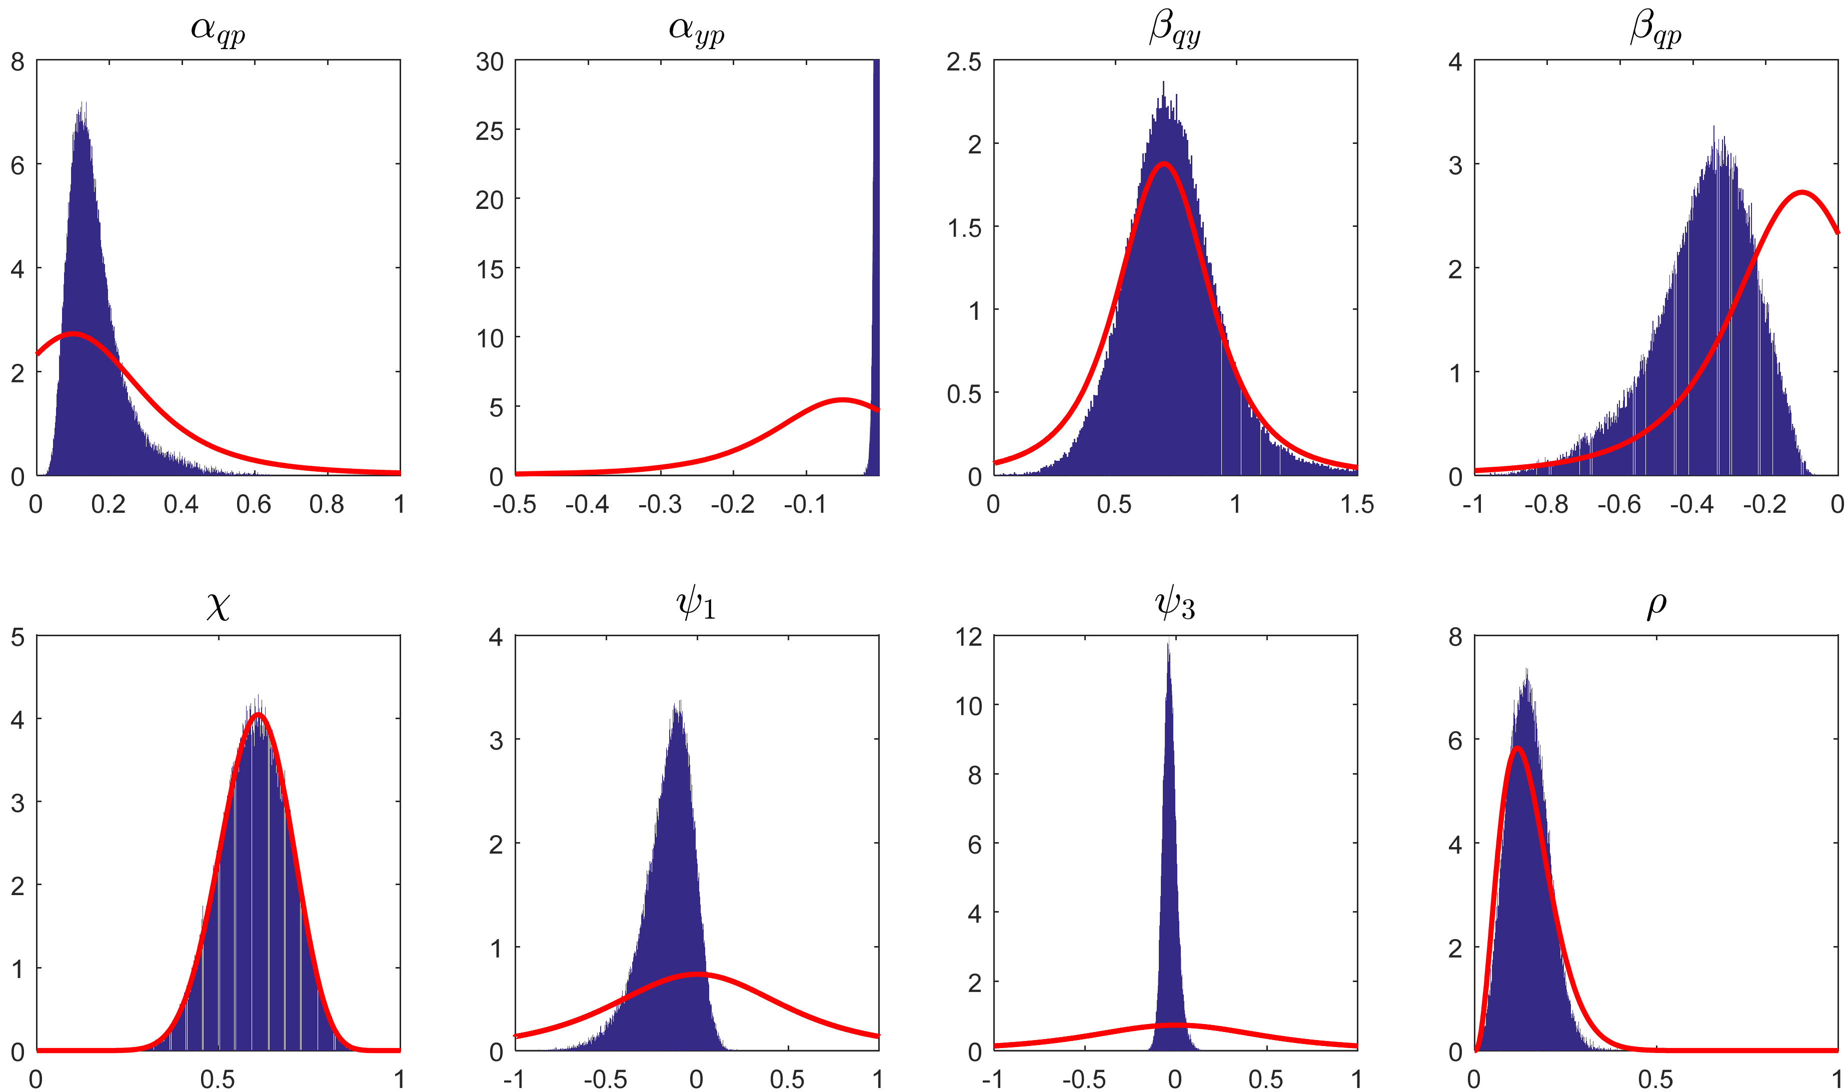

Posterior medians and 68% credibility sets for various magnitudes of interest

Short-run price elasticity of oil supply α_{qp}	**0.15** (0.09, 0.23)
Short-run price elasticity of oil demand β_{qp}	**-0.35** (-0.51, -0.24)
Effect of oil supply shock that raises real oil price 10% on y_{t+12}	**-0.50** (-0.89, -0.16)
Effect of oil demand shock that raises real oil price 10% on y_{t+12}	**0.14** (-0.13, 0.46)
Effect of oil inventory shock that raises real oil price 10% on y_{t+12}	**-0.35** (-0.80, 0.08)

Impulse responses

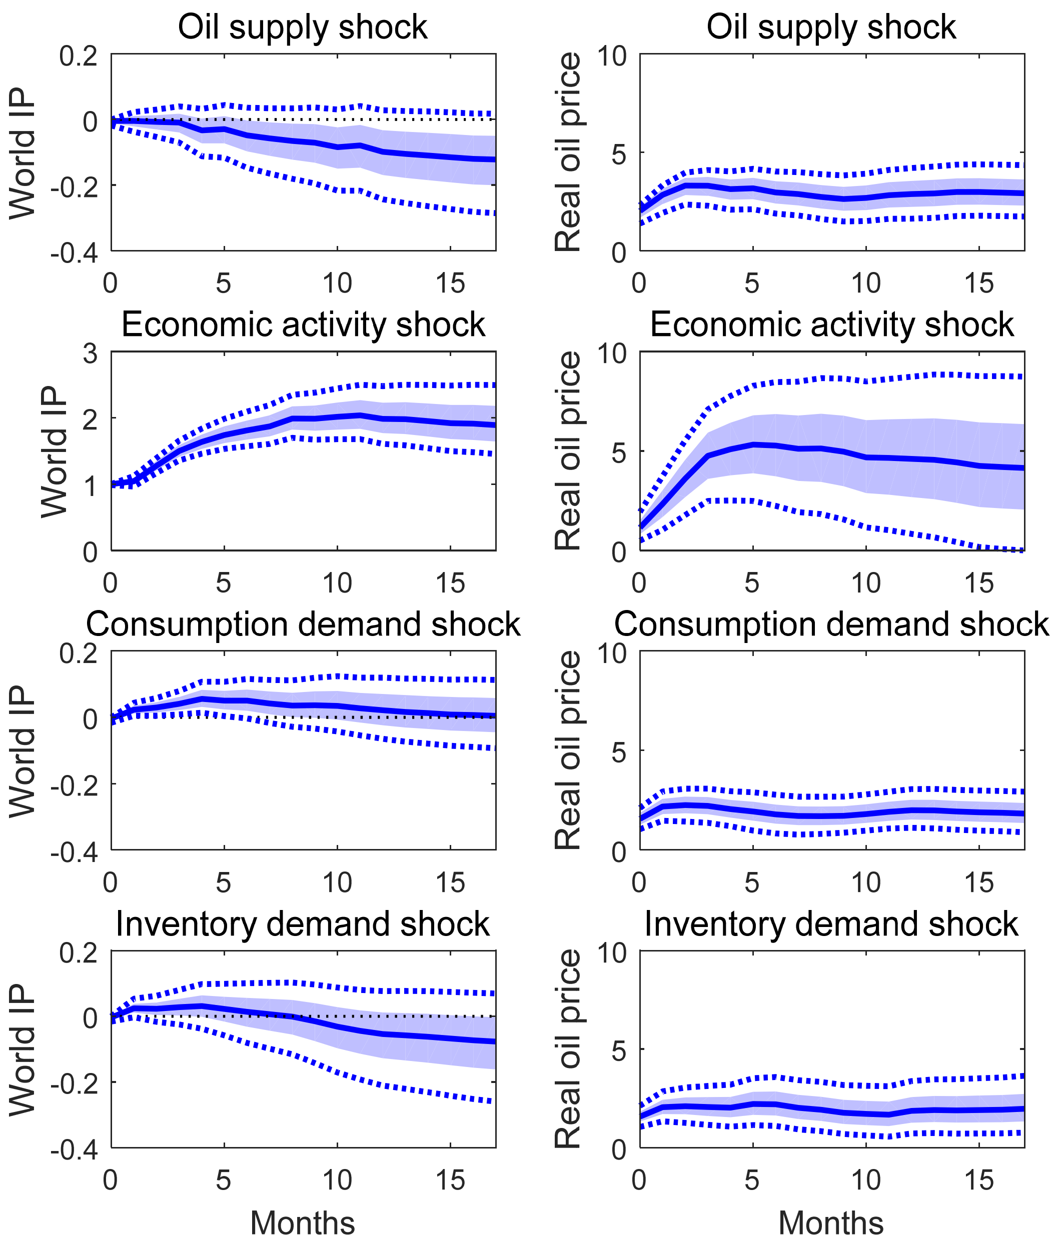

Posterior medians and 68% credibility sets for contribution of oil supply shocks to various historical episodes

Historical episode	Actual real oil price growth (RAC)	Benchmark
(1)	(2)	(3)
June-Oct 90	74.83	**34.47** [46%] (23.84, 46.62)
Jan 07-June 08	86.80	**40.89** [47%] (26.89, 56.84)
June 14-Jan 16	-129.88	**-49.52** [38%] (-71.35, -31.31)
Feb 16-Dec 16	54.01	**16.61** [31%] (8.48, 27.15)

Sensitivity analysis: Less confidence in various components of the prior

Historical episode	Actual real oil price growth (RAC)	Benchmark	Supply and demand elasticities	Lagged structural coefficients	Variances of shocks
		(1)	(2)	(3)	(4)
June 14-Jan 16	-129.88	**-49.52** [38%]	**-33.44** [26%]	**-49.61** [38%]	**-50.58** [39%]
		(-71.35, -31.31)	(-58.70, -15.15)	(-70.87, -31.84)	(-75.23, -31.16)
Feb 16-Dec 16	54.01	**16.61** [31%]	**9.68** [18%]	**16.92** [31%]	**17.06** [32%]
		(8.48, 27.15)	(2.63, 20.98)	(8.69, 27.64)	(8.44, 29.11)

Conclusions

(1) Structural interpretation of correlations only possible by drawing on prior understanding of economic structure

(2) When “incredible” identifying assumptions are relaxed, we find a bigger role for oil supply shocks during key historical episodes

(3) Approach offers many advantages over existing approaches and can be easily applied in many other contexts