Verein für Socialpolitik

Standing Committee on Macroeconomics

Discussion Paper Series

No.: 03/17

Author(s):

Christiane Baumeister, James Hamilton

Title:

Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks

Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks*

Christiane Baumeister
University of Notre Dame
cjsbaumeister@gmail.com

James D. Hamilton
University of California at San Diego
jhamilton@ucsd.edu

October 14, 2015
Revised: November 16, 2015

Abstract

Traditional approaches to structural interpretation of vector autoregressions can be viewed as special cases of Bayesian inference arising from very strong prior beliefs about certain aspects of the model. These traditional methods can be generalized with a less restrictive Bayesian formulation that allows the researcher to summarize uncertainty coming not just from the data but also uncertainty about the model itself. We use this approach to revisit the role of shocks to oil supply and demand and conclude that oil price increases that result from supply shocks lead to a reduction in economic activity after a significant lag, whereas price increases that result from increases in oil consumption demand do not have a significant effect on economic activity.

JEL codes: Q43, C32, E32
keywords: oil prices, vector autoregressions, sign restrictions, Bayesian inference, measurement error

*We are grateful to Lutz Kilian, Dan Murphy, and Eric Swanson for helpful suggestions.

1 Introduction.

Drawing structural inference from vector autoregressions requires making use of prior information. In the usual approach to just-identified models, the researcher proceeds as if he or she is absolutely certain that some structural coefficients are zero (assumptions viewed as necessary to achieve identification) while having no useful information about other magnitudes. Though this approach is extremely prevalent in the literature, most users of these methods are willing to acknowledge that the identifying assumptions are rarely fully persuasive. The structural assumptions that the usual econometric approach regards as known with certainty are in fact almost always subject to serious challenge.

For this reason, it has recently become quite common to try to draw structural conclusions from vector autoregressions using more minimal assumptions such as prior beliefs about the signs of certain structural responses. Although the popular impression is that these methods do not impose meaningful prior beliefs other than the sign restrictions, Baumeister and Hamilton (2015a) demonstrated that the algorithms currently in use by applied researchers in fact imply nonuniform prior distributions for both structural parameters and impulse-response functions and that the influence of the priors does not vanish as the sample size becomes infinite.

Once we acknowledge that prior beliefs necessarily play a role in any structural interpretation of correlations, it becomes helpful to characterize the contribution of prior information using the formal tools of Bayesian analysis. Baumeister and Hamilton (2015a) developed algorithms for Bayesian inference that can be used whether the structural model is over-identified, just-identified, or under-identified. In the latter case, the researcher is explicitly acknowledging some doubts about the credibility of some of the identifying assumptions, and these doubts will be accurately reflected in the posterior inference.

In this paper we use this approach to revisit efforts by Kilian (2009) and Kilian and Murphy (2012) to assess the consequences of supply and demand shocks. We demonstrate that their results can be obtained as special cases of formal Bayesian posterior inference under particular specifications of prior beliefs, namely, near certainty about some functions of parameters and near ignorance about others. We argue that the assumptions that they impose with certainty are implausible, and further that we have additional information about other objects that can significantly improve the inference. We show how these ideas can be adapted in a generalization that allows for relaxation of the identifying assumptions, acknowledges the role of measurement error in variables, and makes use of relations found in earlier data sets. Notwithstanding, our core conclusions are in line with those in previous studies. We find that an oil supply shock leads to a reduction in economic activity after a substantial lag, whereas oil consumption demand shocks are not associated with significant effects.

The plan of the paper is as follows. Section 2 summarizes the Bayesian approach for a model that need not be identified in the frequentist sense. Section 3 uses this framework to revisit earlier studies on the role of oil supply and demand shocks. Section 4 shows how we can allow for measurement error, while Section 5 illustrates how we can bring in information from a variety of sources and data sets to help inform the inference about the role of shocks to oil supply and demand.

2 Bayesian inference for structural vector autoregressions.

Our interest is in dynamic structural models of the form

$$\mathbf{A}\mathbf{y}_t = \mathbf{B}\mathbf{x}_{t-1} + \mathbf{u}_t \tag{1}$$

for $\mathbf{y}_t$ an $(n \times 1)$ vector of observed variables, $\mathbf{A}$ an $(n \times n)$ matrix summarizing their contemporaneous structural relations, $\mathbf{x}_{t-1}$ a $(k \times 1)$ vector (with $k = mn + 1$) containing a constant and m lags of $\mathbf{y}$ ($\mathbf{x}'_{t-1} = (\mathbf{y}'_{t-1}, \mathbf{y}'_{t-2}, ..., \mathbf{y}'_{t-m}, 1)'$), and $\mathbf{u}_t$ an $(n \times 1)$ vector of structural disturbances. As in the vast majority of applied studies, we take the variance matrix of $\mathbf{u}_t$ (denoted $\mathbf{D}$) to be diagonal. To obtain a formal Bayesian solution we treat $\mathbf{u}_t$ as Gaussian, though Baumeister and Hamilton (2015a) showed that the resulting Bayesian posterior distribution can more generally be interpreted as inference about population second moments even if the true innovations are not Gaussian.

2.1 Representing prior beliefs.

From a Bayesian perspective, a researcher's prior beliefs about $\mathbf{A}$ would be represented in the form of a density $p(\mathbf{A})$, where values of $\mathbf{A}$ that are regarded as more plausible a priori are associated with a larger value for $p(\mathbf{A})$, while $p(\mathbf{A}) = 0$ for any values of $\mathbf{A}$ that are completely ruled out. Implementation of our procedure requires only that $p(\mathbf{A})$ be a proper density that integrates to unity.[1]

[1] Actually our algorithm can be implemented even if one doesn't know the constant of integration, so the practical requirement is simply that $p(\mathbf{A})$ is everywhere nonnegative and when integrated over the set of all allowable $\mathbf{A}$ produces a finite positive number. Asymptotic results require $p(\mathbf{A})$ to associate a positive probability with any open neighborhood around the true value $\mathbf{A}_0$.

While we allow the researcher to have arbitrary prior beliefs about $\mathbf{A}$, to reduce computational demands we assume that prior beliefs about the other parameters can be represented by particular families of parametric distributions that allow many features of the Bayesian posterior distribution to be calculated with closed-form analytic expressions. Specifically, we assume that prior beliefs about $\mathbf{D}$ conditional on $\mathbf{A}$ can be represented using $\Gamma(\kappa_i, \tau_i)$ distributions for d_{ii}^{-1},[2]

$$p(\mathbf{D}|\mathbf{A}) = \textstyle\prod_{i=1}^{n} p(d_{ii}|\mathbf{A}) \tag{2}$$

$$p(d_{ii}^{-1}|\mathbf{A}) = \begin{cases} \frac{\tau_i^{\kappa_i}}{\Gamma(\kappa_i)}(d_{ii}^{-1})^{\kappa_i-1}\exp(-\tau_i d_{ii}^{-1}) & \text{for } d_{ii}^{-1} \geq 0 \\ 0 & \text{otherwise} \end{cases},$$

where d_{ii} denotes the row i, column i element of $\mathbf{D}$. Thus κ_i/τ_i denotes the analyst's expected value of d_{ii}^{-1} before seeing the data, while κ_i/τ_i^2 is the variance of this prior distribution. If we have a lot of confidence in these prior beliefs, we would choose κ_i and τ_i to be large numbers to get a prior distribution tightly centered around κ_i/τ_i. A complete absence of useful information about $\mathbf{D}$ could be represented as the limiting case as κ_i and τ_i approach 0. In the applications and formulas below we allow τ_i to depend on $\mathbf{A}$ but assume that κ_i does not; for the more general case when both τ_i and κ_i depend on $\mathbf{A}$ see the treatment in Baumeister and Hamilton (2015a). We offer some suggestions for how to choose the values for κ_i and τ_i in Appendix A.

Prior beliefs about the lagged structural coefficients $\mathbf{B}$ are represented with conditional Gaussian distributions, $\mathbf{b}_i|\mathbf{A}, \mathbf{D} \sim N(\mathbf{m}_i, d_{ii}\mathbf{M}_i)$:

$$p(\mathbf{B}|\mathbf{D}, \mathbf{A}) = \textstyle\prod_{i=1}^{n} p(\mathbf{b}_i|\mathbf{D}, \mathbf{A}) \tag{3}$$

$$p(\mathbf{b}_i|\mathbf{D}, \mathbf{A}) = \frac{1}{(2\pi)^{k/2}|d_{ii}\mathbf{M}_i|^{1/2}} \exp[-(1/2)(\mathbf{b}_i - \mathbf{m}_i)'(d_{ii}\mathbf{M}_i)^{-1}(\mathbf{b}_i - \mathbf{m}_i)]. \tag{4}$$

The vector $\mathbf{m}_i$ denotes our best guess before seeing the data as to the value of $\mathbf{b}_i$, where $\mathbf{b}_i'$ denotes row i of $\mathbf{B}$, that is, $\mathbf{b}_i$ contains the lagged coefficients for the ith structural equation. The matrix $\mathbf{M}_i$ characterizes our confidence in these prior beliefs. A large variance would represent much uncertainty, while having no useful prior information could be regarded as the limiting case when $\mathbf{M}_i^{-1}$ goes to zero. Again the applications and results in this paper allow $\mathbf{m}_i$ to depend on $\mathbf{A}$ but assume that $\mathbf{M}_i$ does not. Suggestions for specifying $\mathbf{m}_i$ and $\mathbf{M}_i$ are offered in Appendix A.

[2] We will follow the notational convention of using $p(.)$ to denote an arbitrary density, where the particular density being referred is implicit by the argument. Thus $p(\mathbf{A})$ is shorthand notation for $p_{\mathbf{A}}(\mathbf{A})$ and represents a different function from $p(\mathbf{D}|\mathbf{A})$, which in more careful notation would be denoted $p_{\mathbf{D}|\mathbf{A}}(\mathbf{D}|\mathbf{A})$.

2.2 Sampling from the posterior distribution.

Although the Bayesian begins with prior beliefs about parameters $p(\mathbf{A}, \mathbf{D}, \mathbf{B})$ represented by the product of $p(\mathbf{A})$ with (2) and (3), the goal is of course to see how observation of the data $\mathbf{Y}_T = (\mathbf{y}_1', \mathbf{y}_2', ..., \mathbf{y}_T')'$ causes us to change these beliefs. The particular distributions recommended above prove to be the natural conjugates. That is, if the prior for d_{ii}^{-1} given $\mathbf{A}$ is $\Gamma(\kappa_i, \tau_i(\mathbf{A}))$, then the posterior for d_{ii}^{-1} given $\mathbf{A}$ and the data $\mathbf{Y}_T$ is $\Gamma(\kappa_i^*, \tau_i^*(\mathbf{A}))$ where

$$\kappa_i^* = \kappa_i + T/2 \tag{5}$$

$$\tau_i^*(\mathbf{A}) = \tau_i(\mathbf{A}) + (1/2)\zeta_i^*(\mathbf{A}) \tag{6}$$

and where $\zeta_i^*(\mathbf{A})$ can be calculated from the sum of squared residuals of a regression of $\tilde{\mathbf{Y}}_i(\mathbf{A})$ on $\tilde{\mathbf{X}}_i$:

$$\zeta_i^*(\mathbf{A}) = \left(\tilde{\mathbf{Y}}_i'(\mathbf{A})\tilde{\mathbf{Y}}_i(\mathbf{A})\right) - \left(\tilde{\mathbf{Y}}_i'(\mathbf{A})\tilde{\mathbf{X}}_i\right)\left(\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i\right)^{-1}\left(\tilde{\mathbf{X}}_i'\tilde{\mathbf{Y}}_i(\mathbf{A})\right) \tag{7}$$

$$\underset{[(T+k)\times 1]}{\tilde{\mathbf{Y}}_i(\mathbf{A})} = \begin{bmatrix} \mathbf{y}_1'\mathbf{a}_i \\ \vdots \\ \mathbf{y}_T'\mathbf{a}_i \\ \mathbf{P}_i'\mathbf{m}_i(\mathbf{A}) \end{bmatrix} \tag{8}$$

$$\underset{[(T+k)\times k]}{\tilde{\mathbf{X}}_i} = \begin{bmatrix} \mathbf{x}_0' \\ \vdots \\ \mathbf{x}_{T-1}' \\ \mathbf{P}_i' \end{bmatrix} \tag{9}$$

for $\mathbf{P}_i$ the Cholesky factor of $\mathbf{M}_i^{-1} = \mathbf{P}_i\mathbf{P}_i'$.

Likewise with a $N(\mathbf{m}_i(\mathbf{A}), d_{ii}\mathbf{M}_i)$ prior for $\mathbf{b}_i|\mathbf{A}, \mathbf{D}$, the posterior for $\mathbf{b}_i$ given $\mathbf{A}$, $\mathbf{D}$, and the data $\mathbf{Y}_T$ is found to be $N(\mathbf{m}_i^*(\mathbf{A}), d_{ii}\mathbf{M}_i^*)$ with

$$\mathbf{m}_i^*(\mathbf{A}) = \left(\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i\right)^{-1}\left(\tilde{\mathbf{X}}_i'\tilde{\mathbf{Y}}_i(\mathbf{A})\right) \tag{10}$$

$$\mathbf{M}_i^* = \left(\tilde{\mathbf{X}}_i'\tilde{\mathbf{X}}_i\right)^{-1}. \tag{11}$$

Baumeister and Hamilton (2015a) showed that the posterior marginal distribution for $\mathbf{A}$ is given by

$$p(\mathbf{A}|\mathbf{Y}_T) = \frac{k_T p(\mathbf{A})[\det(\mathbf{A}\hat{\boldsymbol{\Omega}}_T\mathbf{A}')]^{T/2}}{\prod_{i=1}^n[(2/T)\tau_i^*(\mathbf{A})]^{\kappa_i^*}}\prod_{i=1}^n \tau_i(\mathbf{A})^{\kappa_i}. \tag{12}$$

Here $p(\mathbf{A})$ denotes the original prior density for $\mathbf{A}$, $\hat{\boldsymbol{\Omega}}_T$ is the sample variance matrix for the reduced-form VAR residuals,

$$\hat{\boldsymbol{\Omega}}_T = T^{-1}\left\{\sum_{t=1}^{T}\mathbf{y}_t\mathbf{y}_t' - \left(\sum_{t=1}^{T}\mathbf{y}_t\mathbf{x}_{t-1}'\right)\left(\sum_{t=1}^{T}\mathbf{x}_{t-1}\mathbf{x}_{t-1}'\right)^{-1}\left(\sum_{t=1}^{T}\mathbf{x}_{t-1}\mathbf{x}_{t-1}'\right)\right\}, \tag{13}$$

and k_T is a function of the data and prior parameters (but not dependent on $\mathbf{A}$, $\mathbf{D}$, or $\mathbf{B}$) such that the posterior density integrates to unity over the set of allowable values for $\mathbf{A}$. The value of k_T does not need to be calculated in order to form posterior inference.

The posterior distribution

$$p(\mathbf{A}, \mathbf{D}, \mathbf{B}|\mathbf{Y}_T) = p(\mathbf{A}|\mathbf{Y}_T)p(\mathbf{D}|\mathbf{A}, \mathbf{Y}_T)p(\mathbf{B}|\mathbf{A}, \mathbf{D}, \mathbf{Y}_T)$$

summarizes the researcher's uncertainty about parameters conditional on having observed the sample $\mathbf{Y}_T$. If the model is under-identified, some uncertainty will remain even if the sample size T is infinite, as discussed in detail in Baumeister and Hamilton (2015a). Appendix B describes an algorithm that can be used to generate N different draws from this joint posterior distribution:

$$\{\mathbf{A}^{(\ell)}, \mathbf{D}^{(\ell)}, \mathbf{B}^{(\ell)}\}_{\ell=1}^{N}.$$

Our applications in this paper all set N equal to one million.

2.3 Impulse-response functions.

The structural model (1) has the reduced-form representation

$$\begin{aligned} \mathbf{y}_t &= \boldsymbol{\Phi}\mathbf{x}_{t-1} + \boldsymbol{\varepsilon}_t \qquad (14)\\ &= \boldsymbol{\Phi}_1\mathbf{y}_{t-1} + \boldsymbol{\Phi}_2\mathbf{y}_{t-2} + \cdots + \boldsymbol{\Phi}_m\mathbf{y}_{t-m} + \mathbf{c}+\boldsymbol{\varepsilon}_t \end{aligned}$$

$$\boldsymbol{\Phi} = \mathbf{A}^{-1}\mathbf{B} \tag{15}$$

$$\boldsymbol{\varepsilon}_t = \mathbf{A}^{-1}\mathbf{u}_t. \tag{16}$$

The $(n \times n)$ nonorthogonalized impulse-response matrix at horizon s,

$$\boldsymbol{\Psi}_s = \frac{\partial \mathbf{y}_{t+s}}{\partial \boldsymbol{\varepsilon}_t'}, \tag{17}$$

is then found from the first n rows and columns of $\mathbf{F}^s$, where $\mathbf{F}$ is given by

$$\underset{(nm\times nm)}{\mathbf{F}} = \begin{bmatrix} \mathbf{\Phi}_1 & \mathbf{\Phi}_2 & \cdots & \mathbf{\Phi}_{m-1} & \mathbf{\Phi}_m \\ \mathbf{I}_n & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_n & \cdots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \cdots & \mathbf{I}_n & \mathbf{0} \end{bmatrix}.$$

The dynamic effects of the structural shocks at horizon s are given by

$$\mathbf{H}_s = \frac{\partial \mathbf{y}_{t+s}}{\partial \mathbf{u}_t'} = \mathbf{\Psi}_s \mathbf{A}^{-1}; \tag{18}$$

see for example Hamilton (1994, pages 260 and 331).

Collect the elements of $\mathbf{A}$ and $\mathbf{B}$ in a vector $\boldsymbol{\theta}$ and let $h_{ij}(s;\boldsymbol{\theta})$ denote the row i, column j element of the matrix in (18). Baumeister and Hamilton (2015b) demonstrated that with an L^1 loss function over the impulse-response function $\mathbf{H}_0, \mathbf{H}_1, ..., \mathbf{H}_s$, the optimal estimate of the function from a statistical decision theory perspective is found by calculating the median value of $h_{ij}(s;\boldsymbol{\theta}^{(\ell)})$ over $\ell = 1, ..., N$ separately for each i, j, and s. A 95% posterior credibility set around this optimal point estimate is found from the upper and lower 2.5% quantiles of $h_{ij}(s;\boldsymbol{\theta}^{(\ell)})$ across draws of $\boldsymbol{\theta}^{(\ell)}$, again separately for each i, j, and s.

2.4 Variance decompositions.

We also might be interested in how much of the variability in each of the series could be attributed to the various structural shocks. As in Hamilton (1994, equation [11.5.6]), for the parameters associated with draw ℓ the $(n \times n)$ variance matrix of the s-period-ahead forecast error for all the variables in the system can be written as

$$\mathbf{Q}_s^{(\ell)} = \textstyle\sum_{j=1}^{n} \mathbf{Q}_{js}^{(\ell)} \tag{19}$$

$$\mathbf{Q}_{js}^{(\ell)} = d_{jj}^{(\ell)} \textstyle\sum_{k=0}^{s-1} \mathbf{h}_j(s;\boldsymbol{\theta}^{(\ell)}) \mathbf{h}_j(s;\boldsymbol{\theta}^{(\ell)})' \tag{20}$$

where $\mathbf{h}_j(s;\boldsymbol{\theta})$ denotes the jth column of the impulse-response matrix $\mathbf{H}_s(\boldsymbol{\theta})$ in (18). We calculate the median value of the (i,i) element of (20) as the estimate of the contribution of structural shock j to the MSE of an s-period-ahead forecast of variable i along with its upper and lower 2.5% quantiles.

2.5 Historical decompositions.

Another feature of interest is the contribution of the various structural shocks to different historical episodes. Recall that the value of $\mathbf{y}_{t+s}$ can be written as a function of initial conditions at time t plus the reduced-form innovations between $t+1$ and $t+s$ (e.g., Hamilton, 1994, equation [10.1.14]):

$$\mathbf{y}_{t+s} = \boldsymbol{\varepsilon}_{t+s} + \boldsymbol{\Psi}_1(\boldsymbol{\theta})\boldsymbol{\varepsilon}_{t+s-1} + \boldsymbol{\Psi}_2(\boldsymbol{\theta})\boldsymbol{\varepsilon}_{t+s-1} + \cdots + \boldsymbol{\Psi}_{s-1}(\boldsymbol{\theta})\boldsymbol{\varepsilon}_{t+1} + \mathbf{K}_s(\boldsymbol{\theta})\mathbf{x}_t.$$

Using (16), the reduced-form innovations can in turn be written as $\mathbf{A}^{-1}(\boldsymbol{\theta})\mathbf{u}_t(\mathbf{Y}_T;\boldsymbol{\theta})$ for $\mathbf{u}_t(\mathbf{Y}_T;\boldsymbol{\theta}) = \mathbf{A}\mathbf{y}_t - \mathbf{B}\mathbf{x}_{t-1}$. The contribution of structural shocks between $t+1$ and $t+s$ to the value of $\mathbf{y}_{t+s}$ can then be written as

$$\begin{aligned}
&\mathbf{A}^{-1}(\boldsymbol{\theta})\mathbf{u}_{t+s}(\mathbf{Y}_T;\boldsymbol{\theta}) + \boldsymbol{\Psi}_1(\boldsymbol{\theta})\mathbf{A}^{-1}(\boldsymbol{\theta})\mathbf{u}_{t+s-1}(\mathbf{Y}_T;\boldsymbol{\theta}) + \cdots + \boldsymbol{\Psi}_{s-1}(\boldsymbol{\theta})\mathbf{A}^{-1}(\boldsymbol{\theta})\mathbf{u}_{t+1}(\mathbf{Y}_T;\boldsymbol{\theta}) \\
&= \mathbf{H}_0(\boldsymbol{\theta})\mathbf{u}_{t+s}(\mathbf{Y}_T;\boldsymbol{\theta}) + \mathbf{H}_1(\boldsymbol{\theta})\mathbf{u}_{t+s-1}(\mathbf{Y}_T;\boldsymbol{\theta}) + \cdots + \mathbf{H}_{s-1}(\boldsymbol{\theta})\mathbf{u}_{t+1}(\mathbf{Y}_T;\boldsymbol{\theta})
\end{aligned}$$

for $\mathbf{H}_s(\boldsymbol{\theta})$ the matrix in (18). The contribution of the jth structural shock between dates $t-s$ and t to the value of $\mathbf{y}_t$ is thus given by

$$\boldsymbol{\xi}_{j,t,t-s}(\mathbf{Y}_T;\boldsymbol{\theta}) = \mathbf{h}_j(0;\boldsymbol{\theta})u_{jt}(\mathbf{Y}_T;\boldsymbol{\theta}) + \mathbf{h}_j(1;\boldsymbol{\theta})u_{j,t-1}(\mathbf{Y}_T;\boldsymbol{\theta}) + \cdots + \mathbf{h}_j(s;\boldsymbol{\theta})u_{j,t-s}(\mathbf{Y}_T;\boldsymbol{\theta}) \quad (21)$$

for $\mathbf{h}_j(s;\boldsymbol{\theta})$ the jth column of $\mathbf{H}(s;\boldsymbol{\theta})$. Baumeister and Hamilton (2015b) demonstrated that the optimal estimate of the ith element of (21) is obtained from the median value for each i, j, and t across draws of $\boldsymbol{\theta}^{(\ell)}$, and a 95% credibility set from its 2.5% and 97.5% quantiles.

3 Applications: The role of shocks to oil supply and demand in a 3-variable model.

In this section we illustrate the Bayesian approach with a simple 3-variable description of the global oil market. The first element of the observed vector $\mathbf{y}_t$ is the quantity of oil produced, the second is a measure of real economic activity, and the third captures the real price of oil:

$$\mathbf{y}_t = (q_t, y_t, p_t)'.$$

The estimates reported in this section use the data sets from Kilian (2009) and Kilian and Murphy (2012), in which q_t is the growth rate of monthly world crude oil production, y_t is a cost of international shipping deflated by the U.S. CPI and then reported in deviations from a linear trend, and p_t is the log difference between the EIA series for the refiner acquisition cost of crude oil imports and the U.S. CPI. For details on the various data sets used in this paper see Appendix C.

The structural model of interest consists of the following three equations:

$$q_t = \alpha_{qy} y_t + \alpha_{qp} p_t + \mathbf{b}_1' \mathbf{x}_{t-1} + u_{1t} \tag{22}$$

$$y_t = \alpha_{yq} q_t + \alpha_{yp} p_t + \mathbf{b}_2' \mathbf{x}_{t-1} + u_{2t} \tag{23}$$

$$p_t = \alpha_{pq} q_t + \alpha_{py} y_t + \mathbf{b}_3' \mathbf{x}_{t-1} + u_{3t}. \tag{24}$$

Equation (22) is the oil supply curve, in which α_{qp} is the short-run price elasticity of supply and α_{qy} allows for the possibility that economic activity could enter into the supply decision for reasons other than its effect on price. Oil supply is also presumed to be influenced by lagged values of all the variables over the preceding 2 years, with $\mathbf{x}_{t-1} = (\mathbf{y}_{t-1}', \mathbf{y}_{t-2}', ..., \mathbf{y}_{t-24}', 1)'$. Equation (23) models the determinants of economic activity, with the contemporaneous effects of oil production and oil prices given by α_{yq} and α_{yp}, respectively. Equation (24) governs oil demand, written here in inverse form so that α_{pq} is the reciprocal of the short-run price elasticity of demand and α_{py} is negative one times the ratio of the short-run income elasticity to the short-run price elasticity. One of the goals of the investigation is to distinguish between the consequences of shocks to oil supply u_{1t} and shocks to oil demand u_{3t}.

3.1 A Bayesian interpretation of traditional identification.

As our first example we consider the model estimated by Kilian (2009), who used a familiar recursive interpretation of the structural system with variables ordered as given. From a Bayesian perspective, this amounts to assuming that we know with certainty that production has no contemporaneous response to either price or economic activity, so that $\alpha_{qy} = \alpha_{qp} = 0$, and further that there is no contemporaneous effect of oil prices on economic activity ($\alpha_{yp} = 0$). In contrast to this certainty, the researcher acts as though he or she knows nothing at all about how oil production might affect economic activity (α_{yq}) or the demand parameters (α_{pq} or α_{py}).

We could represent this from a Bayesian perspective using extremely flat priors for the last 3 parameters. For this purpose we used independent Student t distributions with location parameter $c = 0$, scale parameter $\sigma = 100$, and $\nu = 3$ degrees of freedom:

$$p(\alpha_{yq}) = \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\pi\nu}\sigma}\left[1 + \frac{1}{\nu}\left(\frac{\alpha_{yq} - c}{\sigma}\right)^2\right]^{-\frac{\nu+1}{2}}.$$

The specification is then a special case of the model described in Section 2 with

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 \\ -\alpha_{yq} & 1 & 0 \\ -\alpha_{pq} & -\alpha_{py} & 1 \end{bmatrix} \tag{25}$$

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

We also set $\kappa_i = 0.5$, selected τ_i as described in Appendix A, and put a very weak weight on the Doan, Litterman and Sims (1984) random walk prior for the lagged coefficients ($\lambda_0 = 10^9$) to represent essentially no useful prior information about $\mathbf{D}$ and $\mathbf{B}$.

We calculated impulse-response functions for the above model in two ways, first using the traditional Cholesky decomposition of Kilian (2009), with point estimates shown in red in Figure 1. We also show the posterior median (in blue) calculated using the Bayesian algorithm described in Appendix B using the above prior distributions. The two inferences are identical, as of course they should be.

Is there any benefit to giving a Bayesian interpretation to this familiar method? One interesting detail is the implied posterior distributions for α_{yq}, α_{pq}, and α_{py} which are shown in Figure 2. Of particular interest are the prior (shown as a red curve) and posterior (blue histogram) for α_{pq} which is the reciprocal of the short-run price elasticity of demand. The prior distribution is essentially a flat line when viewed on this scale, while the posterior is fairly concentrated between -0.6 and $+0.2$, implying a short-run price elasticity of demand that is contained in $(-\infty, -1.67) \cup (+5, \infty)$. One is thus forced by this identification scheme to conclude that the demand curve is extremely elastic in the short run or possibly even upward sloping.

The claim that we know for certain that supply has no response to price at all within a month, and yet have no reason to doubt that the response of demand could easily be $\pm\infty$ is hardly the place we would have started if we had catalogued from first principles what we expected to find and how surprised we would be at various outcomes. The only reason that thousands of previous researchers have done exactly this kind of thing is that the traditional approach required us to choose some parameters whose values we pretend to know for certain while acting as if we know nothing at all about plausible values for others. Scholars have unfortunately been trained to believe that such a dichotomization is the only way that one could approach these matters scientifically.

The key feature in the data that forces us to impute such unlikely values for the demand elasticity is the very low correlation between the reduced-form residuals for q_t and p_t. If we assume that innovations in q_t represent pure supply shifts, the lack of response of price would force us to conclude that the demand curve is extremely flat and possibly even upward sloping.

3.2 A Bayesian interpretation of sign-restricted VARs.

Many researchers have recognized some of these unappealing aspects of the traditional approach to identification, and as a result have opted instead to use assumptions such as sign restrictions to try to draw a structural inference in VARs. Examples include Baumeister and Peersman (2013a) and Kilian and Murphy (2012), who began with the primitive assumptions that (1) a favorable supply shock (increase in u_{1t}) leads to an increase in oil production, increase in economic activity, and decrease in oil price; (2) an increase in aggregate demand or productivity (increase in u_{2t}) leads to higher oil production, higher economic activity, and higher oil price; and (3) an increase in oil-specific demand leads to higher oil production, lower economic activity, and higher oil price. The assumption is thus that the signs of the elements of $\mathbf{H} = \mathbf{A}^{-1}$ are characterized by

$$\begin{bmatrix} + & + & + \\ + & + & - \\ - & + & + \end{bmatrix}. \tag{26}$$

This is more than an assumption about the signs of all the elements in $\mathbf{A}$ in that it further imposes some complicated constraints on their joint magnitudes. There is, however, a very simple way to implement the above sign restrictions on $\mathbf{H}$ if we impose on $\mathbf{A}$ the restrictions that economic activity only affects oil production contemporaneously through its effect on the price of oil ($\alpha_{qy} = 0$) and that oil production does not enter directly in the economic activity equation ($\alpha_{yq} = 0$). Under these restrictions we would have

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} \\ 0 & 1 & -\alpha_{yp} \\ -\alpha_{pq} & -\alpha_{py} & 1 \end{bmatrix}. \tag{27}$$

Note that although we have imposed two zero restrictions, the model is still unidentified– there is an infinite number of values for $\{\alpha_{qp}, \alpha_{yp}, \alpha_{pq}, \alpha_{py}\}$ that all can achieve the identical maximum value for the likelihood function of the observed data. We can also see that with these two zero restrictions,

$$\mathbf{A}^{-1} = \frac{1}{1-\alpha_{qp}\alpha_{pq}-\alpha_{py}\alpha_{yp}} \begin{bmatrix} 1-\alpha_{py}\alpha_{yp} & \alpha_{qp}\alpha_{py} & \alpha_{qp} \\ \alpha_{yp}\alpha_{pq} & 1-\alpha_{qp}\alpha_{pq} & \alpha_{yp} \\ \alpha_{pq} & \alpha_{py} & 1 \end{bmatrix}. \tag{28}$$

If we believed that the supply curve slopes up ($\alpha_{qp} > 0$), an oil price increase depresses economic activity ($\alpha_{yp} < 0$), the demand curve slopes down ($\alpha_{pq} < 0$), and that higher income boosts oil demand ($\alpha_{py} > 0$), the elements in (28) will always satisfy (26). The under-identified system (27) with these sign restrictions is thus one way of describing the class of models considered by earlier authors.

One of Kilian and Murphy's contributions was to demonstrate that sign restrictions alone are not enough to pin down the magnitudes of interest. They argued that the supply elasticity, although likely not literally zero as assumed in (25), is nevertheless known to be small, which they represented with the bound $\alpha_{qp} \in [0, 0.0258]$. However, they used no other information about the supply elasticity, only imposing that it must fall within this interval. This will be recognized as an essentially Bayesian idea in which the prior density is

$$p(\alpha_{qp}) = \begin{cases} (0.0258)^{-1} & \text{if } \alpha_{qp} \in [0, 0.0258] \\ 0 & \text{otherwise} \end{cases}.$$

This density is plotted in Panel A of Figure 3.

Although we agree with Kilian and Murphy that such prior information can be useful, we find the claim that we are absolutely certain the elasticity could not be as large as 0.0259, and yet regard the values of 0.0257 and 0.001 as equally plausible, as not the most natural way to represent this kind of information.[3] We would recommend instead a prior specification that holds that a value as large as 0.0257 is unlikely, with 0.0259 somewhat less likely, and a continuously declining prior density rather than one with an abrupt cutoff. We also plot in panel A a $\Gamma(\kappa, \tau)$ density with $\kappa = 1.3$ and $\tau = 50$, which we propose as a better representation of the prior information on which Kilian and Murphy are drawing:

$$p(\alpha_{qp}) = \begin{cases} \frac{\tau^{\kappa}}{\Gamma(\kappa)}(\alpha_{qp})^{\kappa-1}\exp(-\tau\alpha_{qp}) & \text{for } \alpha_{qp} \geq 0 \\ 0 & \text{otherwise} \end{cases}. \tag{29}$$

[3]Indeed, Kilian and Murphy themselves seem to share this view, as later in their paper they also reported the results of alternative estimates that instead imposed the bounds $\alpha_{qp} \in [0, 0.0516]$ or $\alpha_{qp} \in [0, 0.0774]$. But in each special case they adopted an estimation method that proceeded as if the researcher knew with certainty that α_{qp} could not fall outside the specified range. By contrast, our favored approach uses a single distribution to summarize the uncertainty that motivates such exploration of different sets of bounds for α_{qp}.

Kilian and Murphy also explored the benefits of using prior information about the response of economic activity to a change in price resulting from the third shock, arguing that we should not expect this magnitude to be large, and presented results in which the product of the (2,3) element of (28) with $\sqrt{d_{33}}$ was restricted to fall in $[-1.5, 0]$. In terms of our motivating structural model, this amounts to a complicated joint restriction on d_{33} and all the elements of $\mathbf{A}$, with some combinations ruled out with certainty but all combinations satisfying the restrictions deemed equally likely. Again it seems more natural to approach such an idea with a simple prior belief that the parameter α_{yp} is small. We implemented this using a gamma distribution with $\kappa = 2.5$ and $\tau = 10$, shown in Panel B of Figure 3:

$$p(\alpha_{yp}) = \begin{cases} \frac{\tau^{\kappa}}{\Gamma(\kappa)}(-\alpha_{qp})^{\kappa-1}\exp(\tau\alpha_{yp}) & \text{for } \alpha_{yp} \leq 0 \\ 0 & \text{otherwise} \end{cases}. \tag{30}$$

The figure also compares this with a prior distribution that insists that $\alpha_{yp}\hat{\omega}_{pp} \in [-1.5, 0]$, where $\hat{\omega}_{pp} = 5.44$ is the square root of the (3,3) element of $\hat{\mathbf{\Omega}}_T$. Although this will not exactly reproduce Kilian and Murphy's method, it should be very similar.

By contrast, Kilian and Murphy do not use any prior information at all about the other parameters (α_{pq} and α_{py}), other than the sign restrictions mentioned above, for which we again adopt the very uninformative Student t priors used in Section 3.1 now truncated by sign restrictions. We then used the algorithm described in Appendix B to calculate the posterior distribution given the prior

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

for $p(a_{qp})$ given by (29), $p(\alpha_{yp})$ given by (30), $p(\alpha_{pq})$ a Student t (0,100,3) density truncated by $\alpha_{pq} < 0$, and $p(\alpha_{py})$ a Student t (0,100,3) density truncated by $\alpha_{py} > 0$. The resulting posterior medians for the impulse-response functions are shown in blue in Figure 4, and indeed coincide almost exactly with the inference reported in Kilian and Murphy's article calculated using their original methodology, which is reproduced as the dashed red lines in Figure 4.[4]

Figure 5 compares prior and posterior distributions for the elements of $\mathbf{A}$. The priors used for α_{qp} and α_{yp} are quite tight and as a result the posterior distributions for these parameters differ very little from the priors. By contrast, we have again used uninformative priors (apart from sign restrictions) for α_{pq} and α_{py}. We have now ruled out an upward-sloping demand curve by assumption, but the specification would still lead the researcher to conclude that

[4] The dashed red lines were produced using the exact methodology of their paper, which is not the posterior median from their model.

monthly demand is extremely sensitive to the current price, with a 60% posterior probability that the on-impact demand elasticity is greater than two in absolute value.

3.3 Do we really know nothing about the elasticity of demand?

In the two examples just discussed, the researchers implicitly proceeded as if we had extremely reliable prior information about the supply elasticity but know nothing about the demand elasticity in the first case and nothing but its sign in the second. In this section we argue that in fact we have a great deal of useful information about the price-elasticity of petroleum demand from other sources.

Figure 6 compares petroleum use per dollar of GDP with the price of gasoline in a cross-section of 23 OECD countries.[5] The relative price of gasoline differs substantially across countries primarily due to differences in taxes. Residents in countries with higher taxes use petroleum less, a finding that is well documented in the literature.[6] The regression line in the first panel of Figure 6 implies an absolute value for the demand elasticity of 0.51 with a standard error of 0.23, statistically significantly greater than zero and less than one. Since tax differentials tend to be stable over time, this coefficient is usually interpreted as a long-run elasticity. For example, one obtains virtually the same regression if 2004 consumption is regressed on 2000 prices, as seen in the second panel of Figure 6.

There also can be no doubt that the short-run price elasticity of demand is substantially less than the long-run elasticity. Consider for example the evidence in Figure 7. The average real retail price of gasoline doubled in the United States between 2002 and 2008, with the price in 2013 still about what it was in 2008. But the average fuel economy of new vehicles sold only increased 9% between 2002 and 2008, and increased an additional 14% after 2008. Of course changes in the fuel economy of the average car on the road increased much more slowly than that for new vehicles sold, as seen in the bottom panel of Figure 7. The changes in petroleum demand that we would expect within one month of a change in price should be significantly smaller than the long-run adjustments implied by the regression lines in Figure 6.

A huge literature has tried to estimate the long-run elasticity of gasoline demand. Hausman and Newey's (1995) estimate from a cross-section of U.S. households was -0.81, while

[5] Data for the price of gasoline and real GDP are from worldbank.org and data for petroleum consumption are from the EIA's *Monthly Energy Review* (Table 11.2). Countries included are Australia, Austria, Belgium, Canada, Denmark, Finland, France, Germany, Greece, Iceland, Ireland, Italy, Japan, the Netherlands, New Zealand, Norway, Portugal, South Korea, Spain, Sweden, Switzerland, the United Kingdom and the United States.

[6] See for example Darmstadter, Dunkerly, and Alterman (1977), Drollas (1984), and Davis (2014).

Yatchew and No's (2001) study of a cross-section of Canadian households came up with -0.9. Dahl and Sterner's (1991) survey of this literature suggested a consensus value of -0.86. Espey's (1998) literature review came up with -0.58; Graham and Glaister (2004) settled on -0.77, while Brons et al. (2008) proposed -0.84. Insofar as taxes and refining costs are a significant component of the user cost for refined products, a 10% increase in the price of crude petroleum should result in a less than 10% increase in the retail price of gasoline, meaning that the price-elasticity of demand for crude oil should be less than that for gasoline. This is indeed confirmed in the smaller literature on estimating the elasticity of oil demand. For example, Dahl's (1993) survey of the literature on developing countries estimated the long-run elasticity at -0.30. We conclude that short-run oil demand elasticities above two in absolute value, such as were implied by the reciprocals of α_{pq} in Figures 2 and 5, are highly implausible.[7]

On the supply side, while it is true that most producers have limited immediate response to changes in price incentives, some countries like Saudi Arabia historically have made quite significant high-frequency adjustments to changing market conditions. Figure 8 shows that in response to weaker demand during the recession of 1981-82, the kingdom reduced production by 6 million barrels per day, implementing by itself an 11% drop in total global production. The Saudis initiated another production decrease of 1.6 mb/d in December 2000 (a few months before the U.S. recession started in March of 2001) and only started to increase production in March of 2002 (four months after the recession had ended). The 1.6 mb/d drop in Saudi production between June 2008 and February 2009 was another clear response to market conditions in an effort to stabilize prices. Equally dramatic in the graph are the rapid increases in Saudi production beginning in August 1990 and January 2003 which were intended to offset some of the anticipated lost production from Iraq associated with the two Gulf Wars. Thus although the response of global production to price increases within the month is likely small, there does not seem to be a solid a priori basis for assuming that it would be significantly smaller than the monthly response of demand.

4 Inventories and measurement error.

Kilian and Murphy (2014) note that another important factor in interpreting short-run comovements of quantities and prices is the behavior of inventories. Increased oil production

[7] An intriguing recent study by Coglianese et al. (forthcoming) used anticipated changes in gasoline taxes as a clever instrument. They estimated a short-run price-elasticity of U.S. gasoline demand of -0.37, with a 95% confidence interval of $(-0.85, +0.11)$.

in month t does not have to be consumed that month but might instead go into inventories:

$$Q_t^S - Q_t^D = \Delta I_t^*.$$

Here Q_t^D is the quantity of oil demanded globally in month t, Q_t^S is the quantity produced, and ΔI_t^* is the true change in global inventories. We append a * to the latter magnitude in recognition of the fact that we have only imperfect observations on this quantity, the implications of which we will discuss below.

Let $q_t = 100\ln(Q_t/Q_{t-1})$ denote the observed monthly growth rate of production. We can then approximate the growth in consumption demand as $q_t - \Delta i_t^*$ for $\Delta i_t^* = 100\Delta I_t^*/Q_{t-1}$. We are thus led to consider the following generalization of the system considered in Section 3.2:

$$q_t = \alpha_{qp} p_t + \mathbf{b}_1' \mathbf{x}_{t-1} + u_{1t}^* \tag{31}$$

$$y_t = \alpha_{yp} p_t + \mathbf{b}_2' \mathbf{x}_{t-1} + u_{2t}^* \tag{32}$$

$$q_t = \beta_{qy} y_t + \beta_{qp} p_t + \Delta i_t^* + \mathbf{b}_3' \mathbf{x}_{t-1} + u_{3t}^* \tag{33}$$

$$\Delta i_t^* = \psi_1^* q_t + \psi_2^* y_t + \psi_3^* p_t + \mathbf{b}_4^{*\prime} \mathbf{x}_{t-1} + u_{4t}^*. \tag{34}$$

Here u_{1t}^*, u_{2t}^*, and u_{3t}^* as before represent shocks to oil supply, economic activity, and oil-specific demand, with the modification to equation (33) acknowledging that oil produced but not consumed in the current period goes into inventories. The shock u_{4t}^* in (34) represents a separate shock to inventory demand, which has sometimes been described as a "speculative demand shock" in the literature.

As noted above, we do not have good data on global oil inventories. There are data on U.S. crude oil inventories and monthly OECD refined-product inventories, from which we can construct a measure of crude oil inventories for OECD countries as in Kilian and Murphy (2014, footnote 6); for details see Appendix C. We represent the fact that these numbers are an imperfect measure of the true magnitude through a measurement-error equation

$$\Delta i_t = \chi \Delta i_t^* + e_t \tag{35}$$

where Δi_t denotes our estimate of the change in OECD crude-oil inventories as a percent of the previous month's world production, $\chi < 1$ is a parameter representing the fact that OECD inventories are only part of the world total, and e_t reflects measurement error. Although the problem of having imperfect measurements on key variables is endemic in macroeconomics, it has been virtually ignored in most of the large literature on structural vector autoregressions because it was not clear how to allow for it using traditional methods.[8] However, it is straightforward to incorporate measurement error in our Bayesian framework, as we now demonstrate.

We can use (35) to rewrite (33) and (34) in terms of observables:

$$q_t = \beta_{qy} y_t + \beta_{qp} p_t + \chi^{-1} \Delta i_t + \mathbf{b}_3' \mathbf{x}_{t-1} + u_{3t}^* - \chi^{-1} e_t \tag{36}$$

$$\Delta i_t = \psi_1 q_t + \psi_2 y_t + \psi_3 p_t + \mathbf{b}_4' \mathbf{x}_{t-1} + \chi u_{4t}^* + e_t \tag{37}$$

where $\psi_j = \chi \psi_j^*$ for $j = 1, 2, 3$. Equations (31), (32), (36), and (37) will be recognized as a system of the form

$$\tilde{\mathbf{A}} \mathbf{y}_t = \tilde{\mathbf{B}} \mathbf{x}_{t-1} + \tilde{\mathbf{u}}_t \tag{38}$$

$$\mathbf{y}_t = (q_t, y_t, p_t, \Delta i_t)'$$

$$\tilde{\mathbf{A}} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 \\ 0 & 1 & -\alpha_{yp} & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -\chi^{-1} \\ -\psi_1 & -\psi_2 & -\psi_3 & 1 \end{bmatrix}$$

$$\tilde{\mathbf{u}}_t = \begin{bmatrix} u_{1t}^* \\ u_{2t}^* \\ u_{3t}^* - \chi^{-1} e_t \\ \chi u_{4t}^* + e_t \end{bmatrix}.$$

Note that although we have explicitly modeled the role of measurement error in contributing to contemporaneous correlations among the variables, we have greatly simplified the analysis by specifying the lagged dynamics of the structural system directly in terms of the observed variables. That is, we are defining $\mathbf{x}_{t-1}$ in (31)-(34) to be based on lags of Δi_{t-j} rather than Δi_{t-j}^*.

The system (38) is not quite in the form of (1) because even though the structural shocks u_{jt}^* are contemporaneously uncorrelated, the observed analogues $\tilde{u}_{jt}$ will be correlated as a

[8] Notable exceptions are Cogley and Sargent (2015) who allowed for measurement error using a state-space model and Amir-Ahmadi, Matthes, and Wang (2015) who identified measurement error from the difference between preliminary and revised data.

result of the measurement error:

$$\tilde{\mathbf{D}} = E(\tilde{\mathbf{u}}_t\tilde{\mathbf{u}}_t') = \begin{bmatrix} d_{11}^* & 0 & 0 & 0 \\ 0 & d_{22}^* & 0 & 0 \\ 0 & 0 & d_{33}^* + \chi^{-2}\sigma_e^2 & -\chi^{-1}\sigma_e^2 \\ 0 & 0 & -\chi^{-1}\sigma_e^2 & \chi^2 d_{44}^* + \sigma_\varepsilon^2 \end{bmatrix}. \tag{39}$$

However, it's not hard to see that $\mathbf{\Gamma}\tilde{\mathbf{D}}\mathbf{\Gamma}' = \mathbf{D}$ is diagonal for

$$\mathbf{\Gamma} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \rho & 1 \end{bmatrix}$$

$$\rho = \frac{\chi^{-1}\sigma_e^2}{d_{33}^* + \chi^{-2}\sigma_e^2}. \tag{40}$$

We can thus write (38) in the form of (1) simply by defining $\mathbf{A} = \mathbf{\Gamma}\tilde{\mathbf{A}}$, $\mathbf{B} = \mathbf{\Gamma}\tilde{\mathbf{B}}$, and

$$\mathbf{u}_t = \mathbf{\Gamma}\tilde{\mathbf{u}}_t = \begin{bmatrix} u_{1t}^* \\ u_{2t}^* \\ u_{3t}^* - \chi^{-1}e_t \\ \chi u_{4t}^* + \rho u_{3t}^* + (1-\rho/\chi)e_t \end{bmatrix} \tag{41}$$

whose variance matrix we denote $\mathbf{D} = \text{diag}(d_{11}, d_{22}, d_{33}, d_{44})$. This is then exactly in the form of the general class of models discussed in Section 2. Specifically, we summarize any prior information about the contemporaneous coefficients in terms of the prior distribution

$$p(\mathbf{A}) = p(\rho)p(\alpha_{qp})p(\alpha_{yp})p(\beta_{qy})p(\beta_{qp})p(\chi)p(\psi_1)p(\psi_2)p(\psi_3)$$

and then follow exactly the methods in Appendix B for inference about $\mathbf{D}$ and $\mathbf{B}$ to generate posterior draws from $p(\mathbf{A}, \mathbf{D}, \mathbf{B}|\mathbf{Y}_T)$.

5 Inference for a 4-variable model with inventories and measurement error.

In addition to making use of the prior information about price elasticities reviewed in Section 3.3, we propose to use prior information about coefficients involving the economic activity measure y_t. For this purpose it is very helpful to use a more conventional measure of economic activity in place of the proxy based on shipping costs that was used in Kilian (2009) and Kilian and Murphy (2012, 2014); among other benefits this allows us to draw directly on information about income elasticities from previous studies. We developed an extended version of the OECD's index of monthly industrial production in the OECD and 6 major other countries as described in Appendix C.

Of course, even more important than having good prior information is having more data. Kilian (2009) and Kilian and Murphy (2012) used the refiner acquisition cost as the measure of crude oil prices. Their series begins in January 1973. Taking differences and including 24 lags means that the first value for the dependent variable in their regressions is February 1975. Thus their analysis makes no use of the important economic responses in 1973 and 1974 to the large oil price increases at the time, nor any earlier observations.

Kilian and Vigfusson (2011) argued that use of the older data is inappropriate since structural relations may have changed over time, suggesting that this is a reason to ignore the earlier data altogether. Moreover, their preferred oil price measure (U.S. refiner acquisition cost, or RAC) is not available before 1974, which might seem to make use of earlier data infeasible. Here again the Bayesian approach offers a compelling advantage, in that we can use results obtained from estimating the model using earlier data for the price of West Texas Intermediate as a prior for the analysis of the subsequent RAC data, putting as much or as little weight as desired on the earlier data set. We describe how this can be done below.

5.1 Informative priors for structural parameters.

This section summarizes the prior information used in our structural inference about the system (31)-(34).

5.1.1 Priors for A.

The discussion of the evidence from separate data sets in Section 3.3 leads us to expect that both the short-run demand elasticity β_{qp} and the short-run supply elasticity α_{qp} should be small in absolute value. We represent this with a prior for β_{qp} that is a Student $t(c^{\beta}_{qp}, \sigma^{\beta}_{qp}, \nu^{\beta}_{qp})$ with mode at $c^{\beta}_{qp} = -0.1$, scale parameter $\sigma^{\beta}_{qp} = 0.2$, $\nu^{\beta}_{qp} = 3$ degrees of freedom, and truncated to be negative. Our prior for α_{qp} is Student $t(c^{\alpha}_{qp}, \sigma^{\alpha}_{qp}, \nu^{\alpha}_{qp})$ with mode at $c^{\alpha}_{qp} = 0.1$, scale parameter $\sigma^{\alpha}_{qp} = 0.2$, $\nu^{\alpha}_{qp} = 3$ degrees of freedom, and truncated to be positive.

Because we use a conventional measure of industrial production we are able to make use of other evidence about the income elasticity of oil demand. Gately and Huntington (2002) reported a nearly linear relationship between log income and log oil demand in developing countries with elasticities ranging between 0.7 and 1, but smaller income elasticities in industrialized countries with values between 0.4 and 0.5. For oil-exporting countries they found an income elasticity closer to 1. Csereklyei, Rubio, and Stern (2016) found that the income elasticity of energy demand is remarkably stable across countries and across time at a value of around 0.7. For our prior for β_{qy} we use a Student t density with mode at 0.7, scale parameter 0.2, 3 degrees of freedom, and truncated to be positive.

We expect the effect of oil prices on economic activity α_{yp} to be small given the small dollar share of crude oil expenditures compared to total GDP (see for example the discussion in Hamilton, 2013). We represent this with a Student t distribution with mode -0.05, scale 0.1, 3 degrees of freedom, and truncated to be negative.

The parameter χ reflects the fraction of total world inventories that are held in OECD countries. Since OECD countries account for around 60% of world petroleum consumption on average over our sample period, a natural expectation is that they also account for about 60% of global inventory. Since χ is necessarily a fraction between 0 and 1, we use a Beta distribution with parameters $\alpha_\chi = 15$ and $\beta_\chi = 10$, which has mean 0.6 and standard deviation of about 0.1.

For the parameters of the inventory equation, we assume that inventories depend on income only through the effects of income on quantity or price. This allows us to set $\psi_2 = 0$ to help with identification. We use relatively uninformative priors for the other coefficients, taking both ψ_1 and ψ_3 to be unrestricted Student t centered at 0 with scale parameter 0.5, and 3 degrees of freedom.

The parameter ρ in (40) captures the importance of inventory measurement error and is also between 0 and 1 by construction. For this we use a Beta distribution with parameters $\alpha_\rho = 3$ and $\beta_\rho = 9$, which has a mean of 0.25 and standard deviation of 0.12.

The prior for the contemporaneous coefficients is then the product of the above densities:

$$p(\mathbf{A}) = p(\alpha_{qp})p(\alpha_{yp})p(\beta_{qy})p(\beta_{qp})p(\chi)p(\psi_1)p(\psi_3)p(\rho)$$

subject to the sign restrictions described above.

5.1.2 Priors for D given A.

Our priors for the reciprocals of the structural variances are independent Gamma distributions, $d_{ii}^{-1}|\mathbf{A} \sim \Gamma(\kappa, \tau_i(\mathbf{A}))$, that reflect the scale of the data as measured by the standard deviation of 12th-order univariate autoregressions fit to the 4 elements of $\mathbf{y}_t$ over $t = 1, ..., T_1$ for T_1

the number of observations in the earlier sample. Letting $\hat{\mathbf{S}}$ denote the estimated variance-covariance matrix of these univariate residuals, we set $\kappa = 2$ (which give the priors a weight of about 4 observations in the first subsample) and $\tau_i(\mathbf{A}) = \kappa \mathbf{a}_i' \hat{\mathbf{S}} \mathbf{a}_i$ where $\mathbf{a}_i'$ denotes the ith row of $\mathbf{A}$.

5.1.3 Priors for B given A and D.

Our priors for the lagged coefficients in the ith structural equation are independent Normals, $\mathbf{b}_i|\mathbf{A}, \mathbf{D} \sim N(\mathbf{m}_i, d_{ii}\mathbf{M})$. Our prior expectation is that the best predictor of the current oil production, economic activity, oil price, or inventory level is its own lagged value, implying that most coefficients for predicting the first-differences in these variables are zero. We allow for the possibility that the 1-period-lag response of supply or demand to a price increase could be similar to the contemporaneous magnitudes, and for this reason set the third element of $\mathbf{m}_1$ to $+0.1$ and the third element of $\mathbf{m}_3$ to -0.1; this gives us a little more information to try to distinguish supply and demand shocks. All other elements of $\mathbf{m}_i$ are set to zero for $i = 1, ..., 4$. For $\mathbf{M}$, which governs the variances of these priors, we follow Doan, Litterman and Sims (1984) in having more confidence that coefficients on higher lags are zero. We implement this by setting diagonal elements of $\mathbf{M}$ to the values specified in equation (43) and other elements of $\mathbf{M}$ to zero, as detailed in Appendix A. For our baseline analysis, we use a value of $\lambda_0 = 0.5$ to control the overall informativeness of these priors on lagged coefficients, which amounts to weighting the prior on the lag-one coefficients equal to about 2 observations.

5.1.4 Using observations from an earlier sample to further inform the prior.

We propose to use observations over 1958:M1 to 1975:M1 to further inform our prior. The observation vector $\mathbf{y}_t$ for date t in this first sample consists of the growth rate of world oil production, growth rate of OECD+6 industrial production, growth rate of WTI, and change in estimated OECD inventories as a percent of the previous month's oil production. We have T_1 observations in the first sample for this $(n \times 1)$ vector $\{\mathbf{y}_t\}_{t=1}^{T_1}$ and associated $(nm + 1 \times 1)$ vector of lagged values and the constant term $\{\mathbf{x}_{t-1}\}_{t=1}^{T_1}$. For the second sample (1975:M2 to 2014:M12) we use the percent change in the refiner acquisition cost (RAC) for the third element of $\mathbf{y}_t$ for which we have observations $\{\mathbf{y}_t, \mathbf{x}_{t-1}\}_{t=T_1+1}^{T_1+T_2}$. Denote the observations for the first sample by $\mathbf{Y}^{(1)}$ and those for the second sample by $\mathbf{Y}^{(2)}$ and collect all the unknown elements of $\mathbf{A}$, $\mathbf{D}$, and $\mathbf{B}$ in a vector $\boldsymbol{\lambda}$.

If we regarded both samples as equally informative about $\boldsymbol{\lambda}$ we could simply collect all the data in a single sample $\mathbf{Y}_T = \{\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}\}$ and apply our method directly to find $p(\boldsymbol{\lambda}|\mathbf{Y}_T)$. This would be numerically identical to using our method to find the posterior distribution from the first sample alone $p(\boldsymbol{\lambda}|\mathbf{Y}^{(1)})$ and then using this distribution as the prior for analyzing the second sample (see Appendix D for demonstration of this and subsequent claims). We propose instead to use as a prior for the second sample an inference that downweights the influence of the first-sample data $\mathbf{Y}^{(1)}$ by a factor $0 \leq \mu \leq 1$. When $\mu = 1$ the observations in the first sample are regarded as equally important as those in the second, while when $\mu = 0$ the first sample is completely discarded. Our baseline analysis below sets $\mu = 0.5$, which regards observations in the first sample as only half as informative as those in the second.

Implementing this procedure requires a simple modification of the procedure described in Section 2.2. We replace equations (8), (9) and (5) with

$$\underset{(T_1+T_2+k)\times 1}{\tilde{\mathbf{Y}}_i(\mathbf{A})} = (\sqrt{\mu}\mathbf{y}_1'\mathbf{a}_i, ..., \sqrt{\mu}\mathbf{y}_{T_1}'\mathbf{a}_i, \mathbf{y}_{T_1+1}'\mathbf{a}_i, ..., \mathbf{y}_{T_1+T_2}'\mathbf{a}_i, \mathbf{m}_i{}'\mathbf{P})'$$

$$\underset{(T_1+T_2+k)\times k}{\tilde{\mathbf{X}}} = \begin{bmatrix} \sqrt{\mu}\mathbf{x}_0 & \cdots & \sqrt{\mu}\mathbf{x}_{T_1-1} & \mathbf{x}_{T_1} & \cdots & \mathbf{x}_{T_1+T_2-1}' & \mathbf{P} \end{bmatrix}'$$

$$\kappa^* = \kappa + (\mu T_1 + T_2)/2$$

for $\mathbf{P}$ the matrix whose diagonal elements are reciprocals of the square roots of (43). We then calculate $\tau_i^*(\mathbf{A})$ and $\zeta_i^*(\mathbf{A})$ using expressions (6) and (7) and replace (12) and (13) with

$$p(\mathbf{A}|\mathbf{Y}_T) = \frac{k_T p(\mathbf{A})[\det(\mathbf{A}\tilde{\boldsymbol{\Omega}}_T\mathbf{A}')]^{(\mu T_1+T_2)/2}}{\prod_{i=1}^n [2\tau_i^*(\mathbf{A})/(\mu T_1 + T_2)]^{\kappa_i^*}} \prod_{i=1}^n \tau_i(\mathbf{A})^{\kappa_i}$$

$$\tilde{\boldsymbol{\Omega}}_T = (\mu T_1 + T_2)^{-1}(\mu\boldsymbol{\zeta}^{(1)} + \boldsymbol{\zeta}^{(2)})$$

$$\boldsymbol{\zeta}^{(1)} = \sum_{t=1}^{T_1}\mathbf{y}_t\mathbf{y}_t' - \left(\sum_{t=1}^{T_1}\mathbf{y}_t\mathbf{x}_{t-1}'\right)\left(\sum_{t=1}^{T_1}\mathbf{x}_{t-1}\mathbf{x}_{t-1}'\right)^{-1}\left(\sum_{t=1}^{T_1}\mathbf{x}_{t-1}\mathbf{y}_t'\right)$$

$$\boldsymbol{\zeta}^{(2)} = \sum_{t=T_1+1}^{T_2}\mathbf{y}_t\mathbf{y}_t' - \left(\sum_{t=T_1+1}^{T_2}\mathbf{y}_t\mathbf{x}_{t-1}'\right)\left(\sum_{t=T_1+1}^{T_2}\mathbf{x}_{t-1}\mathbf{x}_{t-1}'\right)^{-1}\left(\sum_{t=T_1+1}^{T_2}\mathbf{x}_{t-1}\mathbf{y}_t'\right).$$

For example, if we put zero weight on the Minnesota prior for the lagged structural coefficients ($\mathbf{P} = \mathbf{0}$) this would amount to using as a prior for the second sample

$$\mathbf{b}_i|\mathbf{A}, \mathbf{D} \sim N(\mathbf{a}_i'\hat{\boldsymbol{\Phi}}^{(1)}, \mu^{-1}d_{ii}\mathbf{M}^{(1)})$$

$$\hat{\boldsymbol{\Phi}}^{(1)} = \left(\sum_{t=1}^{T_1}\mathbf{y}_t\mathbf{x}_{t-1}'\right)\left(\sum_{t=1}^{T_1}\mathbf{x}_{t-1}\mathbf{x}_{t-1}'\right)^{-1}$$

$$\mathbf{M}^{(1)} = \left(\sum_{t=1}^{T_1}\mathbf{x}_{t-1}\mathbf{x}_{t-1}'\right)^{-1}.$$

Thus the mean for the prior used to analyze the second sample ($\mathbf{a}_i'\hat{\boldsymbol{\Phi}}^{(1)}$) would be the coefficient from an OLS regression on the first sample. When $\mu = 1$ our confidence in this prior comes from the variance of the OLS regression estimate ($d_{ii}\mathbf{M}^{(1)}$), but the variance increases as μ decreases. As μ approaches 0 the variance of the prior goes to infinity and the information in the first sample would be completely ignored.

Likewise with no information about the structural variances other than the estimates from the first sample ($\kappa = \tau = 0$), the prior for the structural variances that we would use for the second sample would be

$$d_{ii}^{-1}|\mathbf{A} \sim \Gamma(\mu T_1, \mu(\mathbf{a}_i' \boldsymbol{\zeta}^{(1)} \mathbf{a}_i)).$$

Again the mean of this distribution is the first-sample OLS estimate ($T_1/(\mathbf{a}_i'\boldsymbol{\zeta}^{(1)}\mathbf{a}_i)$) but the variance goes to infinity as $\mu \to 0$.

5.2 Empirical results.

The solid red curves in Figure 9 denote the priors in $p(\mathbf{A})$ that we used for the contemporaneous coefficients. The posterior distribution with the pre-1975 observations downweighted by $\mu = 0.5$ are reported as blue histograms in Figure 9.

The posterior median of the short-run supply price elasticity, α_{qp}, is 0.16, a little higher than anticipated by our prior. Values less than 0.05 or greater than 0.5 are substantially less plausible after seeing the data than anticipated by our prior. The posterior median of the short-run demand price elasticity, β_{qp}, is -0.35, significantly more elastic than anticipated by our prior. But the data do not help us make a definitive statement about either parameter beyond telling us that values near zero can likely be ruled out for both parameters.

The data also do not give us much information about the short-run income elasticity of demand, β_{qy}. The fraction χ of world inventories that is captured by our proxy may be slightly larger than we anticipated, and ρ, which summarizes the importance of this measurement error, is a little smaller than anticipated. Our prior for α_{yp}, the short-run impact of oil prices on economic activity, forced this coefficient to be negative. The data evidently favor a positive value, causing the posterior to bunch near values just below zero.

Posterior structural impulse-response functions are plotted in Figure 10.[9] An oil supply shock (first row) lowers oil production and raises oil price on impact, whereas a shock to oil consumption demand (third row) raises production and raises price. These four effects were imposed by assumption as a result of the sign restrictions embedded in our priors. An oil supply shock also leads to a significant decline in economic activity, though only after a lag of a number of months, a result found in a large number of studies going back to Hamilton

[9] Note following standard practice these are accumulated impulse-response functions, plotting elements of $(\mathbf{H}_0 + \mathbf{H}_1 + \cdots + \mathbf{H}_s)$ as a function of s where $\mathbf{H}_s$ is given in (18). For example, panel (1,3) shows the effect on the level of oil prices s periods after an oil supply shock.

(1983). By contrast, if oil prices rise as a consequence of a shock to consumption demand, there seems to be no effect on subsequent economic activity. An increase in oil prices that results from an increase in inventory demand alone, which has sometimes been described as a speculative demand shock, seems to have a persistent effect on both inventories and prices and, after a long lag, a negative effect on economic activity as well.

Figure 11 shows the historical decomposition of oil price movements along with 95% credibility regions.[10] Supply shocks were the biggest factor in the oil price spike in 1990, whereas demand was more important in the price run-up in the first half of 2008. All four structural shocks contributed to the oil price drop in 2008:H2 and rebound in 2009, but apart from this episode, economic activity shocks and speculative demand shocks were usually not a big factor in oil price movements. Interestingly, demand is judged to be a little more important than supply in the price collapse of 2014:H2.[11]

Table 1 summarizes variance decompositions, reporting how much of the mean squared error associated with a 12-month-ahead forecast of each of the four variables in the system (represented by different columns of the table) is attributed to each of the four structural shocks (represented by different rows).[12] Oil supply shocks are the biggest single factor accounting for movements in oil production, while oil supply and oil consumption demand shocks are equally important in accounting for price changes. Consumption demand shocks are often buffered by inventories, but supply shocks much less so. Oil-market factors, as represented by the first, third, and fourth shocks combined, account for only 10% of the variation in world industrial production.

5.3 Sensitivity analysis.

The above results achieved partial identification by drawing on a large number of different sources of information. One benefit of using multiple sources is that we can examine the

[10] To our knowledge, the latter have never before been calculated in the large previous literature on sign-restricted vector autoregressions. The figure plots the contribution of the current and $s = 100$ previous structural shocks to the value of $\mathbf{y}_t$ for each date t plotted.

[11] In comparing Figures 10 and 11, recall that Figure 10 plots the response to a 1-unit increase in the specified structural shock, for example a shock to u_{2t}^* that increases economic activity by 1% or a shock to u_{3t}^* that increases real oil demand by 1%. The former corresponds almost to a 2-standard deviation event (the median posterior value of $(d_{22}^*)^{-0.5}$ is 0.54) while the latter is less than a third of one standard deviation (the median posterior value of $(d_{33}^*)^{-0.5}$ is 3.25). A 2-standard-deviation shock to u_{2t}^* has a bigger effect of p_{t+s} than a 1/3-standard-deviation shock to u_{3t}^* (as seen in the (2,3) and (3,3) panels of Figure 10) even though typically shocks to u_{3t}^* are more important than shocks to u_{2t}^* in accounting for fluctuations in prices (as seen in the second and third panels of Figure 11).

[12] Note that these refer to forecasts of the variables as defined in $\mathbf{y}_t$, e.g., 12-month-ahead forecasts of the rate of growth of oil prices as opposed to forecasts of the level of oil prices. At the 12-month horizon these are very close to decompositions of the unconditional variance of the monthly change in oil prices.

effects of putting less weight on any specified components of the prior to see how it affects the results.

Table 2 presents the posterior median and posterior 95% credibility sets for some of the magnitudes of interest under alternative specifications for the prior. The first column presents results from the baseline specification that were just summarized. Rows 1 and 2 report inference about the short-run supply elasticity α_{qp} and demand elasticity β_{qp}. Row 3 looks at the response of economic activity 12 months after a supply shock, row 4 the response to an oil consumption demand shock, and row 5 the response to a speculative demand shock, with each shock normalized for purposes of the table as an event that leads to a 1% increase in the real oil price at time 0. Note that this is a different normalization from that used in Figure 10, where the effect plotted was that of a one-unit change in the structural shock $\partial y_{i,t+s}/\partial u^*_{jt}$.[13] Rows 6-8 look at the variance decomposition for the error forecasting the real price of oil 12 months ahead, reporting the contribution to this variance of oil supply shocks, economic activity shocks, and oil consumption demand shocks, respectively.

Table 3 reports a few summary statistics for the historical decomposition. Column 2 reports the actual cumulative magnitude of the oil price change (as measured by the refiner acquisition cost) in three important episodes in the sample: the price collapse over January to July 1986, the price run-up over January to June 2008, and the price collapse from July to December 2014. Column 3 reports the posterior median and 95% credibility sets for the predicted change over that interval if the only structural shocks had come from the oil supply equation as inferred using the baseline prior.[14] According to the baseline specification, supply shocks accounted for about 46% of the price drop in 1986 but only 32% of the price drop in 2014:H2. The 95% credibility interval for the latter does not exceed 63%– factors in addition to strong supply growth clearly contributed to the most recent oil price decline. Demand factors were the single most important factor in the first half of 2008, accounting for 54% of the run-up in prices observed at that time.

We next explored the consequences of using a much weaker prior for the short-run supply and demand elasticities, replacing the scale parameter of $\sigma^{\alpha}_{qp} = \sigma^{\beta}_{qp} = 0.2$ that was used in the baseline analysis with the values $\sigma^{\alpha}_{qp} = \sigma^{\beta}_{qp} = 1.0$. This change gives the prior for these two parameters a variance that is 25 times larger than in the baseline specification. The

[13] Calculating the size of a shock that in equilibrium raises the price of oil by 1% potentially requires making use of all the parameters in $\mathbf{A}$ and $\mathbf{D}$. For this reason, credibility sets for the magnitudes defined in rows 3-5 of Table 2 are typically wider than those for the magnitudes plotted in Figure 10.

[14] For example, column 2 reports the third element of $\mathbf{y}_{t_1} + \mathbf{y}_{t_1-1} + \cdots + \mathbf{y}_{t_0}$ for $t_0 =$ 1986:M1 and $t_1 =$ 1986:M7 while column 3 reports the median and 95% credibility interval across draws of $\boldsymbol{\theta}^{(\ell)}$ for the third element of $\boldsymbol{\xi}_{j,t_1,t_0}(\mathbf{Y}_T;\boldsymbol{\theta}^{(\ell)}) + \boldsymbol{\xi}_{j,t_1-1,t_0}(\mathbf{Y}_T;\boldsymbol{\theta}^{(\ell)}) + \cdots + \boldsymbol{\xi}_{j,t_0,t_0}(\mathbf{Y}_T;\boldsymbol{\theta}^{(\ell)})$ for $\boldsymbol{\xi}_{j,t,t-s}(\mathbf{Y}_T;\boldsymbol{\theta})$ calculated from (21).

implications of these changes for inference about key parameters are reported in column 2 of Table 2. If we had very little prior information about the elasticities themselves, we would tend to infer a smaller short-run supply elasticity (row 1) and more elastic demand (row 2). Our key conclusions about impulse responses (rows 3-5) would change very little. If we relied less on prior information about the supply and demand elasticities, our point estimates would imply a smaller role for supply shocks and a bigger role for demand shocks (rows 6 and 8 of Table 2). Note that credibility sets for the contributions to individual episodes become significantly wider if we make less use of information about supply and demand elasticities (see column 4 of Table 3). We nevertheless would still have 95% posterior confidence that supply shocks contributed no more than 66% to the oil price decline in the second half of 2014.

Our prior beliefs about the role of measurement error were represented by the Beta$(\alpha_\chi, \beta_\chi)$ distribution for χ (which summarizes the ratio of OECD inventories to world total) and Beta$(\alpha_\rho, \beta_\rho)$ distribution for ρ (which summarizes the component of the correlation between price and inventory changes that is attributed to measurement error). Our baseline specification used $\alpha_\chi = 15$, $\beta_\chi = 10$, $\alpha_\rho = 3$, $\beta_\rho = 9$, which imply standard deviations for the priors of 0.1 and 0.12, respectively. In our less informative alternative specification we take $\alpha_\chi = 1.5$, $\beta_\chi = 1$, $\alpha_\rho = 1$, $\beta_\rho = 3$, whose standard deviations are 0.26 and 0.19, respectively (recall that each variable by definition is known to fall between 0 and 1). The implications of this weaker prior about the role of measurement error are reported in column 3 of Table 2 and column 5 of Table 3. These results are uniformly very close to those for our baseline specification.

Next we examined the consequences of paying less attention to data prior to 1975. Our baseline specification set $\mu = 0.5$, which gives pre-1975 data half the weight of the more recent data. Our less informative alternative uses $\mu = 0.25$, thus regarding the earlier data as only 1/4 as important as the more recent numbers. The inferences are virtually identical to those under our baseline specification.

The role of prior information about lagged structural coefficients $\mathbf{b}_i$ is summarized by the value of λ_0 in (43). An increase in λ_0 increases the variance on all the priors involving the lagged coefficients. Our baseline specification took $\lambda_0 = 0.5$, whereas the weaker value of $\lambda_0 = 1$ (which implies a variance 4 times as large) was used in column 5 of Table 2 and column 7 of Table 3. This has only a modest effect on any of the inferences.

The weight on prior information about structural variances is captured by the value of κ in equation (5). Our baseline specification used $\kappa = 2$, which gives the prior a weight of about 4 observations. The alternative in column 6 of Table 2 and column 8 of Table 3 uses $\kappa = 0.5$. This does not matter for any of the conclusions.

Finally, in making use of the historical data we relied on West Texas Intermediate as our measure of the crude oil price prior to 1975, since the refiner acquisition cost is unavailable. But our baseline analysis nevertheless used RAC as the oil price measure since 1975. An alternative is to use WTI for both samples. Column 9 of Table 3 reports the measured size of the oil price change recorded by WTI in three episodes of interest. The two oil price measures can give quite different answers for the size of the move in any given month. Nevertheless, the inference about key model parameters (column 7 of Table 2) is the same regardless of which oil price measure we use.

6 Conclusion.

Prior information has played a key role in every structural analysis of vector autoregressions that has ever been done. Typically prior information has been treated as "all or nothing", which from a Bayesian perspective would be described as either dogmatic priors (details that the analyst claims to know with certainty before seeing the data) or completely uninformative priors. In this paper we noted that there is vast middle ground between these two extremes. We advocate that analysts should both relax the dogmatic priors, acknowledging that we have some uncertainty about the identifying assumptions themselves, and strengthen the uninformative priors, drawing on whatever information may be known outside of the dataset being analyzed.

We illustrated these concepts by revisiting the role of supply and demand shocks in oil markets. We demonstrated how previous studies can be viewed as a special case of Bayesian inference and proposed a generalization that draws on a rich set of prior information beyond the data being analyzed while simultaneously relaxing some of the dogmatic priors implicit in traditional identification. Notwithstanding, we end up confirming some of the core conclusions of earlier studies. We find that oil price increases that result from supply shocks lead to a reduction in economic activity after a significant lag, whereas price increases that result from increases in oil consumption demand do not have a significant effect on economic activity. We also examined the sensitivity of our results to the priors used, and found that many of the key conclusions change very little when substantially less weight is placed on various components of the prior.

References

Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang (2015). "Measurement Errors and Monetary Policy: Then and Now," working paper, Goethe University Frankfurt.

Baumeister, Christiane, and James D. Hamilton (2015a). "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," *Econometrica* 83: 1963–1999.

Baumeister, Christiane, and James D. Hamilton (2015b). "Optimal Inference about Impulse-Response Functions and Historical Decompositions in Incompletely Identified Structural Vector Autoregressions," working paper, UCSD.

Baumeister, Christiane, and Gert Peersman (2013a). "The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market," *Journal of Applied Econometrics* 28: 1087-1109.

Baumeister, Christiane, and Gert Peersman (2013b). "Time-Varying Effects of Oil Supply Shocks on the US Economy," *American Economic Journal: Macroeconomics* 5: 1-28.

Brons, Martijn, Peter Nijkamp, Eric Pels, and Piet Rietveld (2008). "A Meta-Analysis of the Price Elasticity of Gasoline Demand: A SUR Approach," *Energy Economics* 30: 2105-2122.

Cogley, Timothy, and Thomas J. Sargent (2015), "Measuring Price-Level Uncertainty and Instability in the U.S., 1850-2012," *Review of Economics and Statistics* 97: 827-838.

Coglianese, John, Lucas W. Davis, Lutz Kilian, and James H. Stock (forthcoming). "Anticipation, Tax Avoidance, and the Price Elasticity of Demand," *Journal of Applied Econometrics*.

Csereklyei, Zsuzsanna, Maria del Mar Rubio Varas and David Stern (2016). "Energy and Economic Growth: The Stylized Facts," *Energy Journal* 37: 223-255.

Dahl, Carol A. (1993). "A Survey of Oil Demand Elasticities for Developing Countries," *OPEC Review* 17: 399-419.

Dahl, Carol and Thomas Sterner (1991). "Analysing Gasoline Demand Elasticities: A Survey," *Energy Economics* 13: 203-210.

Darmstadter, Joel, Joy Dunkerly, and Jack Alterman (1977). *How Industrial Societies Use Energy: A Comparative Analysis*, Baltimore: Johns Hopkins University Press.

Davis, Lucas W. (2014). "The Economic Cost of Global Fuel Subsidies," *American Economic Review: Papers and Proceedings* 104: 581-585.

Doan, Thomas, Robert B. Litterman, and Christopher A. Sims (1984). "Forecasting and Conditional Projection Using Realistic Prior Distributions," *Econometric Reviews* 3: 1-100.

Doan, Thomas (2013). *RATS User's Guide, Version 8.2,* www.estima.com.

Drollas, Leonidas (1984). "The Demand for Gasoline: Further Evidence," *Energy Economics* 6: 71-83.

Espey, Molly (1998). "Gasoline Demand Revisited: An International Meta-Analysis of Elasticities," *Energy Economics* 20: 273-295.

Gately, Dermot, and Hillard G. Huntington (2002). "The Asymmetric Effects of Changes in Price and Income on Energy and Oil Demand," *Energy Journal* 23: 19-55.

Graham, Daniel J., and Stephen Glaister (2004). "Road Traffic Demand Elasticity Estimates: A Review," *Transport Reviews* 24: 261-274.

Hamilton, James D. (1983). "Oil and the Macroeconomy Since World War II," *Journal of Political Economy*, 91(2): 228-248.

Hamilton, James D. (1994). *Time Series Analysis.* Princeton: Princeton University Press.

Hamilton, James D. (2013). "Oil Prices, Exhaustible Resources, and Economic Growth," in *Handbook on Energy and Climate Change,* pp. 29-63, edited by Roger Fouqet, Cheltenham, United Kingdom: Edward Elgar Publishing.

Hausman, Jerry A. and Whitney K. Newey (1995). "Nonparametric Estimation of Exact Consumers Surplus and Deadweight Loss," *Econometrica* 63: 1445-1476.

Kilian, Lutz (2009). "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review* 99: 1053-1069.

Kilian, Lutz, and Daniel P. Murphy (2012). "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models," *Journal of the European Economic Association* 10(5): 1166-1188.

Kilian, Lutz, and Daniel P. Murphy (2014). "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil," *Journal of Applied Econometrics* 29: 454-478.

Kilian, Lutz, and Robert J. Vigfusson (2011). "Nonlinearities in the Oil Price-Output Relationship," *Macroeconomic Dynamics* 15(Supplement 3): 337–363.

Sims, Christopher A., and Tao Zha (1998). "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review* 39: 949-968.

Yatchew, Adonis and Joungyeo Angela No (2001). "Household Gasoline Demand in Canada," *Econometrica* 69: 1697-1709.

Appendix

A. Reference priors for D and B.

Prior for $\mathbf{D}|\mathbf{A}$. Prior beliefs about structural variances should reflect in part the scale of the underlying data. Let $\hat{e}_{it}$ denote the residual of an mth-order univariate autoregression fit to series i and $\mathbf{S}$ the sample variance matrix of these univariate residuals ($s_{ij} = T^{-1}\sum_{t=1}^{T}\hat{e}_{it}\hat{e}_{jt}$). Baumeister and Hamilton (2015a) proposed setting κ_i/τ_i (the prior mean for d_{ii}^{-1}) equal to the reciprocal of the ith diagonal element of $\mathbf{ASA}'$; in other words, $\tau_i(\mathbf{A}) = \kappa_i \mathbf{a}_i'\mathbf{S}\mathbf{a}_i$. Given equation (5), the prior carries a weight equivalent to $2\kappa_i$ observations of data; for example, setting $\kappa_i = 2$ would give the prior as much weight as 4 observations.

Prior for $\mathbf{B}|\mathbf{A},\mathbf{D}$. A standard prior for many data sets suggested by Doan, Litterman and Sims (1984) is that individual series behave like random walks. Baumeister and Hamilton (2015a, equation 45) adapted Sims and Zha's (1998) method for representing this in terms of a particular specification for $\mathbf{m}_i(\mathbf{A})$. For other data sets, such as the one analyzed in Section 5, a more natural prior is that series behave like white noise ($\mathbf{m}_i = \mathbf{0}$). For either case, we recommend following Doan, Litterman and Sims (1984) in placing greater confidence in our expectation that coefficients on higher lags are zero, implemented by using smaller values for the diagonal elements for $\mathbf{M}_i$ associated with higher lags. Define

$$\underset{(1\times m)}{\mathbf{v}_1'} = \left(1/(1^{2\lambda_1}), 1/(2^{2\lambda_1}), ..., 1/(m^{2\lambda_1})\right) \tag{42}$$

$$\underset{(1\times n)}{\mathbf{v}_2'} = (s_{11}^{-1}, s_{22}^{-1}, ..., s_{nn}^{-1})'$$

$$\mathbf{v}_3 = \lambda_0^2 \begin{bmatrix} \mathbf{v}_1 \otimes \mathbf{v}_2 \\ \lambda_3^2 \end{bmatrix}. \tag{43}$$

Then $\mathbf{M}_i$ is taken to be a diagonal matrix whose (r,r) element is the rth element of $\mathbf{v}_3$:

$$M_{i,rr} = v_{3r}. \tag{44}$$

Here λ_0 summarizes the overall confidence in the prior (with smaller λ_0 corresponding to greater weight given to the prior), λ_1 governs how much more confident we are that higher coefficients are zero (with a value of $\lambda_1 = 0$ giving all lags equal weight), and λ_3 is a separate parameter governing the tightness of the prior for the constant term, with all $\lambda_k \geq 0$.

Doan (2013) discussed possible values for these parameters. For the baseline specification in Section 5 we set $\lambda_1 = 1$ (which governs how quickly the prior for lagged coefficients tightens to zero as the lag ℓ increases), $\lambda_3 = 100$ (which makes the prior on the constant term essentially irrelevant), and set λ_0, the parameter controlling the overall tightness of the prior, to 0.5.

B. Details of Bayesian algorithm.

For any numerical value of $\mathbf{A}$ we can calculate $\zeta_i^*(\mathbf{A})$ and $\tau_i^*(\mathbf{A})$ using equations (7) and (6) from which we can calculate the log of the target

$$q(\mathbf{A}) = \log(p(\mathbf{A})) + (T/2)\log\left[\det\left(\mathbf{A}\hat{\boldsymbol{\Omega}}_T\mathbf{A}'\right)\right] \\ - \textstyle\sum_{i=1}^n \kappa_i^* \log[(2/T)\tau_i^*(\mathbf{A})] + \sum_{i=1}^n \kappa_i\tau_i(\mathbf{A}). \tag{45}$$

We can improve the efficiency of the algorithm by using information about the shape of this function calculated as follows. Collect elements of $\mathbf{A}$ that are not known with certainty in an $(n_\alpha \times 1)$ vector $\boldsymbol{\alpha}$, and find the value $\hat{\boldsymbol{\alpha}}$ that maximizes (45) numerically. This value $\hat{\boldsymbol{\alpha}}$ offers a reasonable guess for the posterior mean of $\boldsymbol{\alpha}$, while the matrix of second derivatives (again obtained numerically) gives an idea of the curvature of the posterior distribution:

$$\hat{\boldsymbol{\Lambda}} = -\left.\frac{\partial^2 q(\mathbf{A}(\boldsymbol{\alpha}))}{\partial\boldsymbol{\alpha}\partial\boldsymbol{\alpha}'}\right|_{\boldsymbol{\alpha}=\hat{\boldsymbol{\alpha}}}.$$

We then use this guess to inform a random-walk Metropolis Hastings algorithm to generate candidate draws of $\boldsymbol{\alpha}$ from the posterior distribution, as follows. As a result of step ℓ we have generated a value of $\boldsymbol{\alpha}^{(\ell)}$. For step $\ell+1$ we generate

$$\tilde{\boldsymbol{\alpha}}^{(\ell+1)} = \boldsymbol{\alpha}^{(\ell)} + \xi\left(\hat{\mathbf{Q}}^{-1}\right)'\mathbf{v}_t$$

for $\mathbf{v}_t$ an $(n_\alpha \times 1)$ vector of Student t variables with 2 degrees of freedom, $\hat{\mathbf{Q}}$ the Cholesky factor of $\hat{\boldsymbol{\Lambda}}$ (namely $\hat{\mathbf{Q}}\hat{\mathbf{Q}}' = \hat{\boldsymbol{\Lambda}}$ with $\hat{\mathbf{Q}}$ lower triangular), and ξ a tuning scalar to be described shortly. If $q(\mathbf{A}(\tilde{\boldsymbol{\alpha}}^{(\ell+1)})) < q(\mathbf{A}(\boldsymbol{\alpha}^{(\ell)}))$, we set $\boldsymbol{\alpha}^{(\ell+1)} = \boldsymbol{\alpha}^{(\ell)}$ with probability $1 - \exp\left[q(\mathbf{A}(\tilde{\boldsymbol{\alpha}}^{(\ell+1)})) - q(\mathbf{A}(\boldsymbol{\alpha}^{(\ell)}))\right]$; otherwise, we set $\boldsymbol{\alpha}^{(\ell+1)} = \tilde{\boldsymbol{\alpha}}^{(\ell+1)}$. The parameter ξ is chosen so that about 30% of the newly generated $\tilde{\boldsymbol{\alpha}}^{(\ell+1)}$ get retained. The algorithm can be started by setting $\boldsymbol{\alpha}^{(1)} = \hat{\boldsymbol{\alpha}}$, and the values after the first D burn-in draws $\{\boldsymbol{\alpha}^{(D+1)}, \boldsymbol{\alpha}^{(D+2)}, ..., \boldsymbol{\alpha}^{(D+N)}\}$ represent a sample of size N drawn from the posterior distribution $p(\boldsymbol{\alpha}|\mathbf{Y}_T)$; in our applications we have used $D = N = 10^6$.

For each of these N final values for $\boldsymbol{\alpha}^{(\ell)}$ we further generate $\delta_{ii}^{(\ell)} \sim \Gamma(\kappa_i^*, \tau_i^*(\mathbf{A}(\boldsymbol{\alpha}^{(\ell)})))$ for $i = 1, ..., n$ and take $\mathbf{D}^{(\ell)}$ to be a diagonal matrix whose row i, column i element is given by $1/\delta_{ii}^{(\ell)}$. From these we also generate $\mathbf{b}_i^{(\ell)} \sim N(\mathbf{m}_i^*(\mathbf{A}(\boldsymbol{\alpha}^{(\ell)})), d_{ii}^{(\ell)}\mathbf{M}_i^*)$ for $i = 1, ..., n$ and take $\mathbf{B}^{(\ell)}$ the matrix whose ith row is given by $\mathbf{b}_i^{(\ell)\prime}$. The triple $\{\mathbf{A}(\boldsymbol{\alpha}^{(\ell)}), \mathbf{D}^{(\ell)}, \mathbf{B}^{(\ell)}\}_{\ell=D+1}^{D+N}$ then represents a sample of size N drawn from the posterior distribution $p(\mathbf{A}, \mathbf{D}, \mathbf{B}|\mathbf{Y}_T)$.

C. Data sources.

The data sets used in the original studies by Kilian (2009) and Kilian and Murphy (2012) are available from the public data archives of the *Journal of the European Economic Association (*http://onlinelibrary.wiley.com/doi/10.1111/j.1542-4774.2012.01080.x/suppinfo) and we used these exact same data for the statistical analysis reported in Sections 3.1 and 3.2. We also reconstructed these data sets from the original sources ourselves as part of the process of assembling extended time series as described below.

Monthly world oil production data measured in thousands of barrels of oil per day were obtained from the U.S. Energy Information Administration's (EIA) *Monthly Energy Review* for the period January 1973 to December 2014. Monthly data for global production of crude oil for the period 1958:M1 to 1972:M12 were collected from the weekly *Oil and Gas Journal* (issue of the first week of each month) as in Baumeister and Peersman (2013b).

The nominal spot oil price for West Texas Intermediate (WTI) was retrieved from the Federal Reserve Economic Data (FRED) database maintained by the St. Louis FED (OILPRICE). Prior to 1982 this equals the posted price. This series was discontinued in July 2013. From August 2013 onwards data are obtained from the EIA website (http://www.eia.gov/dnav/pet/hist/LeafHandler.ashx?n=pet&s=rwtc&f=m). To deflate the nominal spot oil price, we use the U.S. consumer price index (CPIAUCSL: consumer price index for all urban consumers: all items, index 1982-1984 = 100) which was taken from the FRED database.

For the extended data set our measure for global economic activity is the industrial production index for OECD countries and six major non-member economies (Brazil, China, India, Indonesia, the Russian Federation and South Africa) obtained from the OECD Main Economic Indicators (MEI) database in 2011. The index covers the period 1958:M1 to 2011:M10 and was subsequently discontinued. To extend the data set after October 2011, we applied the same methodology used by the OCED. Specifically, we use OECD industrial production and industrial production for the individual non-member countries which are available in the MEI database and apply the weights reported by the OECD to aggregate those series into a single index. The source of the weights data is the International Monetary Fund's World Economic Outlook (WEO) database. The weights are updated on a yearly basis and a link to a document containing the weights can be found at http://www.oecd.org/std/compositeleadingindicatorsclifrequentlyaskedquestionsfaqs.htm#11.

Monthly U.S. crude oil stocks in millions of barrels (which include the Stragegic Petroleum Reserve) are available from EIA for the entire period 1958:M1-2014:M12. We obtain an estimate for global stocks as in Kilian and Murphy (2012) by multiplying the U.S. crude oil inventories by the ratio of OECD inventories of crude petroleum and petroleum products to

U.S. inventories of petroleum and petroleum products. Given that OECD petroleum inventories only start in January 1988, we assume that the ratio before January 1988 is the same as in January 1988. To calculate our proxy for Δi_t, the change in OECD inventories as a fraction of last period's oil production, we convert the production data into millions of barrels per month by multiplying the million barrels of crude oil produced per day by 30.

D. Using downweighted observations from an earlier sample.

Let $\mathbf{Y}^{(1)}$ denote observations from the first sample, $\mathbf{Y}^{(2)}$ observations from the second and $\boldsymbol{\lambda}$ the vector of parameters about which we wish to form an inference. If both samples are regarded as equally informative about $\boldsymbol{\lambda}$, the posterior would be calculated as

$$p(\boldsymbol{\lambda}|\mathbf{Y}^{(1)},\mathbf{Y}^{(2)}) = \frac{p(\mathbf{Y}^{(2)}|\mathbf{Y}^{(1)},\boldsymbol{\lambda})p(\mathbf{Y}^{(1)}|\boldsymbol{\lambda})p(\boldsymbol{\lambda})}{\int p(\mathbf{Y}^{(2)}|\mathbf{Y}^{(1)},\boldsymbol{\lambda})p(\mathbf{Y}^{(1)}|\boldsymbol{\lambda})p(\boldsymbol{\lambda})d\boldsymbol{\lambda}}. \tag{46}$$

Define $p(\mathbf{Y}^{(1)}) = \int p(\mathbf{Y}^{(1)}|\boldsymbol{\lambda})p(\boldsymbol{\lambda})d\boldsymbol{\lambda}$. Then the posterior density based on the first sample alone would be

$$p(\boldsymbol{\lambda}|\mathbf{Y}^{(1)}) = \frac{p(\mathbf{Y}^{(1)}|\boldsymbol{\lambda})p(\boldsymbol{\lambda})}{p(\mathbf{Y}^{(1)})}.$$

Dividing numerator and denominator of (46) by $p(\mathbf{Y}^{(1)})$ we see that the full-sample posterior could equivalently be obtained by using the posterior from the first sample as the prior for the second:

$$p(\boldsymbol{\lambda}|\mathbf{Y}^{(1)},\mathbf{Y}^{(2)}) = \frac{p(\mathbf{Y}^{(2)}|\mathbf{Y}^{(1)},\boldsymbol{\lambda})p(\boldsymbol{\lambda}|\mathbf{Y}^{(1)})}{\int p(\mathbf{Y}^{(2)}|\mathbf{Y}^{(1)},\boldsymbol{\lambda})p(\boldsymbol{\lambda}|\mathbf{Y}^{(1)})d\boldsymbol{\lambda}}.$$

We propose instead to use as a prior for the second sample a distribution that downweights the influence of the data from the first sample,

$$p(\boldsymbol{\lambda}|\mathbf{Y}^{(1)}) \propto \left[p(\mathbf{Y}^{(1)}|\boldsymbol{\lambda})\right]^{\mu} p(\boldsymbol{\lambda})$$

for some $0 \leq \mu \leq 1$. In the present instance the likelihood for the first sample is given by

$$p(\mathbf{Y}^{(1)}|\boldsymbol{\lambda}) = (2\pi)^{-T_1 n/2}|\det(\mathbf{A})|^{T_1}|\mathbf{D}|^{-T_1/2}\prod_{i=1}^{n} \exp\left[-\sum_{t=1}^{T_1}\frac{(\mathbf{a}_i'\mathbf{y}_t - \mathbf{b}_i'\mathbf{x}_{t-1})^2}{2d_{ii}}\right]$$

so the downweighted first-sample likelihood is

$$p(\mathbf{Y}^{(1)}|\boldsymbol{\lambda})^{\mu} = (2\pi)^{-\mu T_1 n/2}|\det(\mathbf{A})|^{\mu T_1}|\mathbf{D}|^{-\mu T_1/2}\prod_{i=1}^{n} \exp\left[-\sum_{t=1}^{T_1}\frac{(\mathbf{a}_i'\sqrt{\mu}\mathbf{y}_t - \mathbf{b}_i'\sqrt{\mu}\mathbf{x}_{t-1})^2}{2d_{ii}}\right].$$

Repeating the derivations in Baumeister and Hamilton (2015a) for this downweighted likelihood leads to the algorithm described in Section 5.1.4.

Table 1. Decomposition of variance of 12-month-ahead forecast errors.

	Oil production	Industrial production	Oil price	Inventories
Oil supply	**1.61** [64%] (0.76, 2.17)	**0.01** [3%] (0.00, 0.03)	**17.23** [36%] (6.46, 34.55)	**0.07** [4%] (0.03, 0.12)
Economic activity	**0.09** [3%] (0.04, 0.15)	**0.34** [90%] (0.30, 0.40)	**2.16** [4%] (0.88, 4.44)	**0.05** [3%] (0.02, 0.09)
Oil consumption demand	**0.46** [19%] (0.18, 0.84)	**0.02** [4%] (0.01, 0.03)	**20.16** [42%] (6.91, 35.50)	**0.57** [38%] (0.25, 0.96)
Speculative oil demand	**0.35** [14%] (0.12, 0.96)	**0.01** [3%] (0.00, 0.02)	**8.86** [18%] (4.98, 14.24)	**0.82** [54%] (0.47, 1.16)

Notes to Table 1. Table reports the estimated contribution of each shock to the 12-month-ahead mean squared forecast error of each variable in bold, and expressed as a percent of total MSE in brackets. Parentheses indicate 95% credibility intervals.

Table 2. Sensitivity of parameter inference when less weight is placed on various components of the prior.

Benchmark	Supply and demand elasticities	Measure-ment error	Pre-1975 data	Lagged structural coefficients	Variances of shocks	Replace RAC with WTI
(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) *Short-run price elasticity of oil supply* α_{qp}						
0.16	**0.12**	**0.14**	**0.16**	**0.16**	**0.16**	**0.15**
(0.07, 0.38)	(0.04, 0.36)	(0.06, 0.35)	(0.07, 0.38)	(0.08, 0.36)	(0.07, 0.37)	(0.07, 0.45)
(2) *Short-run price elasticity of oil demand* β_{qp}						
-0.35	**-0.45**	**-0.39**	**-0.35**	**-0.34**	**-0.35**	**-0.30**
(-0.71, -0.16)	(-0.98, -0.17)	(-0.83, -0.17)	(-0.70, -0.16)	(-0.63, -0.16)	(-0.71, -0.16)	(-0.63, -0.09)
(3) *Effect of oil supply shock that raises real oil price by 1% on economic activity 12 months later*						
-0.06	**-0.07**	**-0.07**	**-0.05**	**-0.07**	**-0.06**	**-0.07**
(-0.15, 0.00)	(-0.19, 0.00)	(-0.17, 0.00)	(-0.14, 0.02)	(-0.16, 0.00)	(-0.15, 0.00)	(-0.15, -0.01)
(4) *Effect of oil consumption demand shock that raises real oil price by 1% on economic activity 12 months later*						
0.02	**0.01**	**0.02**	**0.02**	**0.01**	**0.02**	**0.01**
(-0.04, 0.12)	(-0.05, 0.10)	(-0.04, 0.12)	(-0.04, 0.12)	(-0.05, 0.11)	(-0.04, 0.11)	(-0.05, 0.11)
(5) *Effect of oil inventory demand shock that raises real oil price by 1% on economic activity 12 months later*						
-0.04	**-0.05**	**-0.05**	**-0.02**	**-0.07**	**-0.04**	**-0.05**
(-0.14, 0.05)	(-0.17, 0.04)	(-0.18, 0.05)	(-0.12, 0.07)	(-0.17, 0.03)	(-0.14, 0.05)	(-0.14, 0.04)
(6) *12-month-ahead MSE of oil price due to oil supply shock*						
17.23 [36%]	**13.05** [27%]	**15.33** [32%]	**17.10** [35%]	**17.42** [36%]	**17.18** [35%]	**19.15** [35%]
(6.46, 34.55)	(4.26, 33.36)	(5.20, 32.89)	(6.42, 34.83)	(7.54, 33.48)	(6.51, 34.07)	(7.66, 42.43)
(7) *12-month-ahead MSE of oil price due to economic activity shock*						
2.16 [4%]	**2.18** [5%]	**2.13** [4%]	**2.32** [5%]	**2.42** [5%]	**2.16** [4%]	**2.72** [5%]
(0.88, 4.44)	(0.90, 4.50)	(0.86, 4.42)	(0.93, 4.86)	(1.05, 4.83)	(0.89, 4.45)	(1.22, 5.14)
(8) *12-month-ahead MSE of oil price due to oil consumption demand shock*						
20.16 [42%]	**24.85** [51%]	**22.91** [48%]	**20.60** [42%]	**20.05** [41%]	**20.21** [42%]	**22.78** [41%]
(6.91, 35.50)	(7.41, 39.95)	(6.27, 38.63)	(7.18, 35.90)	(7.70, 33.98)	(7.15, 35.34)	(5.71, 39.30)

Notes to Table 2. Table reports posterior median (in bold) and 95% credibility regions (in parentheses) for indicated magnitudes. Baseline uses priors specified in Sections 5.1 and 5.2. Alternatives put less weight on indicated component of the prior as detailed in the text.

Table 3. Sensitivity of importance attributed to structural shocks in historical decomposition of oil prices when less weight is placed on various components of the prior.

Historical episode	Actual real oil price growth (RAC)	Benchmark	Supply and demand elasticities	Measurement error	Pre-1975 data	Lagged structural coefficients	Variances of shocks	Actual real oil price growth (WTI)	Replace RAC with WTI
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Jan-July	-87.65	**-40.29** [46%]	**-27.90** [32%]	**-40.38** [46%]	**-40.11** [46%]	**-41.85** [48%]	**-40.62** [46%]	-85.54	**-40.59** [47%]
1986 (*supply*)		(-70.54, -15.31)	(-73.12, -5.29)	(-73.04, -16.89)	(-70.24, -16.28)	(-71.54, -17.22)	(-72.48, -16.33)		(-71.51, -17.50)
Jan-June	38.86	**21.00** [54%]	**27.82** [72%]	**21.20** [55%]	**22.47** [58%]	**20.42** [53%]	**20.86** [54%]	35.04	**19.57** [56%]
2008 (*consumption*)		(7.53, 37.72)	(7.22, 45.60)	(2.95, 38.57)	(8.35, 39.22)	(7.47, 36.45)	(6.97, 37.15)		(6.65, 34.29)
July-Dec	-55.55	**-17.54** [32%]	**-11.23** [20%]	**-17.57** [32%]	**-17.71** [32%]	**-17.14** [31%]	**-17.72** [32%]	-57.45	**-19.99** [35%]
2014 (*supply*)		(-35.17, -5.15)	(-36.53, -1.10)	(-36.75, -5.80)	(-35.08, -5.80)	(-34.39, -5.29)	(-36.43, -5.61)		(-39.79, -7.27)

Notes to Table 3. Table reports posterior median (in bold) and 95% credibility regions (in parentheses) for indicated magnitudes. Baseline uses priors specified in Sections 5.1 and 5.2. Alternatives put less weight on indicated component of the prior as detailed in the text.

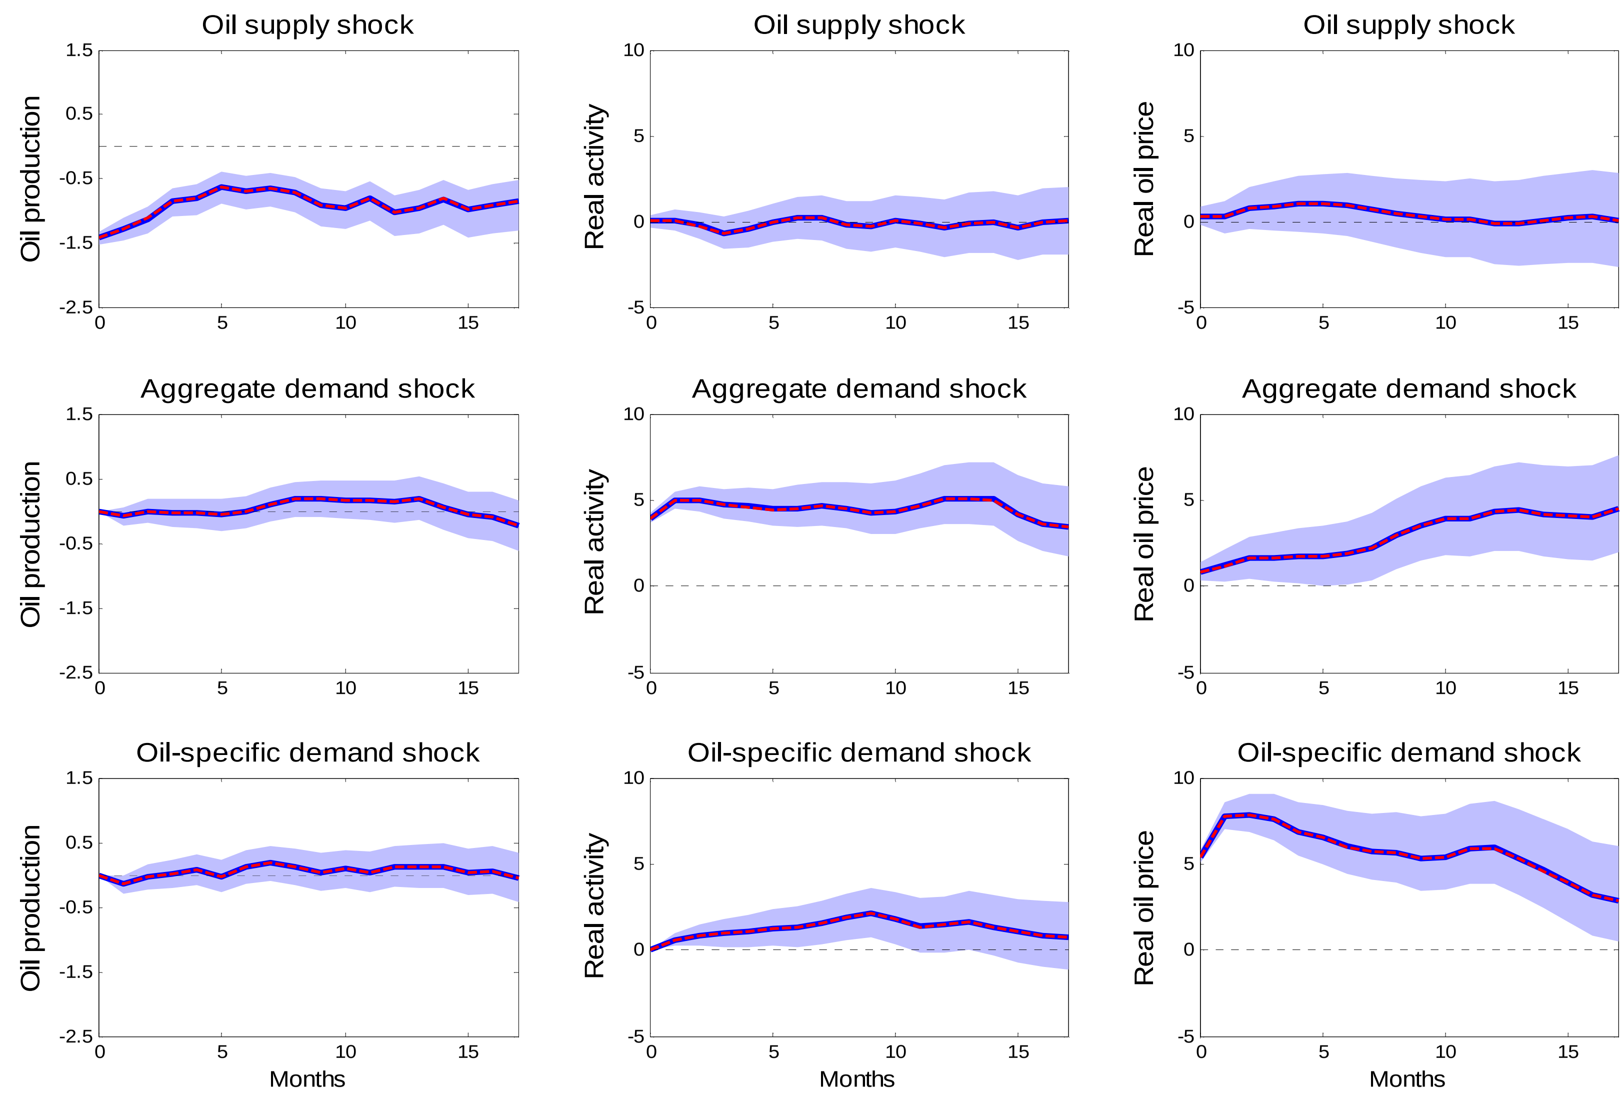

Figure 1. Impulse-response functions for 3-variable model under traditional Cholesky identification. Red dashed lines: point estimates arrived at using Kilian's (2009) original methodology; blue solid lines: median of Bayesian posterior distribution; shaded regions: 95% posterior credible set.

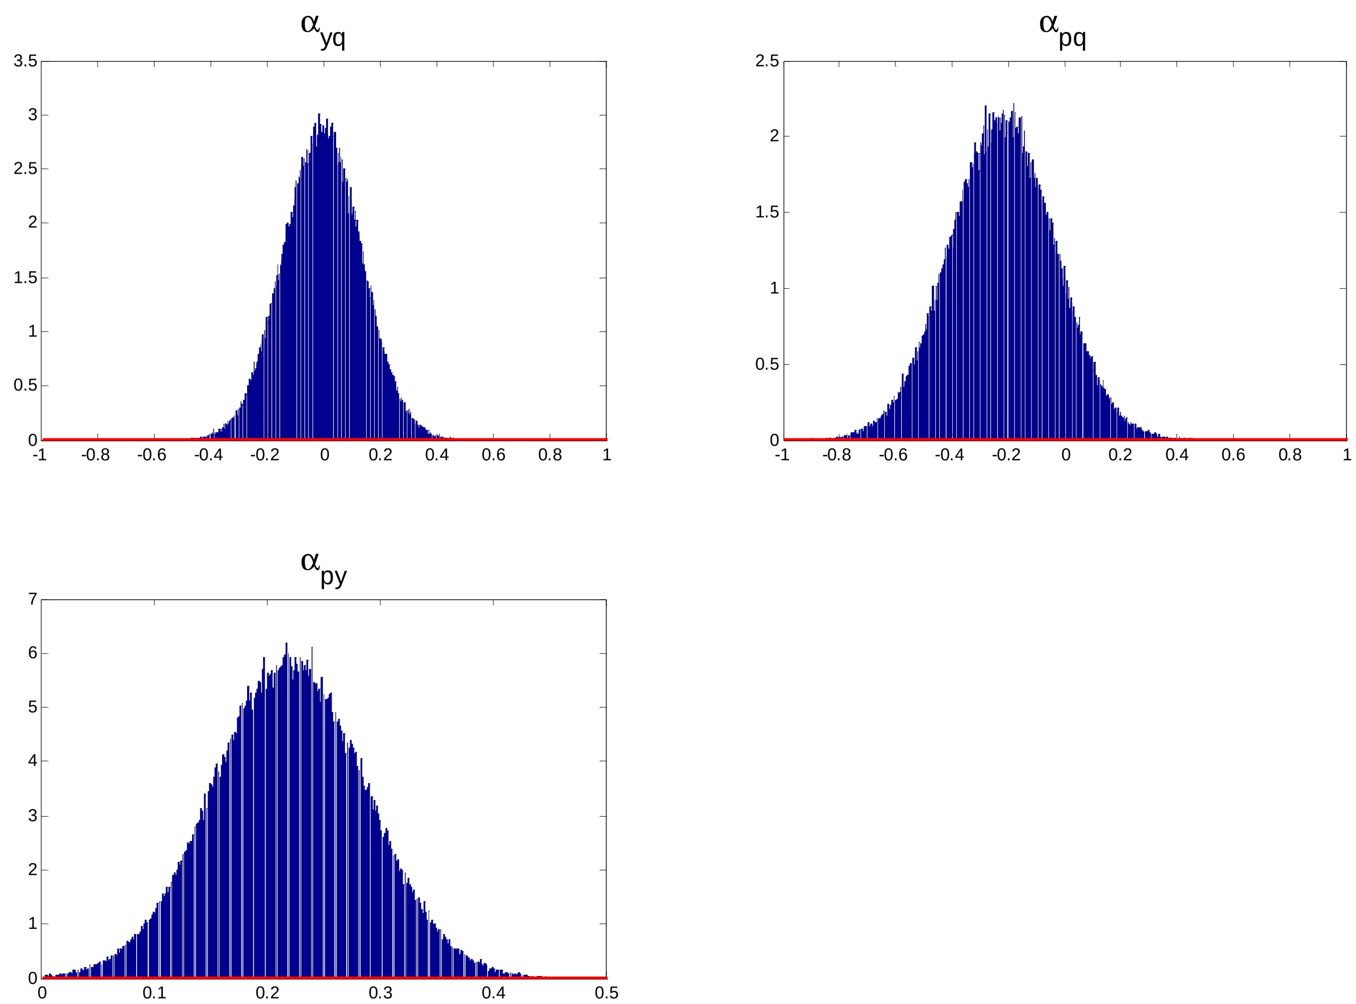

Figure 2. Prior (red lines) and posterior (blue histograms) distributions for the unknown elements in **A** using traditional 3-variable Cholesky-type identification.

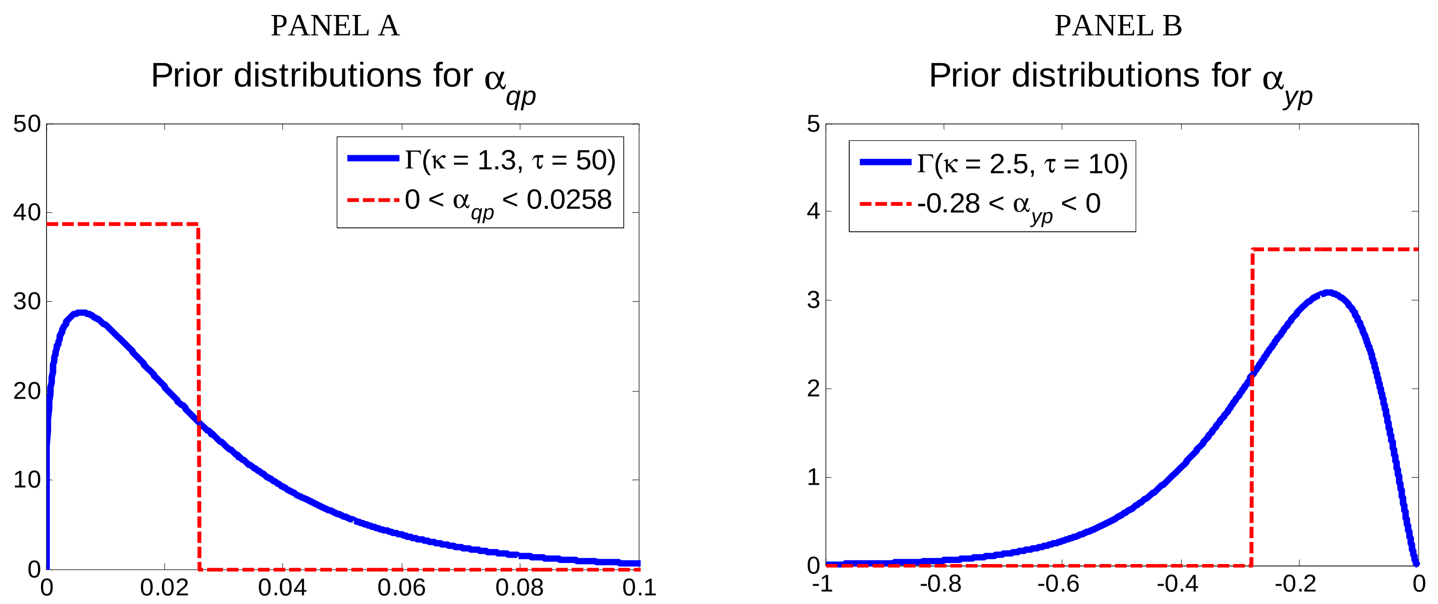

Figure 3. Prior beliefs about α_{qp} and α_{yp} represented by a bound (red dotted line) as in Kilian and Murphy (2012) and by a continuous density (blue solid line).

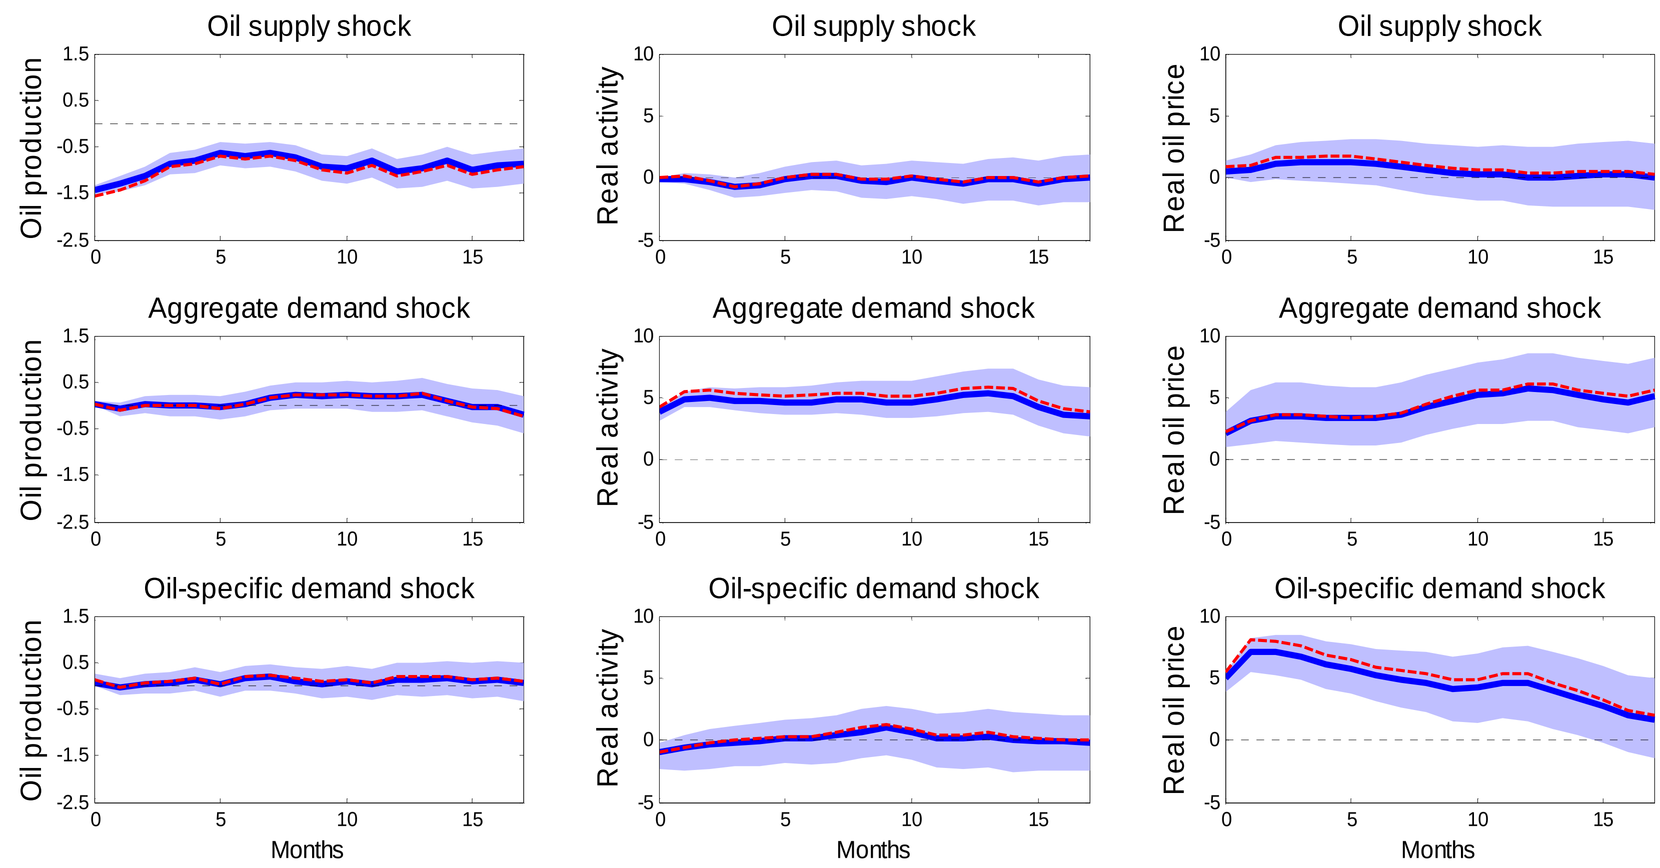

Figure 4. Impulse-response functions for 3-variable model partially identified using sign restrictions and bounds. Red dashed lines: estimates arrived at using Kilian and Murphy's (2012) original methodology; blue solid lines: Bayesian posterior median; shaded regions: 95% posterior credible sets.

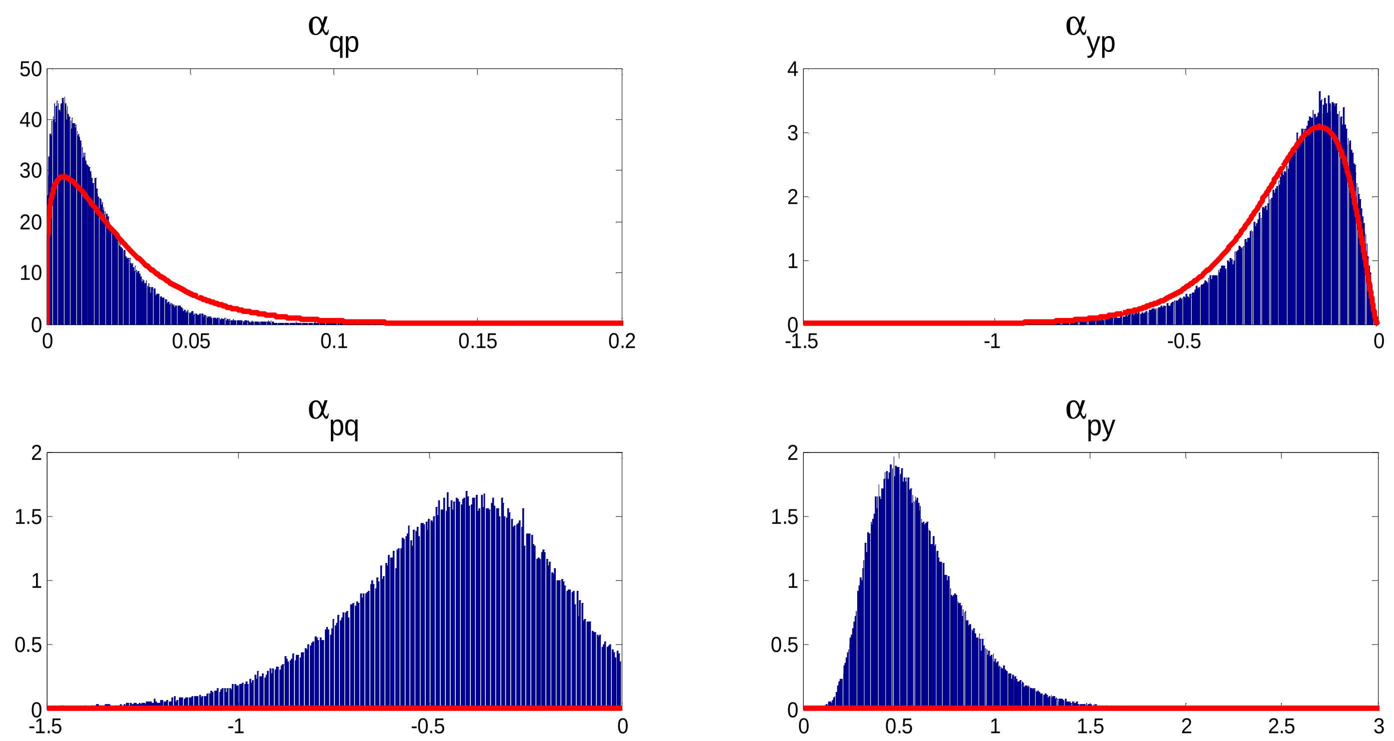

Figure 5. Prior (red lines) and posterior (blue histograms) distributions for the unknown elements in **A** in Bayesian implementation of the 3-variable partially identified model of Kilian and Murphy (2012).

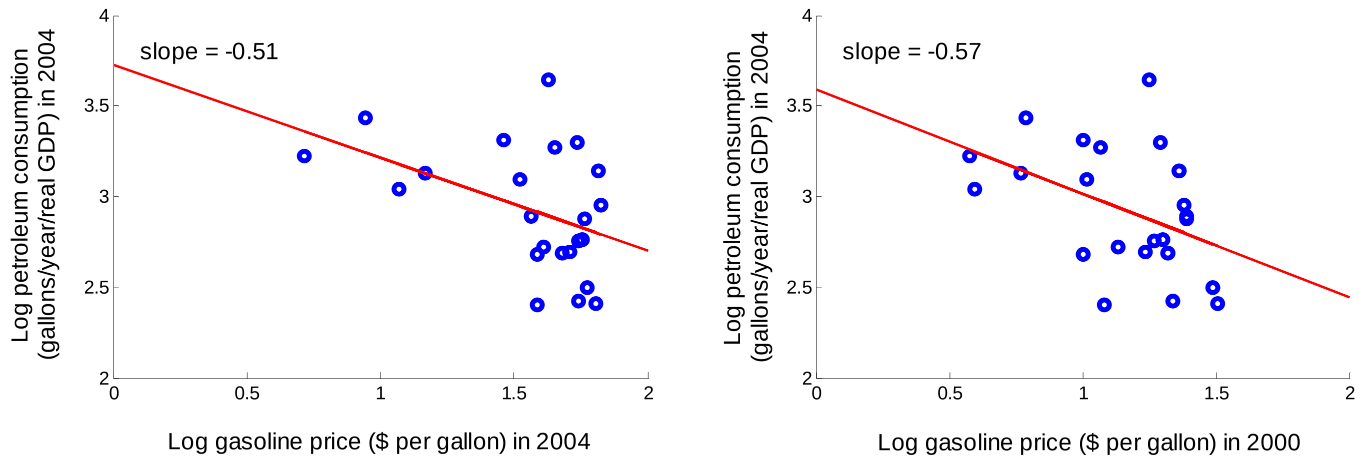

Figure 6. Evidence on demand elasticity from cross-section of 23 countries. Left panel: log of oil consumption per dollar of real GDP in 2004 versus price of gasoline in 2004. Right panel: log of oil consumption per dollar of real GDP in 2004 versus price of gasoline in 2000.

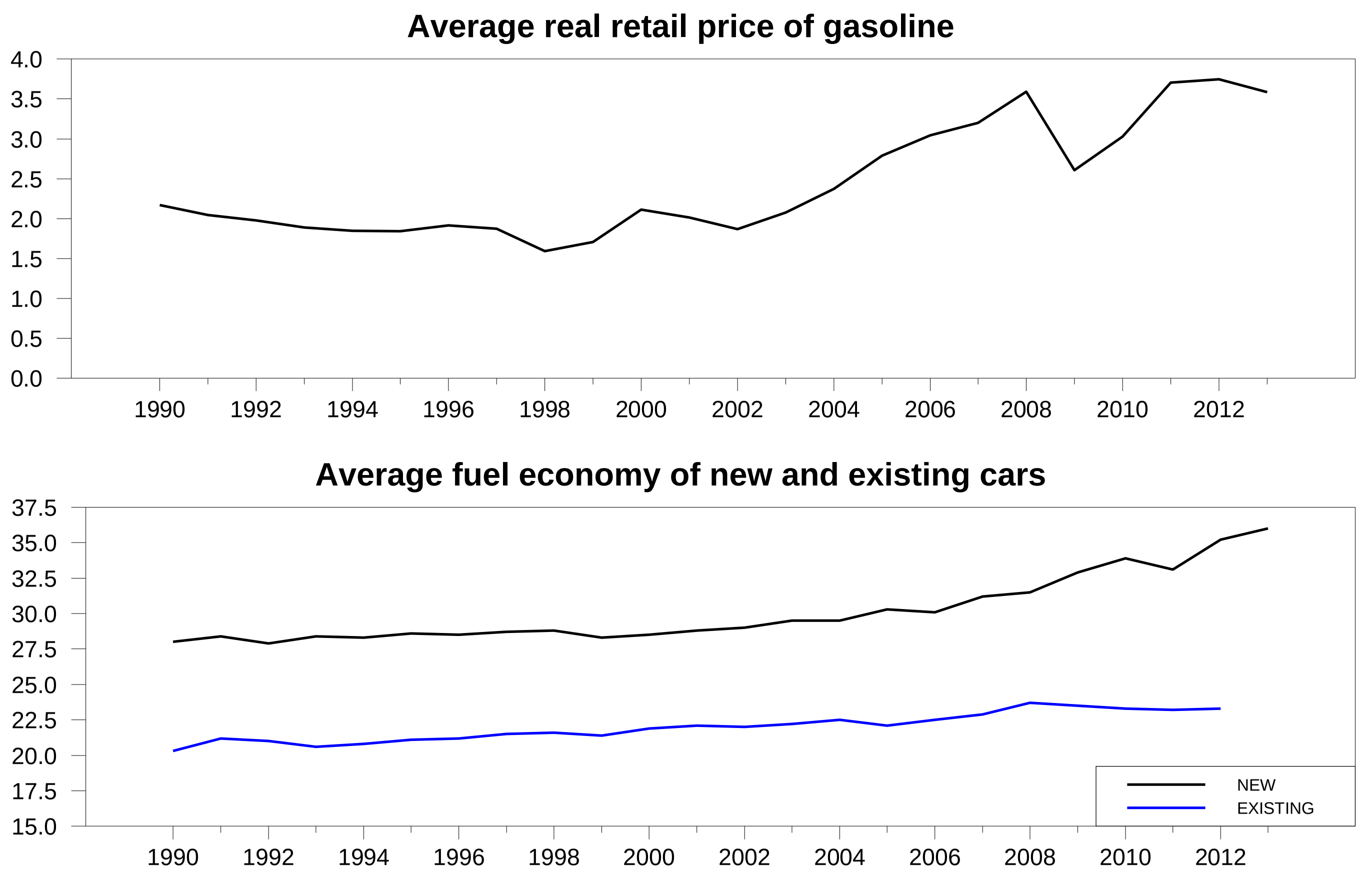

Figure 7. Average U.S. retail gasoline price and fuel economy, 1990-2013. Top panel: average retail price of all grades of gasoline (from Energy Information Administration (EIA), *Monthly Energy Review*, Table 9.4, http://www.eia.gov/totalenergy/data/monthly/#prices) in 2013 dollars per gallon (deflated using annual average of monthly seasonally unadjusted CPI, from http://research.stlouisfed.org/fred2/series/CPIAUCNS). Bottom panel: average fuel economy of new passenger vehicles and existing fleet of short-wheel base light-duty vehicles in miles per gallon (from http://www.rita.dot.gov/bts/sites/rita.dot.gov.bts/files/publications/national_transportation_statistics/html/table_04_23.html).

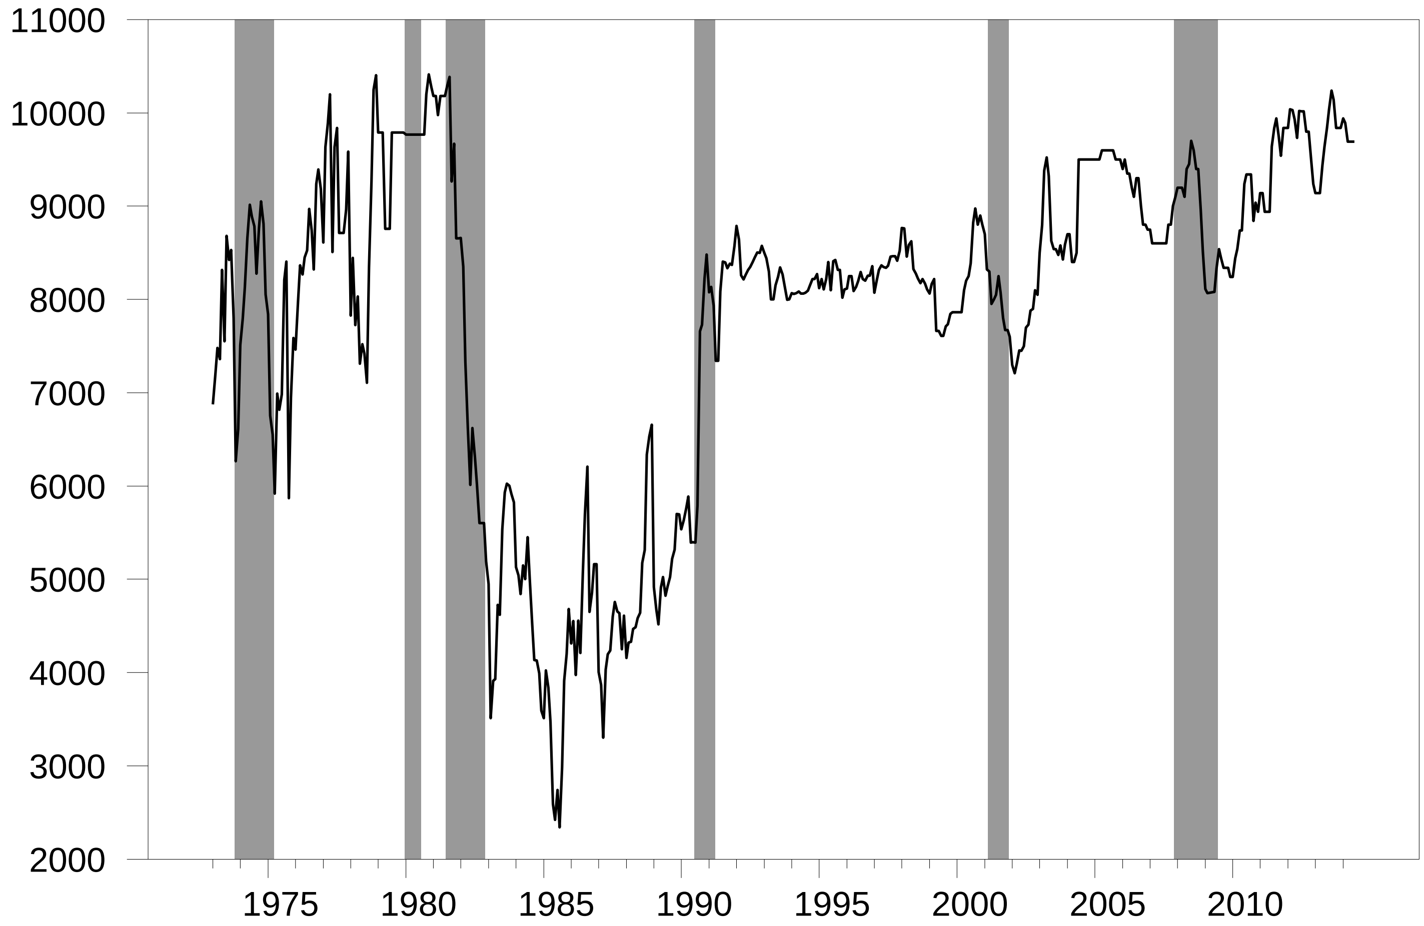

Figure 8. Monthly crude oil production from Saudi Arabia, January 1973 to June 2014, in thousands of barrels per day. Data source: Energy Information Administration (EIA), *Monthly Energy Review*, Table 11.1a (http://www.eia.gov/totalenergy/data/monthly/#international). Shaded regions correspond to U.S. economic recessions as dated by the National Bureau of Economic Research.

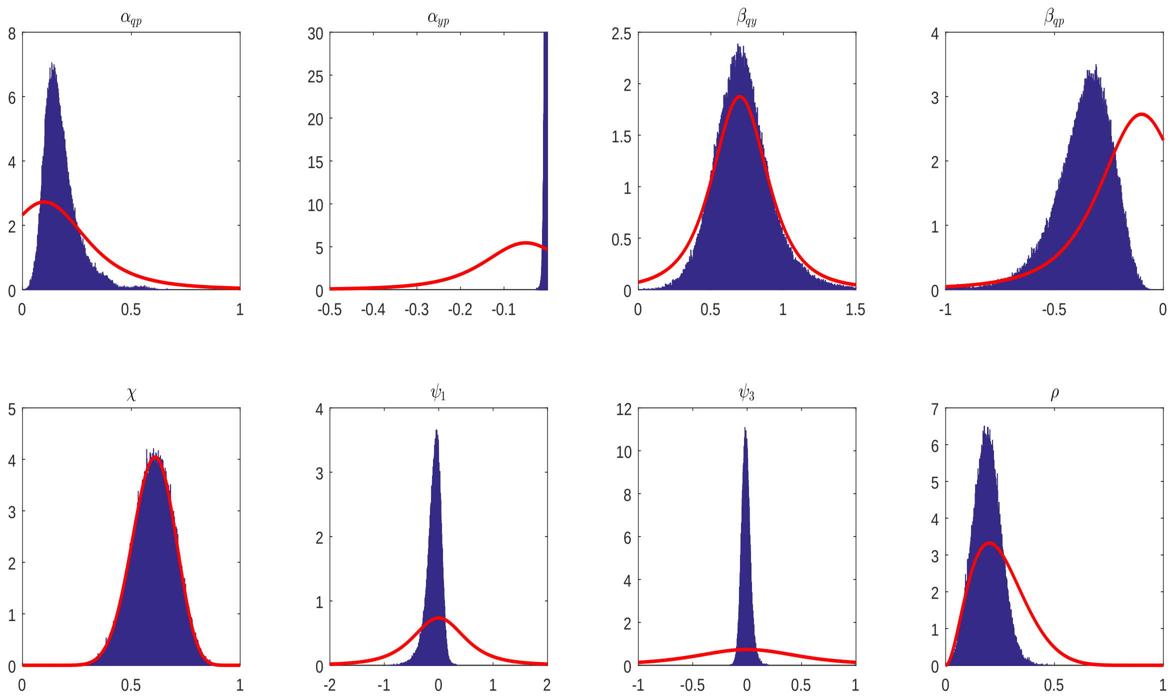

Figure 9. Baseline prior (solid red curves) and posterior (blue histograms) distributions for the unknown elements in **A** in 4-variable model with inventories, measurement error, and weighting data prior to 1975 by $\mu = 0.5$.

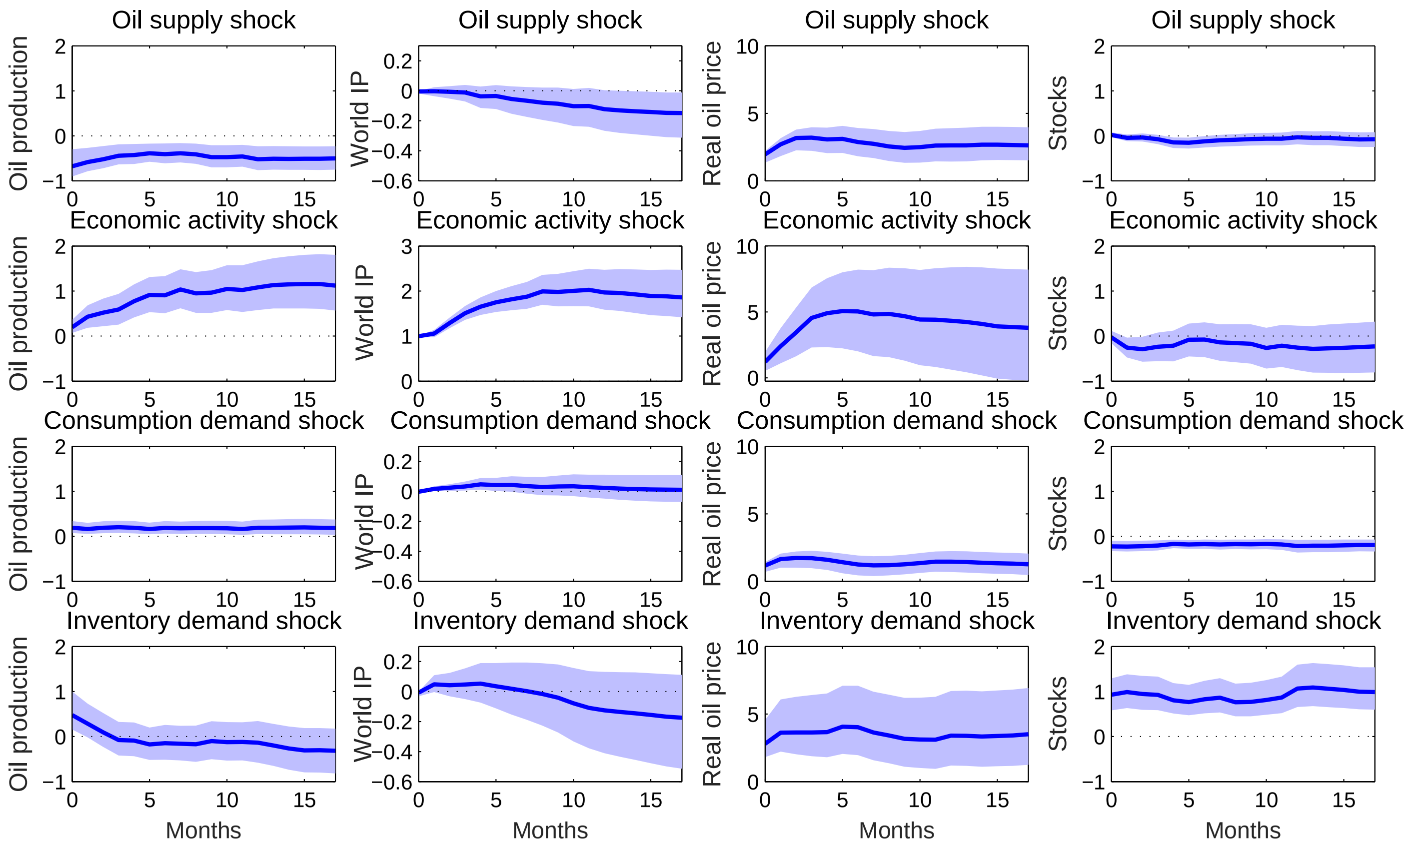

Figure 10. Impulse-response functions for 4-variable model with inventories, measurement error, and weighting data prior to 1975 by $\mu = 0.5$. Blue solid lines: Bayesian posterior median; shaded regions: 95% posterior credible sets.

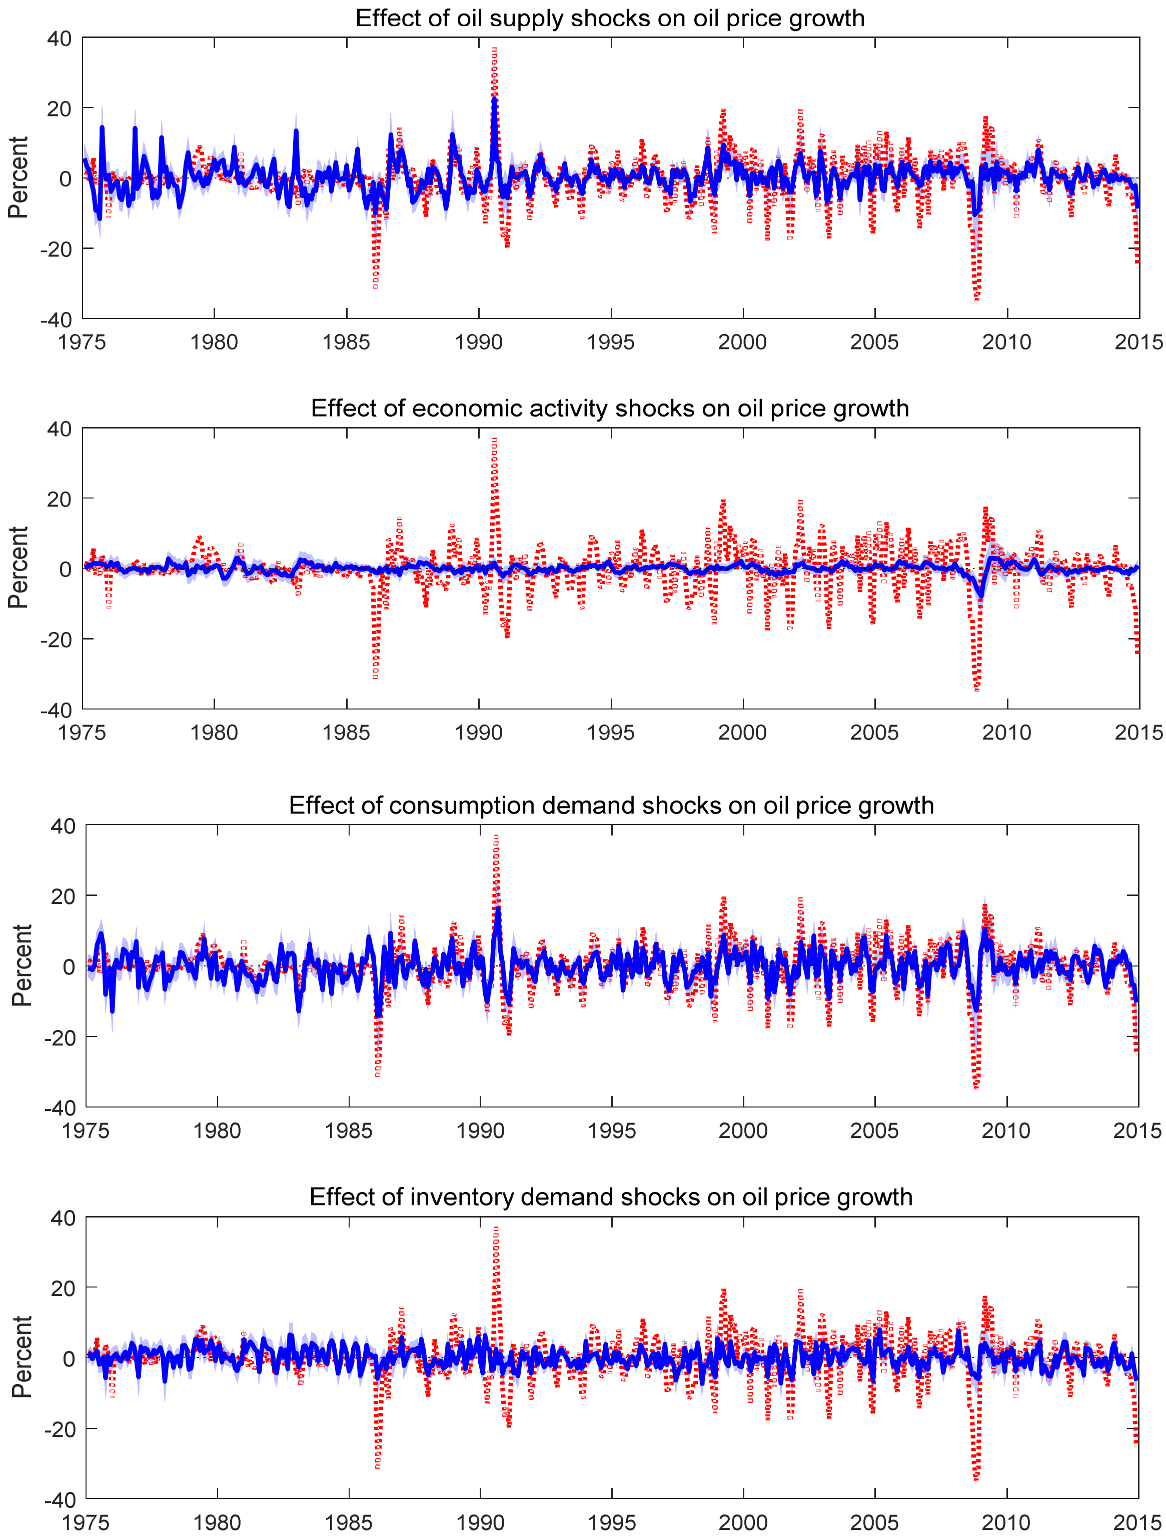

Figure 11. Actual changes in oil prices (red dashed lines) and historical contribution of separate structural shocks with 95% posterior credibility regions (blue and shaded) based on the 4-variable model with inventories, measurement error, and weighting data prior to 1975 by $\mu = 0.5$.