

NBER WORKING PAPER SERIES

STRUCTURAL INTERPRETATION OF VECTOR AUTOREGRESSIONS WITH INCOMPLETE IDENTIFICATION:
REVISITING THE ROLE OF OIL SUPPLY AND DEMAND SHOCKS

Christiane J.S. Baumeister
James D. Hamilton

Working Paper 24167
<http://www.nber.org/papers/w24167>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2017

We are grateful to Lutz Kilian, John Leahy, Dan Murphy, Eric Swanson and anonymous referees for helpful suggestions. Hamilton also thanks the JP Morgan Center for Commodities for financial support as a JPMCC Distinguished Visiting Fellow. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Christiane J.S. Baumeister and James D. Hamilton. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

Christiane J.S. Baumeister and James D. Hamilton

NBER Working Paper No. 24167

December 2017, Revised August 2018

JEL No. C32,E32,Q43

ABSTRACT

Traditional approaches to structural vector autoregressions can be viewed as special cases of Bayesian inference arising from very strong prior beliefs. These methods can be generalized with a less restrictive formulation that incorporates uncertainty about the identifying assumptions themselves. We use this approach to revisit the importance of shocks to oil supply and demand. Supply disruptions turn out to be a bigger factor in historical oil price movements and inventory accumulation a smaller factor than implied by earlier estimates. Supply shocks lead to a reduction in global economic activity after a significant lag, whereas shocks to oil demand do not.

Christiane J.S. Baumeister
Department of Economics
University of Notre Dame
3028 Jenkins Nanovic Hall
Notre Dame, IN 46556
and NBER
cbaumeis@nd.edu

James D. Hamilton
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER
jhamilton@ucsd.edu

A data appendix is available at <http://www.nber.org/data-appendix/w24167>

1 Introduction.

Structural interpretation of vector autoregressions has typically involved an all-or-nothing approach to the use of prior information, treating some features of the underlying structure as if known with certainty (often regarded as identifying assumptions) while claiming to be completely ignorant about other features. In this paper we argue that both aspects of the traditional approach can be improved on. Researchers need to acknowledge openly that a reasonable person would have doubts about the restrictions that are typically viewed as identifying assumptions. But we can make up for this in part by drawing on all available information about the structure, while acknowledging that this information, too, is imperfect. Under our approach, error bands incorporate not just uncertainty that is a result of having a finite sample of data but also reflect our doubts about the structure itself.

We illustrate these ideas by revisiting the role of supply and demand in generating historical fluctuations in the price of oil. As examples of excessive confidence about certain features of the underlying structural model, we revisit Kilian's (2009) assumption that we know with certainty that there is no short-run response of oil supply to the price and Kilian and Murphy's (2012) assumption that the short-run price elasticity of oil supply is known to be less than 0.0258. On the other hand, those studies made no use of information about the oil demand elasticity, and indeed their estimates imply demand elasticities that are far too large to be economically reasonable. We use this setting to illustrate how one can relax the strong restrictions about supply but supplement it with imperfect information about demand and other features of the economic structure to answer the kinds of questions researchers have studied with earlier methods.

Our paper makes a number of other methodological contributions. First, we show how to use prior information about both elasticities and the equilibrium impacts of structural shocks. Second, we show how to generalize structural vector autoregressions to allow for measurement error. Third, we show how one can downweight earlier data if the researcher has doubts about structural stability over time.

Among the new insights that emerge from our analysis is an estimate of the short-run oil supply elasticity of 0.15, consistent with the conclusion of Caldara, Cavallo and Iacoviello (2017) but considerably larger than the upper bound assumed in Kilian and Murphy (2012, 2014). We are also led to conclude that supply shocks were more important in accounting for historical oil price movements than was found in studies that assumed very precise prior information about the size of the supply elasticity. We attribute the run-up in oil prices in 2007-2008 to strong demand confronting stagnating supply. Our results suggest that weak demand and strong supply were both important in the oil price collapse in 2014-2016, while attributing most of the rebound in oil prices in 2016 to stronger demand. Our analysis further suggests that there is considerable error in measuring world inventories of oil. Once we allow

for this measurement error, we find little evidence for a contribution of speculation or changes in inventory demand to most historical oil price movements, in contrast for example to the conclusion of Juvenal and Petrella (2015).

The plan of the paper is as follows. Section 2 summarizes the Bayesian approach for a model that may be incompletely identified. Section 3 uses this framework to revisit earlier studies on the role of oil supply and demand shocks. Section 4 shows how we can incorporate a role for inventories while acknowledging the possibility of considerable error in estimates of global oil inventories. Section 5 summarizes the prior information we rely on and describes the conclusions that follow. Section 6 investigates how results change when we relax reliance on individual sources of prior information, while Section 7 concludes.

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} ($\mathbf{x}'_{t-1} = (\mathbf{y}'_{t-1}, \mathbf{y}'_{t-2}, \dots, \mathbf{y}'_{t-m}, 1)'$), and \mathbf{u}_t an $(n \times 1)$ vector of structural disturbances. 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 (2015) 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 information.

From a Bayesian perspective, a researcher's prior information 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. Information may pertain to individual elements of \mathbf{A} or to nonlinear combinations such as specified elements of \mathbf{A}^{-1} , the equilibrium effects of structural shocks. Our applications in this paper draw on both sources of prior information. Implementation of our procedure requires only that $p(\mathbf{A})$ be a proper density that integrates to unity.¹

¹Actually our algorithm can be implemented even if one does not 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.

While we allow the researcher to have arbitrary prior information about \mathbf{A} , to reduce computational demands we assume that prior information 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 information about \mathbf{D} conditional on \mathbf{A} can be represented using $\Gamma(\kappa_i, \tau_i)$ distributions for d_{ii}^{-1} ,

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

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

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 this prior information, we would choose κ_i and τ_i to be large numbers to get a prior distribution tightly centered around κ_i/τ_i . In the formulas below we allow τ_i to depend on \mathbf{A} but assume that κ_i does not. Appendix A offers some suggestions for how to choose the values for κ_i and τ_i .

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

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

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

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 i th structural equation. The matrix \mathbf{M}_i characterizes our confidence in this prior information. 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. The applications in this paper allow \mathbf{m}_i to depend on \mathbf{A} but assume that \mathbf{M}_i does not. Appendix A offers suggestions for specifying \mathbf{m}_i and \mathbf{M}_i .

2.2 Sampling from the posterior distribution.

The Bayesian begins with prior information about parameters $p(\mathbf{A}, \mathbf{D}, \mathbf{B})$ represented by the product of $p(\mathbf{A})$ with (2) and (3). The objective is to see how observation of the data $\mathbf{Y}_T = (\mathbf{y}'_1, \mathbf{y}'_2, \dots, \mathbf{y}'_T)'$ causes us to revise these beliefs. 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 turns out to be $\Gamma(\kappa_i^*, \tau_i^*(\mathbf{A}))$ where

$$\kappa_i^* = \kappa_i + T/2 \quad (5)$$

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

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

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

$$\tilde{\mathbf{Y}}_i(\mathbf{A}) = \left[\begin{array}{cccc} \mathbf{a}'_i\mathbf{y}_1 & \cdots & \mathbf{a}'_i\mathbf{y}_T & \mathbf{m}_i(\mathbf{A})'\mathbf{P}_i \end{array} \right]'_{[(T+k)\times 1]} \quad (8)$$

$$\tilde{\mathbf{X}}_i = \left[\begin{array}{cccc} \mathbf{x}_0 & \cdots & \mathbf{x}'_{T-1} & \mathbf{P}_i \end{array} \right]'_{[(T+k)\times k]} \quad (9)$$

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 turns out to be $N(\mathbf{m}_i^*(\mathbf{A}), d_{ii}\mathbf{M}_i^*)$ with

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

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

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

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

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,

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

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

$$p(\mathbf{A}, \mathbf{D}, \mathbf{B}|\mathbf{Y}_T) = p(\mathbf{A}|\mathbf{Y}_T) p(\mathbf{D}|\mathbf{A}, \mathbf{Y}_T) p(\mathbf{B}|\mathbf{A}, \mathbf{D}, \mathbf{Y}_T) \quad (14)$$

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 (2015). 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

$$\mathbf{y}_t = \mathbf{\Phi} \mathbf{x}_{t-1} + \boldsymbol{\epsilon}_t \quad (15)$$

$$= \mathbf{\Phi}_1 \mathbf{y}_{t-1} + \mathbf{\Phi}_2 \mathbf{y}_{t-2} + \cdots + \mathbf{\Phi}_m \mathbf{y}_{t-m} + \mathbf{c} + \boldsymbol{\epsilon}_t$$

$$\mathbf{\Phi} = \mathbf{A}^{-1} \mathbf{B} \quad (16)$$

$$\boldsymbol{\epsilon}_t = \mathbf{A}^{-1} \mathbf{u}_t. \quad (17)$$

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

$$\boldsymbol{\Psi}_s = \frac{\partial \mathbf{y}_{t+s}}{\partial \boldsymbol{\epsilon}_t'}, \quad (18)$$

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

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

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

$$\mathbf{H}_s = \frac{\partial \mathbf{y}_{t+s}}{\partial \mathbf{u}_t'} = \boldsymbol{\Psi}_s \mathbf{A}^{-1}; \quad (19)$$

see for example Hamilton (1994, pages 260 and 331). We report the pointwise median value of draws of these magnitudes; see Baumeister and Hamilton (2017) for a discussion of the optimality properties of these estimates.

3 Bayesian interpretation of traditional approaches to structural inference.

In this section we show how previous approaches can be given a Bayesian interpretation, using a 3-variable description of the global oil market for illustration. 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: $\mathbf{y}_t = (q_t, y_t, p_t)'$. For this section we 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 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:

$$q_t = \alpha_{qy}y_t + \alpha_{qp}p_t + \mathbf{b}'_1\mathbf{x}_{t-1} + u_{1t} \quad (20)$$

$$y_t = \alpha_{yq}q_t + \alpha_{yp}p_t + \mathbf{b}'_2\mathbf{x}_{t-1} + u_{2t} \quad (21)$$

$$p_t = \alpha_{pq}q_t + \alpha_{py}y_t + \mathbf{b}'_3\mathbf{x}_{t-1} + u_{3t}. \quad (22)$$

Equation (20) 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 $\mathbf{x}_{t-1} = (\mathbf{y}'_{t-1}, \mathbf{y}'_{t-2}, \dots, \mathbf{y}'_{t-24}, 1)'$. Equation (21) 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 (22) governs oil demand, written here in inverse form so that α_{pq} is the reciprocal of the short-run price elasticity of demand. 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 Cholesky identification.

As our first example we consider the analysis by Kilian (2009), who used a familiar recursive interpretation of the structural system with variables ordered as (q_t, y_t, p_t) . 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

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & 0 \\ -\alpha_{yq} & 1 & 0 \\ -\alpha_{pq} & -\alpha_{py} & 1 \end{bmatrix} \quad (23)$$

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

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 as dashed red curves in Figure 1. We also show the posterior median (solid blue) calculated using the Bayesian algorithm described in Appendix B using the above prior distributions. The two inferences are identical.

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 (see the upper right panel of Figure 2). The prior distribution is essentially a flat line when viewed on this scale, while the posterior has most of its mass between -0.6 and $+0.2$, implying a short-run price elasticity of demand that is concentrated within $(-\infty, -1.67) \cup (+5, \infty)$. The latter distribution is plotted in the last panel of Figure 2. 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 an all-or-nothing approach is the only way that one could study these questions 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

$$\begin{bmatrix} + & + & + \\ + & + & - \\ - & + & + \end{bmatrix}. \quad (24)$$

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, requiring that feedback effects arising from a possible direct response of oil production to economic activity (α_{qy}) or economic activity to oil production (α_{yq}) must be small. One simple way to guarantee the sign restrictions is to set these two parameters to zero:

$$\mathbf{A} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} \\ 0 & 1 & -\alpha_{yp} \\ -\alpha_{pq} & -\alpha_{py} & 1 \end{bmatrix}. \quad (25)$$

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,

$$\mathbf{A}^{-1} = \frac{1}{1 - \alpha_{qp}\alpha_{pq} - \alpha_{py}\alpha_{yp}} \begin{bmatrix} 1 - \alpha_{py}\alpha_{yp} & \alpha_{qp}\alpha_{py} & \alpha_{qp} \\ \alpha_{yp}\alpha_{pq} & 1 - \alpha_{qp}\alpha_{pq} & \alpha_{yp} \\ \alpha_{pq} & \alpha_{py} & 1 \end{bmatrix}. \quad (26)$$

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 (26) will always satisfy (24). The under-identified system (25) 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 (23), is nevertheless known to be small, which they represented with the bounds $\alpha_{qp} \in [0, 0.0258]$. 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 the uniform distribution

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

This density is plotted as the red curve in the upper left panel of Figure 3.

Kilian and Murphy also explored the benefits of using prior information about the (2,3) element of (26), $\alpha_{yp}/\det(\mathbf{A})$, which corresponds to the equilibrium effect on economic activity of a shock to demand. They imposed that the effect of a one-standard-deviation shock was restricted to fall in $[-1.5, 0]$. This prior is plotted in the bottom left panel of Figure 3.²

By contrast, Kilian and Murphy did not use any prior information at all about the other parameters other than the sign restrictions mentioned above. We again represent this with the very uninformative Student t priors used in Section 3.1 now truncated by sign restrictions. We used the algorithm described in Appendix B to form posterior inference resulting from the prior

$$p(\mathbf{A}) \propto \begin{cases} p(\alpha_{yp})p(\alpha_{pq})p(\alpha_{py}) & \text{if } \alpha_{qp} \in [0, 0.0258] \text{ and } \sqrt{d_{33}}\alpha_{yp}/\det(\mathbf{A}) \in [-1.5, 0] \\ 0 & \text{otherwise} \end{cases}$$

for $p(\alpha_{yp})$ and $p(\alpha_{pq})$ Student t (0,100,3) densities truncated to be negative and $p(\alpha_{py})$ a Student t (0,100,3) density truncated to be positive. The resulting posterior medians for the impulse-response functions are shown in blue in Figure 4, and coincide almost exactly with the

²Specifically, we imposed $-1.5 < h_{23} < 0$ for $h_{23} = \sqrt{d_{33}}\alpha_{yp}/\det(\mathbf{A})$ and $\sqrt{d_{33}} = 5.44$, the standard deviation of the error of the reduced-form forecast of p_t .

inference reported in Kilian and Murphy’s article calculated using their original methodology, which is reproduced as the dashed red lines in Figure 4.³

Kilian and Murphy’s approach, like the more traditional approach to identification discussed in Section 3.1, can thus again be given a Bayesian interpretation. But as in the previous example, once we do so, we see prior information being used in some strange ways. It is very odd to regard a supply elasticity of 0.0257 as perfectly plausible while maintaining that an elasticity of 0.0259 is completely impossible. A more natural and honest representation of prior information would allow at least some possibility of larger values and would not involve a sharp drop-off in the probability at any fixed value. And it is very strange to claim to have such precise information about the supply elasticity but know nothing at all about the demand elasticity. Posterior distributions for the elements of \mathbf{A} are plotted as blue histograms in Figure 3, with the bottom right panel displaying the implied posterior distribution for the short-run price elasticity of oil demand. The underlying model would lead the researcher to conclude that monthly demand is extremely sensitive to the current price, with a 60% posterior probability that a 10% increase in price leads to more than a 20% drop in quantity demanded within a single month.

In the following sections we review the literature on what we actually know and what we don’t know about supply and demand elasticities.

3.3 Do we really know for certain that the oil supply elasticity is less than 0.0258?

The online appendix to Kilian and Murphy (2012) justifies their 0.0258 bound on the supply elasticity from the following reasoning. When Iraq invaded Kuwait in August of 1990 oil production from both countries fell dramatically and the price went up 45.3%. But production outside of Iraq and Kuwait increased 1.17% in August, suggesting a short-run supply elasticity of $1.17/45.3 = 0.0258$. Kilian and Murphy regarded this as an upper bound on what we might expect in normal times due to excess capacity in 1990 and because “there was rare unanimity among oil producers in 1990 that it was essential to offset market fears about a wider war in the Middle East.”

Just prior to the invasion, on July 17 Iraq’s President Saddam Hussein had threatened to use military force on Arab nations that did not curb oil production. Caldara, Cavallo, and Iacoviello (2017) noted that the *New York Times* reported this threat with the headline, “Iraq Threatens Emirates and Kuwait on Oil Glut.” Within a week, the United Arab Emirates announced they would implement a significant cut in production, and indeed U.A.E. production in August was 19.5% lower than in July. Oil production outside of Iraq, Kuwait,

³The dashed red lines were produced using the exact methodology of their paper, which is not the posterior median from their model.

and U.A.E. actually increased 1.95% in August, which would imply an elasticity of $1.95/45.3 = 0.043$, almost twice the Kilian-Murphy estimate, once we acknowledge the effect of the implicit military threat on U.A.E. production.

Caldara, Cavallo and Iacoviello further noted that August 1990 was but one of dozens of historical episodes like this that could have been used for such calculations. Other examples include strikes by Norwegian oil workers in 1986, attacks on Libyan oil fields in 2011, and hurricanes disrupting Mexican production in 1995 and U.S. production in 2005 and 2008. Estimating an average monthly supply elasticity for countries excluded from each episode using instrumental variables, the authors came up with a short-run supply elasticity of 0.077 (three times the Kilian-Murphy upper bound) with a standard error of 0.037.

Figure 5 plots monthly production from Saudi Arabia in which one clearly sees high-frequency adjustments to changing market conditions. 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.

Bjørnland, Nordvik and Rohrer (2017) analyzed monthly crude oil production from 15,000 individual wells in North Dakota over 1986 to 2015. They found producers varied both the timing of completion of new wells as well as production flows from existing wells in response to monthly changes in prices, consistent with a short-run supply elasticity as high as 0.2 for some shale producers.

Even Kilian and Murphy (2012) implicitly acknowledged that they themselves were not certain that the supply elasticity is below 0.0258, insofar as they also reported some results using an upper bound two or three times this magnitude. If we were to give that procedure of reporting multiple estimates a Bayesian interpretation, it amounts to first producing an answer as if we were certain the elasticity is below 0.0258, then producing an alternative answer as if we were certain the elasticity is below 0.0516. Of course the correct Bayesian way to use this kind of prior information is to acknowledge that we are not certain that the elasticity is below 0.0258, nor are we certain that it is below 0.0516. But we do attach lower and lower prior plausibility to higher and higher elasticities. This is the approach we will follow in the analysis conducted below.

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

Hundreds of studies have looked at the price elasticity of oil demand using all kinds of different data sources and methods. Studies using cross-section data include Hausman and Newey's (1995) estimate from a cross-section of U.S. households of a long-run price elasticity of gasoline demand of -0.81 and Yatchew and No's (2001) estimate of -0.9 from a cross-section of Canadian households. Figure 6 displays some cross-country evidence, comparing petroleum use per dollar of GDP with the price of gasoline for 23 OECD countries.⁴ 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.⁵ 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 demand 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.

Dahl and Sterner (1991) surveyed 296 different estimates of the long-run price elasticity of gasoline demand based on cross-section, time-series, and panel data, and found an average 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.

And there is no doubt that the short-run elasticity is significantly less than the long run. For example, it takes more than a decade for the stock of automobiles to turn over. Dahl and Sterner's (1991) survey found an average short-run elasticity of -0.26 . Hughes, Knittel and Sperling (2008) used exogenous petroleum supply disruptions as an instrument to conclude that the short-run gasoline price elasticity was below 0.08 in absolute value for U.S. data over 2001-2006. Gelman et al. (2017) estimated a short-run elasticity of -0.22 with a standard error of 0.05 from observations on individual financial transactions of a half million consumers, while Coglianese et al. (2017) used state tax changes as an instrument to arrive at an estimate of -0.37 with a standard error of 0.24.

We conclude that short-run oil demand elasticities above two in absolute value, such as

⁴Data 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.

⁵See for example Darmstadter, Dunkerly, and Alterman (1977), Drollas (1984), and Davis (2014).

were implied by the bottom right panels in Figures 2 and 3, are highly implausible.

4 Inventories and measurement error.

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

$$q_t = \alpha_{qp}p_t + \mathbf{b}'_1\mathbf{x}_{t-1} + u_{1t}^* \quad (27)$$

$$y_t = \alpha_{yp}p_t + \mathbf{b}'_2\mathbf{x}_{t-1} + u_{2t}^* \quad (28)$$

$$q_t = \beta_{qy}y_t + \beta_{qp}p_t + \Delta i_t^* + \mathbf{b}'_3\mathbf{x}_{t-1} + u_{3t}^* \quad (29)$$

$$\Delta i_t^* = \psi_1^*q_t + \psi_2^*y_t + \psi_3^*p_t + \mathbf{b}'_4\mathbf{x}_{t-1} + u_{4t}^*. \quad (30)$$

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 (29) acknowledging that oil produced but not consumed in the current period goes into inventories. The shock u_{4t}^* in (30) 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

$$\Delta i_t = \chi \Delta i_t^* + e_t \quad (31)$$

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 which we assume to be serially uncorrelated and uncorrelated with \mathbf{u}_t^* . 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.⁶ However, it is straightforward to incorporate measurement error in our Bayesian framework, as we now demonstrate.

We can use (31) to rewrite (29) and (30) in terms of observables:

$$q_t = \beta_{qy}y_t + \beta_{qp}p_t + \chi^{-1}\Delta i_t + \mathbf{b}'_3\mathbf{x}_{t-1} + u_{3t}^* - \chi^{-1}e_t \quad (32)$$

$$\Delta i_t = \psi_1q_t + \psi_2y_t + \psi_3p_t + \mathbf{b}'_4\mathbf{x}_{t-1} + \chi u_{4t}^* + e_t \quad (33)$$

where $\psi_j = \chi\psi_j^*$ for $j = 1, 2, 3$. Equations (27), (28), (32), and (33) will be recognized as a system of the form

$$\tilde{\mathbf{A}}\mathbf{y}_t = \tilde{\mathbf{B}}\mathbf{x}_{t-1} + \tilde{\mathbf{u}}_t \quad (34)$$

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

$$\tilde{\mathbf{A}} = \begin{bmatrix} 1 & 0 & -\alpha_{qp} & 0 \\ 0 & 1 & -\alpha_{yp} & 0 \\ 1 & -\beta_{qy} & -\beta_{qp} & -\chi^{-1} \\ -\psi_1 & -\psi_2 & -\psi_3 & 1 \end{bmatrix} \quad (35)$$

$$\tilde{\mathbf{u}}_t = \begin{bmatrix} u_{1t}^* \\ u_{2t}^* \\ u_{3t}^* - \chi^{-1}e_t \\ \chi u_{4t}^* + e_t \end{bmatrix}. \quad (36)$$

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 (27)-(30) to be based on lags of Δi_{t-j} rather than Δi_{t-j}^* .

The residuals \tilde{u}_{3t} and \tilde{u}_{4t} in (34) are contemporaneously correlated. We show in Appendix D that by premultiplying (34) by the matrix $\mathbf{\Gamma}$ in expression (48), we can transform it into a representation of the form of (1), in which the shocks are uncorrelated. The matrix $\mathbf{\Gamma}$ is a function of ρ , the negative of a coefficient from a regression of \tilde{u}_{4t} on \tilde{u}_{3t} :

$$\rho = \frac{\chi^{-1}\sigma_e^2}{d_{33}^* + \chi^{-2}\sigma_e^2}. \quad (37)$$

⁶Notable exceptions are Cogley and Sargent (2015) who allowed for measurement error using a state-space model and Amir-Ahmadi, Matthes, and Wang (2017) who identified measurement error from the difference between preliminary and revised data.

For the transformed system we can generate draws for the parameters $\mathbf{A}, \mathbf{D}, \mathbf{B}$ in (1) using the algorithm described in Appendix B.

From these we can then go back and calculate the values of $\tilde{\mathbf{A}}, \tilde{\mathbf{D}}, \tilde{\mathbf{B}}$ for the parameters in (34) to make structural inference. To calculate structural impulse-response functions, note that premultiplying (34) by $\tilde{\mathbf{A}}^{-1}$ puts the system in the reduced form (15) with $\boldsymbol{\epsilon}_t = \tilde{\mathbf{A}}^{-1}\tilde{\mathbf{u}}_t$. Thus

$$\frac{\partial \mathbf{y}_{t+s}}{\partial \tilde{\mathbf{u}}_t'} = \frac{\partial \mathbf{y}_{t+s}}{\partial \boldsymbol{\epsilon}_t'} \frac{\partial \boldsymbol{\epsilon}_t}{\partial \tilde{\mathbf{u}}_t'} = \boldsymbol{\Psi}_s \tilde{\mathbf{A}}^{-1} \quad (38)$$

for $\boldsymbol{\Psi}_s$ the nonorthogonalized impulse-response function in (18). From (36) we further know that

$$\frac{\partial \tilde{\mathbf{u}}_t}{\partial \mathbf{u}_t^{*'}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & -\chi^{-1} \\ 0 & 0 & 0 & \chi & 1 \end{bmatrix} = \boldsymbol{\Xi}$$

$$\frac{\partial \mathbf{y}_{t+s}}{\partial \mathbf{u}_t^{*'}} = \boldsymbol{\Psi}_s \tilde{\mathbf{A}}^{-1} \boldsymbol{\Xi} = \mathbf{H}_s^* \quad (39)$$

for $\mathbf{u}_t^* = (u_{1t}^*, u_{2t}^*, u_{3t}^*, u_{4t}^*, e_t)'$.

For purposes of calculating a historical decomposition, conditional on a draw of the parameters, $\tilde{\mathbf{u}}_t$ can be uncovered from (34) as $\tilde{\mathbf{u}}_t = \tilde{\mathbf{A}}\mathbf{y}_t - \tilde{\mathbf{B}}\mathbf{x}_{t-1}$. However, observation of the four elements of $\tilde{\mathbf{u}}_t$ is not enough to know the value of the five shocks $u_{1t}^*, u_{2t}^*, u_{3t}^*, u_{4t}^*, e_t$. Nevertheless, we can form an optimal estimate of those 5 magnitudes for each historical date t . This gives us an estimate of the contribution of each of the five shocks to the observed historical values for \mathbf{Y}_T . Again see Appendix D for details.

5 A true Bayesian analysis of the shocks to oil supply and demand.

In addition to making use of the prior information about price elasticities reviewed in Sections 3.3 and 3.4, 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, 2014) 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 (WTI) 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 discusses the prior information used in our structural analysis.

5.1.1 Priors for A.

The discussion in Sections 3.3 and 3.4 leads us to conclude that the absolute values of the short-run demand elasticity β_{qp} and the short-run supply elasticity α_{qp} are unlikely to be much bigger than 0.5. We represent this with a prior for β_{qp} that is a Student $t(c_{qp}^\beta, \sigma_{qp}^\beta, \nu_{qp}^\beta)$ with mode at $c_{qp}^\beta = -0.1$, scale parameter $\sigma_{qp}^\beta = 0.2$, $\nu_{qp}^\beta = 3$ degrees of freedom, and truncated to be negative. This allows a 10% probability that $\beta_{qp} < -0.5$. Our prior for α_{qp} is Student $t(c_{qp}^\alpha, \sigma_{qp}^\alpha, \nu_{qp}^\alpha)$ with mode at $c_{qp}^\alpha = 0.1$, scale parameter $\sigma_{qp}^\alpha = 0.2$, $\nu_{qp}^\alpha = 3$ degrees of freedom, and truncated to be positive. These along with our other priors are summarized in Table 1 and displayed as the red curves in Figure 7. Appendix F provides additional demonstration that these are relatively weak priors that are perfectly consistent with values for the supply elasticity as low as those maintained by Kilian and Murphy (2012, 2014), but also allow the possibility of a supply elasticity greater than 0.0258.

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_{yy} 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 (37) captures the importance of inventory measurement error and is between 0 and χ by construction.⁷ We accordingly use a prior for ρ conditional on χ that is χ times a Beta-distributed variable with parameters $\alpha_\rho = 3$ and $\beta_\rho = 9$, which has a mean of 0.25χ and standard deviation of 0.12χ .

We can also make use of prior information about the likely equilibrium impacts of various shocks, which amount to prior beliefs about how the various elements of \mathbf{A} may be related. From (38) and (35) the equilibrium impacts of structural shocks are given by the matrix

$$\frac{\partial \mathbf{y}_t}{\partial \tilde{\mathbf{u}}_t'} = \tilde{\mathbf{A}}^{-1} = \frac{1}{\det(\tilde{\mathbf{A}})} \mathbf{C}$$

$$\mathbf{C} = \begin{bmatrix} -\psi_3\chi^{-1} - \alpha_{yp}\beta_{qy} - \beta_{qp} & \alpha_{qp}\beta_{qy} & \alpha_{qp} & \alpha_{qp}\chi^{-1} \\ \alpha_{yp}(\psi_1\chi^{-1} - 1) & \chi^{-1}(-\alpha_{qp}\psi_1 - \psi_3) + (\alpha_{qp} - \beta_{qp}) & \alpha_{yp} & \alpha_{yp}\chi^{-1} \\ \psi_1\chi^{-1} - 1 & \beta_{qy} & 1 & \chi^{-1} \\ -\psi_1(\alpha_{yp}\beta_{qy} + \beta_{qp}) - \psi_3 & \beta_{qy}(\alpha_{qp}\psi_1 + \psi_3) & \alpha_{qp}\psi_1 + \psi_3 & \alpha_{qp} - \alpha_{yp}\beta_{qy} - \beta_{qp} \end{bmatrix} \quad (40)$$

$$\det(\tilde{\mathbf{A}}) = \chi^{-1}(-\alpha_{qp}\psi_1 - \psi_3) + (\alpha_{qp} - \alpha_{yp}\beta_{qy} - \beta_{qp}).$$

Unless the determinant is restricted to be positive, all shocks could have either positive or negative effects on any variable. We could put as much or as little weight as we like on the prior belief that $h_1 = \det(\tilde{\mathbf{A}}) > 0$ using the asymmetric t distribution introduced by

⁷To see this, divide (37) by χ to verify that $\rho/\chi = \chi^{-2}\sigma_e^2/(d_{33}^* + \chi^{-2}\sigma_e^2) < 1$.

Baumeister and Hamilton (2017):

$$p(h_1) = k_1 \sigma_1^{-1} \tilde{\phi}_{\nu_1}((h_1 - \mu_1)/\sigma_1) \Phi(\lambda_1 h_1/\sigma_1). \quad (41)$$

Here $\tilde{\phi}_{\nu_1}(x)$ denotes the probability density function of a standard Student t variable with ν_1 degrees of freedom evaluated at the point x , $\Phi(x)$ is the cumulative distribution function for a standard $N(0, 1)$ variable, and k_1 is a constant to make the density integrate to unity⁸. The parameter λ_1 governs how strongly the distribution of h_1 is skewed to be positive. When $\lambda_1 = 0$ the density (41) is a symmetric Student $t(\mu_1, \sigma_1, \nu_1)$ distribution, while when $\lambda_1 \rightarrow \infty$ it becomes a Student t distribution truncated to be positive. To determine the location parameter μ_1 , we generated 50,000 draws for $\boldsymbol{\theta}_{\mathbf{A}} = (\alpha_{qp}, \alpha_{yp}, \beta_{qy}, \beta_{qp}, \chi, \psi_1, \psi_3, \rho)'$ from the densities described in Table 1 and used the average value across these draws to obtain $\mu_1 = 0.6$. We set $\sigma_1 = 1.6$, the standard deviation of $\det(\tilde{\mathbf{A}})$ across these draws. Setting $\lambda_1 = 2$ and $\nu_1 = 3$ associates a 91.2% prior probability to $\det(\tilde{\mathbf{A}}) > 0$. This prior distribution for h_1 is plotted in red in the next-to-last panel in Figure 7.

Even if $\det(\tilde{\mathbf{A}}) > 0$ and the other restrictions we have used are all imposed, the signs of some elements of the impact matrix $\tilde{\mathbf{A}}^{-1}$ are still ambiguous:

$$\tilde{\mathbf{A}}^{-1} = \begin{bmatrix} ? & + & + & + \\ ? & ? & - & - \\ ? & + & + & + \\ ? & ? & ? & + \end{bmatrix}.$$

These ambiguities arise from equilibrium feedback effects. For example, from (40) the (2,2) element of $\tilde{\mathbf{A}}^{-1}$ can be written as

$$h_2 = \frac{\det(\tilde{\mathbf{A}}) + \alpha_{yp} \beta_{qy}}{\det(\tilde{\mathbf{A}})}.$$

If oil demand increases sufficiently much in response to higher economic activity (β_{qy} large) and higher oil prices depress economic activity sufficiently much (α_{yp} a big negative number), it is possible in principle for h_2 to be a negative number. We can represent a belief that these feedback effects are modest with a prior for h_2 that is a symmetric Student t distribution with $\mu_2 = 0.8$, $\sigma_2 = 0.2$, and $\nu_2 = 3$, which imply a 98.6% prior probability that $h_2 > 0$.

We thus use the prior

$$p(\boldsymbol{\theta}_{\mathbf{A}}) \propto p(\alpha_{qp})p(\alpha_{yp})p(\beta_{qy})p(\beta_{qp})p(\chi)p(\psi_1)p(\psi_3)p(\rho|\chi)p(h_1(\boldsymbol{\theta}_{\mathbf{A}}))p(h_2(\boldsymbol{\theta}_{\mathbf{A}})) \quad (42)$$

which gives the numerical value for $p(\mathbf{A})$ used in the numerator of (12) that is associated with

⁸We don't need to know the value of k_1 for purposes of the algorithm described in Appendix B.

any proposed value for $\theta_{\mathbf{A}}$. Odd-numbered columns of Table 2 report the prior probabilities implied by (42) that the equilibrium impact of any given shock on any given variable is positive.

Our priors also imply probable signs for the effects s periods after a given shock based on the dynamics incorporated in the prior $p(\mathbf{B}|\mathbf{A}, \mathbf{D})$ described below. More persistence in \mathbf{B} implies more persistence in the effects of shocks.

5.1.2 Priors for \mathbf{D} given \mathbf{A} .

Our priors for the reciprocals of the structural variances are independent Gamma distributions, $d_{ii}^{-1}|\mathbf{A} \sim \Gamma(\kappa_i, \tau_i(\mathbf{A}))$, 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 $t = 1, \dots, T_1$ 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 $\kappa_i = 2$ (which give the priors a weight of about 4 observations in the first subsample) and $\tau_i(\mathbf{A}) = \kappa_i \mathbf{a}'_i \hat{\mathbf{S}} \mathbf{a}_i$ where \mathbf{a}'_i denotes the i th row of \mathbf{A} .

5.1.3 Priors for \mathbf{B} given \mathbf{A} and \mathbf{D} .

Our priors for the lagged coefficients in the i th 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 changes in oil production, economic activity, oil prices, and inventories are all hard to forecast, meaning our prior expected value for most coefficients is $\mathbf{m}_i = \mathbf{0}$ for $i = 1, \dots, 4$. 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 $\mathbf{m}_1, \dots, \mathbf{m}_4$ are set to zero. 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 (44) 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

$\{\mathbf{x}_{t-1}\}_{t=1}^{T_1}$ containing $m = 12$ lagged values of \mathbf{y} and the constant term. For the second sample (1975:M2 to 2016: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 E 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

$$\begin{aligned} \tilde{\mathbf{Y}}_i(\mathbf{A}) &= (\sqrt{\mu}\mathbf{y}'_1\mathbf{a}_i, \dots, \sqrt{\mu}\mathbf{y}'_{T_1}\mathbf{a}_i, \mathbf{y}'_{T_1+1}\mathbf{a}_i, \dots, \mathbf{y}'_{T_1+T_2}\mathbf{a}_i, \mathbf{m}_i'\mathbf{P})' \\ & \quad (T_1+T_2+k) \times 1 \\ \tilde{\mathbf{X}} &= \begin{bmatrix} \sqrt{\mu}\mathbf{x}_0 & \cdots & \sqrt{\mu}\mathbf{x}_{T_1-1} & \mathbf{x}_{T_1} & \cdots & \mathbf{x}'_{T_1+T_2-1} & \mathbf{P} \end{bmatrix}' \\ & \quad (T_1+T_2+k) \times k \\ \kappa_i^* &= \kappa_i + (\mu T_1 + T_2)/2 \end{aligned}$$

for \mathbf{P} the matrix whose diagonal elements are reciprocals of the square roots of (44). We then calculate $\tau_i^*(\mathbf{A})$ and $\zeta_i^*(\mathbf{A})$ using expressions (6) and (7) and replace (12) and (13) with

$$\begin{aligned} p(\mathbf{A}|\mathbf{Y}_T) &= \frac{k_T p(\mathbf{A}) [\det(\mathbf{A}\tilde{\boldsymbol{\Omega}}_T\mathbf{A}')]^{(\mu T_1+T_2)/2}}{\prod_{i=1}^n [2\tau_i^*(\mathbf{A})/(\mu T_1 + T_2)]^{\kappa_i^*}} \prod_{i=1}^n \tau_i(\mathbf{A})^{\kappa_i} \\ \tilde{\boldsymbol{\Omega}}_T &= (\mu T_1 + T_2)^{-1} (\mu \boldsymbol{\zeta}^{(1)} + \boldsymbol{\zeta}^{(2)}) \\ \boldsymbol{\zeta}^{(1)} &= \sum_{t=1}^{T_1} \mathbf{y}_t \mathbf{y}'_t - \left(\sum_{t=1}^{T_1} \mathbf{y}_t \mathbf{x}'_{t-1} \right) \left(\sum_{t=1}^{T_1} \mathbf{x}_{t-1} \mathbf{x}'_{t-1} \right)^{-1} \left(\sum_{t=1}^{T_1} \mathbf{x}_{t-1} \mathbf{y}'_t \right) \\ \boldsymbol{\zeta}^{(2)} &= \sum_{t=T_1+1}^{T_2} \mathbf{y}_t \mathbf{y}'_t - \left(\sum_{t=T_1+1}^{T_2} \mathbf{y}_t \mathbf{x}'_{t-1} \right) \left(\sum_{t=T_1+1}^{T_2} \mathbf{x}_{t-1} \mathbf{x}'_{t-1} \right)^{-1} \left(\sum_{t=T_1+1}^{T_2} \mathbf{x}_{t-1} \mathbf{y}'_t \right). \end{aligned}$$

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

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

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

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

Thus the mean for the prior used to analyze the second sample ($\mathbf{a}'_i \hat{\Phi}^{(1)}$) 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

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

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

5.1.5 Comparison with Kilian and Murphy (2014).

Kilian and Murphy (2014) also used a 4-variable model of the global oil market to address some of the questions considered in our paper. Here we summarize the key differences between our approaches.

First, Kilian and Murphy (2014), like Kilian and Murphy (2012), assumed that the supply elasticity is known with certainty to be below 0.0258, a claim of prior knowledge that we regard as unfounded for reasons detailed in Section 3.3. Kilian and Murphy (2014) imposed a large number of such restrictions as if they were all known with certainty. Running their code exactly as posted on the journal's website generates 5 million draws of which only 16 satisfy all of the restrictions they imposed. The figures in their article came from somewhat arbitrarily choosing a single "representative" model among these 16. Even if one accepted that notion, it poses practical problems for researchers wishing to replicate or update their findings; in fact, one arrives at different results even on their original data set if the code is simply rerun with a different starting seed for the random number generator (see Figure 8). By contrast, all of the results we report below are based on 1 million draws, for which magnitudes like the posterior median and posterior credibility sets are well motivated and well estimated summary statistics. Furthermore, our summary statistics directly acknowledge both estimation error and uncertainty about the identification itself.

Second, Kilian and Murphy measured aggregate economic activity using a detrended average commodity shipping cost. This series has exhibited erratic behavior in the data since their study was conducted that likely reflects volatility in the ship-building and scrapping cycle since the Great Recession and the collapse in iron ore shipments. Indeed, if one insisted

that this remains a reliable measure of world economic activity, it would force us to conclude that the decline in world economic activity in 2015-2016 was even more severe than in the financial crisis of the Great Recession (see Figure 1 in Kilian and Zhou, 2017). Our preferred variable, the log of world industrial production, is measured in standard economic units. This allows us to estimate the percent change in world industrial production that would result from a supply shock of specified size, something that Kilian and Murphy could not do.

Other important differences are Kilian and Murphy (2014) treated global oil inventories as if observed without measurement error. Our analysis below suggests that measurement error in this series is quite important. They also completely threw out data before 1975. Our approach downweights these data but does not ignore them altogether.

Despite these major methodological differences, some of our conclusions are similar to those reached by Kilian and Murphy (2014). We highlight below points on which our findings reinforce and where they raise doubts about conclusions from the previous literature.

5.2 Empirical results.

The solid red curves in Figure 7 denote different components of the prior information about the contemporaneous coefficients in \mathbf{A} on which our analysis draws. The posterior distributions with pre-1975 observations downweighted by $\mu = 0.5$ are reported as blue histograms.⁹

The posterior median of the short-run price elasticity of oil supply, α_{qp} , is 0.15, a little above Caldara, Cavallo and Iacoviello’s (2017) estimate of 0.11. 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 price elasticity of oil demand, β_{qp} , is -0.35 , significantly more elastic than anticipated by our prior. Values for these and several other magnitudes of interest are reported in column 1 of Table 3.

Posterior structural impulse-response functions are plotted in Figure 9.¹⁰ 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. An oil supply shock also leads to a decline in economic activity. The effect on impact is practically zero (see the (1,2) panel of Figure 7), but accumulates over time (the (1,2) panel of Figure 9), a conclusion consistent with a large number of studies going back to Hamilton (1983). Our estimates imply that a reduction in oil production that raises the oil price by 10% would lower world economic activity by 0.5% after a year. 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. A similar conclusion was reached by Kilian (2009) and Kilian and Murphy (2012, 2014), though it is

⁹Data and code to replicate these results are available at https://sites.google.com/site/cjsbaumeister/BH2_code_web.zip.

¹⁰Note following standard practice these are accumulated impulse-response functions, plotting elements of $(\mathbf{H}_0^* + \mathbf{H}_1^* + \dots + \mathbf{H}_s^*)$ as a function of s where \mathbf{H}_s^* is given in (39). For example, panel (1,3) shows the effect on the level of oil prices s periods after an oil supply shock.

a little harder to interpret the finding in their exercise due to the indirect nature of their proxy for world 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 a negative effect on economic activity as well.

Figure 10 shows the historical decomposition of oil price movements along with 95% credibility regions.¹¹ Column 3 of Table 4 summarizes the contribution of supply shocks to several historical episodes of interest. Whereas Kilian (2009) and Kilian and Murphy (2012) concluded that the supply disruptions associated with the First Persian Gulf War played little direct role in the price increase, we find supply and demand shocks to have been equally important in this episode; a similar conclusion was reached by Kilian and Murphy (2014). We also find that accumulated supply shocks over 2007-2008 (showing up as an unexpected stagnation in global oil production) also accounted for much of the oil price run-up over that period, consistent with the analysis in Hamilton (2009), but in contrast to the conclusion in Kilian and Murphy (2012, 2014). We find supply shocks were less than half the story behind the oil price collapse in 2014-2016. That finding differs from that of Baffes et al. (2015) but is consistent with the conclusions of Hamilton (2015) and Baumeister and Kilian (2016). Our estimates attribute most of the oil price rebound in 2016 to strong demand.

The results in the last panel of Figure 10 suggest that inventory demand shocks have played a much smaller role in price fluctuations than implied by the analysis in Kilian and Murphy (2012, 2014) and Juvenal and Petrella (2014). A key reason we reach a different conclusion from these earlier researchers is our allowance for the possibility that our measure of world inventories contains a lot of error. The posterior median of σ_e^2 is 0.99,¹² which is a sizeable fraction of the total variance of the reduced-form VAR forecast of observed inventory changes ($\hat{\omega}_{44} = 1.06$).¹³

Our overall conclusion is that speculation is less important, and shocks to fundamentals, in particular, shocks to supply, were more important, than was found in several previous studies.

6 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 effects of putting less weight on any particular components of the prior to see how it affects

¹¹The 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.

¹²We used equations (55) and (50) to calculate the value of $\sigma_e^2 = \rho\chi d_{33}$ associated with each draw of $\boldsymbol{\theta}_A$.

¹³Nor is it the case that our implied measurement error is imputed to have any significant effects. The share of the 4 price movements in Table 4 attributed to measurement error shocks from equation (61) is 4.7%, 1.4%, 3.1%, and 1.7%, respectively.

the results.

Table 3 presents the posterior median and posterior 68% credibility sets for some of the magnitudes of interest when we weaken different components of the prior. The first column presents results from the baseline specification that were just summarized. Panels A and B report inference about the short-run supply elasticity α_{qp} and demand elasticity β_{qp} . Panel C looks at the response of economic activity 12 months after a supply shock, panel D the response to an oil consumption demand shock, and panel E the response to an inventory demand shock, with each shock normalized for purposes of the table as an event that leads to a 10% increase in the real oil price at time 0. Note that this is a different normalization from that used in Figure 9, where the effect plotted was that of a one-unit change in the structural shock $\partial y_{i,t+s}/\partial u_{jt}^*$.

Table 4 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 four important episodes in the sample. Column 3 reports the posterior median and 68% 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

We next explored the consequences of using a much weaker prior for the short-run supply and demand elasticities, replacing the scale parameters $\sigma_{qp}^\alpha = \sigma_{qp}^\beta = 0.2$ that were used in the baseline analysis with the alternative values $\sigma_{qp}^\alpha = \sigma_{qp}^\beta = 1.0$. This change gives the prior for these two parameters a variance that is 25 times larger than in the baseline specification, with the result that prior information about these parameters is allowed to have very little influence on any of our conclusions. The implications of these changes for inference about key magnitudes are reported in column 2 of Table 3. If we had very little prior information about the elasticities themselves, we would tend to infer a smaller short-run supply elasticity (panel A) and more elastic demand (panel B). Our core conclusions about impulse responses (panels C-E) do not change.

Our prior beliefs about the role of measurement error were represented by the $\text{Beta}(\alpha_\chi, \beta_\chi)$ distribution for χ (which summarizes the ratio of OECD inventories to world total) and χ times a $\text{Beta}(\alpha_\rho, \beta_\rho)$ variable 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 0.12χ , 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 0.19χ , respectively. The implications of this weaker prior about the role of measurement error are reported in column 3 of Table 3 and column 5 of Table 4. These results are virtually identical to those for our

¹⁴Let p_t denote the 100 times the change in log oil price in month t (the dashed line in the top panel of Figure 10) and \hat{p}_{t_1} the value of the solid line for that date. The number reported in column 2 is $p_{t_0} + p_{t_0+1} + \dots + p_{t_1}$ for $t_0 = \text{July } 1990$ and $t_1 = \text{October } 1990$. The number reported in column 3 is $\hat{p}_{t_0,1} + \hat{p}_{t_0+1,1} + \dots + \hat{p}_{t_1,1}$.

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 differ only slightly from 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 (44). 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 3 and column 7 of Table 4. 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 3 and column 8 of Table 4 uses $\kappa = 0.5$. This does not matter for any of the conclusions.

Finally, in making use of the historical data we relied on WTI prices 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 4 reports the measured size of the oil price change recorded by WTI in four episodes of interest. The two oil prices 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 3) is the same regardless of which measure we use.

7 Conclusion.

Prior information has played a key role in any structural analysis of vector autoregressions. 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 data set 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 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. A key difference from earlier analyses is that supply shocks appear to be more important and speculative demand shocks less important than found by some earlier researchers. 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 information.

References

- Amir-Ahmadi, Pooyan, Christian Matthes, and Mu-Chun Wang (2017). "Measurement Errors and Monetary Policy: Then and Now," *Journal of Economic Dynamics and Control* 79: 66-78.
- Baffes, John, M. Ayhan Kose, Franziska Ohnsorge, and Marc Stocker (2015). "The Great Plunge in Oil Prices: Causes, Consequences, and Policy Responses," World Bank Policy Research Note.
- Baumeister, Christiane, and James D. Hamilton (2015). "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information," *Econometrica* 83: 1963-1999.
- Baumeister, Christiane, and James D. Hamilton (2017). "Optimal Inference about Impulse-Response Functions and Historical Decompositions in Incompletely Identified Structural Vector Autoregressions," working paper, UCSD.
- Baumeister, Christiane, and Lutz Kilian (2016). "Understanding the Decline in the Price of Oil Since June 2014," *Journal of the Association of Environmental and Resource Economists* 3: 131-158.
- 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.
- Bjørnland, Hilde C., Frode Martin Nordvik, and Maximilian Rohrer (2017). "Supply Flexibility in the Shale Patch: Evidence from North Dakota," working paper, Norwegian Business School.
- 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.
- Caldara, Dario, Michele Cavallo, and Matteo Iacoviello (2017). "Oil Price Elasticities and Oil Price Fluctuations," working paper, Federal Reserve Board.
- 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 (2017). "Anticipation, Tax Avoidance, and the Price Elasticity of Demand," *Journal of Applied Econometrics* 32: 1-15.
- 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, 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.

Gelman, Michael et al. (2017). "The Response of Consumer Spending to Changes in Gasoline Prices," NBER Working Paper 22969.

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. (2009). "Causes and consequences of the oil shock of 2007-08," *Brookings Papers on Economic Activity* Spring 2009: 215-261.

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

Hamilton, James D. (2015). "Common Factors in Commodity and Asset Markets," <http://econbrowser.com/archives/2015/09/common-factors-in-commodity-and-asset-markets>.

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

Hughes, Jonathan E., Christopher R. Knittel, and Daniel Sperling (2008), "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand," *Energy Journal* 29: 93-114.

Juvenal, Luciana, and Ivan Petrella (2014). "Speculation in the Oil Market," *Journal of Applied Econometrics* 30: 621-649.

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.

Kilian, Lutz, and Xiaoqing Zhou (2017). "Modeling Fluctuations in the Global Demand for Commodities," working paper, University of Michigan.

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 D|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 m th-order univariate autoregression fit to series i and $\hat{\mathbf{S}}$ 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 (2015) proposed setting κ_i/τ_i (the prior mean for d_{ii}^{-1}) equal to the reciprocal of the i th diagonal element of $\mathbf{A}'\hat{\mathbf{S}}\mathbf{A}$; in other words, $\tau_i(\mathbf{A}) = \kappa_i \mathbf{a}'_i \hat{\mathbf{S}} \mathbf{a}_i$. 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 B|A, 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 (2015, 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

$$\mathbf{v}'_1 = (1/(1^{2\lambda_1}), 1/(2^{2\lambda_1}), \dots, 1/(m^{2\lambda_1})) \quad (43)$$

(1×m)

$$\mathbf{v}'_2 = (s_{11}^{-1}, s_{22}^{-1}, \dots, s_{nn}^{-1})'$$

(1×n)

$$\mathbf{v}_3 = \lambda_0^2 \begin{bmatrix} \mathbf{v}_1 \otimes \mathbf{v}_2 \\ \lambda_3^2 \end{bmatrix}. \quad (44)$$

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

$$M_{i,rr} = v_{3r}. \quad (45)$$

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

$$q(\mathbf{A}) = \log(p(\mathbf{A})) + (T/2) \log \left[\det \left(\mathbf{A} \hat{\boldsymbol{\Omega}}_T \mathbf{A}' \right) \right] - \sum_{i=1}^n \kappa_i^* \log[(2/T) \tau_i^*(\mathbf{A})] + \sum_{i=1}^n \kappa_i \log \tau_i(\mathbf{A}). \quad (46)$$

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 (46) 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 $\{\boldsymbol{\alpha}^{(D+1)}, \boldsymbol{\alpha}^{(D+2)}, \dots, \boldsymbol{\alpha}^{(D+N)}\}$ 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 $i = 1, \dots, n$ 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 $i = 1, \dots, n$ and take $\mathbf{B}^{(\ell)}$ the matrix whose i th row is given by $\mathbf{b}_i^{(\ell)'$. 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/supinfo>) 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 Strategic Petroleum Reserve) are available from EIA for the entire period 1958:M1-2016: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. Adapting the algorithms in Baumeister and Hamilton (2015) to allow for measurement error.

Rewriting (34) in the form of (1). The variance matrix for the structural shocks in (34) is given by

$$\tilde{\mathbf{D}} = E(\tilde{\mathbf{u}}_t \tilde{\mathbf{u}}_t') = \begin{bmatrix} d_{11}^* & 0 & 0 & 0 \\ 0 & d_{22}^* & 0 & 0 \\ 0 & 0 & d_{33}^* + \chi^{-2}\sigma_e^2 & -\chi^{-1}\sigma_e^2 \\ 0 & 0 & -\chi^{-1}\sigma_e^2 & \chi^2 d_{44}^* + \sigma_e^2 \end{bmatrix}. \quad (47)$$

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

$$\mathbf{\Gamma} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \rho & 1 \end{bmatrix} \quad (48)$$

with ρ given by (37). Thus if we premultiply (34) by $\mathbf{\Gamma}$ we arrive at a system in the form of (1) for which $\mathbf{A} = \mathbf{\Gamma} \tilde{\mathbf{A}}$, $\mathbf{B} = \mathbf{\Gamma} \tilde{\mathbf{B}}$, and

$$\mathbf{u}_t = \mathbf{\Gamma} \tilde{\mathbf{u}}_t = \begin{bmatrix} u_{1t}^* \\ u_{2t}^* \\ u_{3t}^* - \chi^{-1}e_t \\ \chi u_{4t}^* + \rho u_{3t}^* + (1 - \rho/\chi)e_t \end{bmatrix} \quad (49)$$

whose variance matrix we denote $\mathbf{D} = \text{diag}(d_{11}, d_{22}, d_{33}, d_{44})$. This is exactly in the form of the general class of models discussed in Section 2 with the elements of the matrix \mathbf{A} determined by $\boldsymbol{\theta}_{\mathbf{A}} = (\alpha_{qp}, \alpha_{yp}, \beta_{qy}, \beta_{qp}, \chi, \psi_1, \psi_3, \rho)'$. Given a prior distribution for these parameters, we can then draw from the posterior distribution $p(\boldsymbol{\theta}_{\mathbf{A}}, \mathbf{D}, \mathbf{B} | \mathbf{Y}_T)$ as described in Appendix B.

Given a draw from $p(\boldsymbol{\theta}_{\mathbf{A}}, \mathbf{D}, \mathbf{B} | \mathbf{Y}_T)$ we immediately have a draw for $\tilde{\mathbf{A}} = \mathbf{\Gamma}^{-1} \mathbf{A}$, $\tilde{\mathbf{B}} = \mathbf{\Gamma}^{-1} \mathbf{B}$, and $\tilde{\mathbf{D}} = \mathbf{\Gamma}^{-1} \mathbf{D} (\mathbf{\Gamma}^{-1})'$. The lower-right (2×2) block of the last equation is

$$\begin{bmatrix} \tilde{d}_{33} & \tilde{d}_{34} \\ \tilde{d}_{43} & \tilde{d}_{44} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ -\rho & 1 \end{bmatrix} \begin{bmatrix} d_{33} & 0 \\ 0 & d_{44} \end{bmatrix} \begin{bmatrix} 1 & -\rho \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} d_{33} & -\rho d_{33} \\ -\rho d_{33} & d_{44} + \rho^2 d_{33} \end{bmatrix}. \quad (50)$$

Historical decompositions. Recall $\boldsymbol{\epsilon}_t = \tilde{\mathbf{A}}^{-1} \tilde{\mathbf{u}}_t$ for $\tilde{\mathbf{u}}_t$ defined in (36). Taking expectations of both sides conditional on the data and on a draw of the parameters, the following

equation holds exactly for every t ,

$$\boldsymbol{\epsilon}_t = \tilde{\mathbf{A}}^{-1} \begin{bmatrix} u_{1t}^* \\ u_{2t}^* \\ E(u_{3t}^* | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) - \chi^{-1} E(e_t | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \\ \chi E(u_{4t}^* | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) + E(e_t | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \end{bmatrix}, \quad (51)$$

where

$$\begin{aligned} \begin{bmatrix} E(u_{3t}^* | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \\ E(u_{4t}^* | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \\ E(e_t | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \end{bmatrix} &= \begin{bmatrix} E(u_{3t}^* | \tilde{u}_{3t}, \tilde{u}_{4t}, \boldsymbol{\theta}_A, \mathbf{D}) \\ E(u_{4t}^* | \tilde{u}_{3t}, \tilde{u}_{4t}, \boldsymbol{\theta}_A, \mathbf{D}) \\ E(e_t | \tilde{u}_{3t}, \tilde{u}_{4t}, \boldsymbol{\theta}_A, \mathbf{D}) \end{bmatrix} \\ &= E \left(\begin{bmatrix} u_{3t}^* \tilde{u}_{3t} & u_{3t}^* \tilde{u}_{4t} \\ u_{4t}^* \tilde{u}_{3t} & u_{4t}^* \tilde{u}_{4t} \\ e_t \tilde{u}_{3t} & e_t \tilde{u}_{4t} \end{bmatrix} \middle| \boldsymbol{\theta}_A, \mathbf{D} \right) \left\{ E \left(\begin{bmatrix} \tilde{u}_{3t}^2 & \tilde{u}_{3t} \tilde{u}_{4t} \\ \tilde{u}_{4t} \tilde{u}_{3t} & \tilde{u}_{4t}^2 \end{bmatrix} \middle| \boldsymbol{\theta}_A, \mathbf{D} \right) \right\}^{-1} \begin{bmatrix} \tilde{u}_{3t} \\ \tilde{u}_{4t} \end{bmatrix}. \end{aligned} \quad (52)$$

The second matrix in (52) is known from (50) while the first matrix can be calculated from the equations

$$\tilde{u}_{3t} = u_{3t}^* - \chi^{-1} e_t$$

$$\tilde{u}_{4t} = \chi u_{4t}^* + e_t$$

and the fact that the three disturbances u_{3t}^* , u_{4t}^* , and e_t are mutually uncorrelated:

$$E(\tilde{u}_{3t} u_{4t}^*) = E(\tilde{u}_{4t} u_{3t}^*) = 0 \quad (53)$$

$$E(\tilde{u}_{3t} e_t) = -\chi^{-1} \sigma_e^2 = E(\tilde{u}_{3t} \tilde{u}_{4t}) \quad (54)$$

$$E(\tilde{u}_{4t} e_t) = \sigma_e^2 = -\chi E(\tilde{u}_{3t} \tilde{u}_{4t}) \quad (55)$$

$$E(\tilde{u}_{3t} u_{3t}^*) = E(\tilde{u}_{3t})(\tilde{u}_{3t} + \chi^{-1} e_t) = E(\tilde{u}_{3t}^2) + \chi^{-1} E(\tilde{u}_{3t} \tilde{u}_{4t}) \quad (56)$$

$$\chi E(\tilde{u}_{4t} u_{4t}^*) = E(\tilde{u}_{4t})(\tilde{u}_{4t} - e_t) = E(\tilde{u}_{4t}^2) + \chi E(\tilde{u}_{3t} \tilde{u}_{4t}). \quad (57)$$

Collecting (53)-(57) into a matrix equation,

$$\begin{aligned} E \left(\begin{bmatrix} u_{3t}^* \tilde{u}_{3t} & u_{3t}^* \tilde{u}_{4t} \\ u_{4t}^* \tilde{u}_{3t} & u_{4t}^* \tilde{u}_{4t} \\ e_t \tilde{u}_{3t} & e_t \tilde{u}_{4t} \end{bmatrix} \middle| \boldsymbol{\theta}_A, \mathbf{D} \right) &= \begin{bmatrix} \tilde{d}_{33} + \chi^{-1} \tilde{d}_{34} & 0 \\ 0 & (\tilde{d}_{44} + \chi \tilde{d}_{34}) / \chi \\ \tilde{d}_{34} & -\chi \tilde{d}_{34} \end{bmatrix} \\ &= \begin{bmatrix} d_{33}(1 - \rho / \chi) & 0 \\ 0 & \frac{d_{44} + \rho(\rho - \chi)d_{33}}{\chi} \\ -\rho d_{33} & \rho \chi d_{33} \end{bmatrix} \end{aligned} \quad (58)$$

where the last equation follows from (50). Substituting (50) and (58) into (52),

$$\begin{aligned} \begin{bmatrix} E(u_{3t}^* | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \\ E(u_{4t}^* | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \\ E(e_t | \tilde{\mathbf{u}}_t, \boldsymbol{\theta}_A, \mathbf{D}) \end{bmatrix} &= \begin{bmatrix} d_{33}(1 - \frac{\rho}{\chi}) & 0 \\ 0 & \frac{d_{44} + \rho(\rho - \chi)d_{33}}{\chi} \\ -\rho d_{33} & \rho\chi d_{33} \end{bmatrix} \begin{bmatrix} d_{33} & -\rho d_{33} \\ -\rho d_{33} & d_{44} + \rho^2 d_{33} \end{bmatrix}^{-1} \begin{bmatrix} \tilde{u}_{3t} \\ \tilde{u}_{4t} \end{bmatrix} \\ &= \begin{bmatrix} h_{33} & h_{34} \\ h_{43} & h_{44} \\ h_{e3} & h_{e4} \end{bmatrix} \begin{bmatrix} \tilde{u}_{3t} \\ \tilde{u}_{4t} \end{bmatrix} \end{aligned} \quad (59)$$

Historical decompositions can thus be calculated as follows. The value of \mathbf{y}_t can be written as the r -period-ahead forecast plus a known function of the forecast errors between $t - r$ and t :

$$\mathbf{y}_t = \hat{\mathbf{y}}_{t|t-r} + \sum_{s=0}^{r-1} \boldsymbol{\Psi}_s \boldsymbol{\epsilon}_{t-s}. \quad (60)$$

We have inferred values for $\tilde{\mathbf{u}}_t$ for each date from $\tilde{\mathbf{u}}_t = \tilde{\mathbf{A}}\mathbf{y}_t - \tilde{\mathbf{B}}\mathbf{x}_{t-1}$. Hence

$$\mathbf{y}_t = \hat{\mathbf{y}}_{t|t-r} + \sum_{s=0}^{r-1} \boldsymbol{\Psi}_s \tilde{\mathbf{A}}^{-1} \begin{bmatrix} u_{1,t-s}^* \\ u_{2,t-s}^* \\ \hat{u}_{3,t-s}^* - \chi^{-1} \hat{e}_{t-s} \\ \chi \hat{u}_{4,t-s}^* + \hat{e}_{t-s} \end{bmatrix}$$

where $(u_{1t}^*, u_{2t}^*)' = (\tilde{u}_{1t}, \tilde{u}_{2t})'$ and the vector $(\hat{u}_{3t}^*, \hat{u}_{4t}^*, \hat{e}_t)'$ is calculated from (59). With this expression we can calculate the contribution of each of the shocks $(u_1^*, u_2^*, u_3^*, u_4^*, e)$ to the historical value of \mathbf{y}_t . For example, the historical contribution of inventory demand shocks (u_{4t}^*) to \mathbf{y}_t is found from $\sum_{s=0}^{r-1} \boldsymbol{\Psi}_s \tilde{\mathbf{A}}^{-1} \mathbf{h}_{t-s}^{4*}$ for $\mathbf{h}_t^{4*} = (0, 0, 0, \chi(h_{43}\tilde{u}_{3t} + h_{44}\tilde{u}_{4t}))'$, while the contribution of measurement error is

$$\sum_{s=0}^{r-1} \boldsymbol{\Psi}_s \tilde{\mathbf{A}}^{-1} \mathbf{h}_{t-s}^e \quad (61)$$

for $\mathbf{h}_t^e = (0, 0, -\chi^{-1}(h_{e3}\tilde{u}_{3t} + h_{e4}\tilde{u}_{4t}), h_{e3}\tilde{u}_{3t} + h_{e4}\tilde{u}_{4t})'$.

E. 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

$$p(\boldsymbol{\lambda} | \mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}) = \frac{p(\mathbf{Y}^{(2)} | \mathbf{Y}^{(1)}, \boldsymbol{\lambda}) p(\mathbf{Y}^{(1)} | \boldsymbol{\lambda}) p(\boldsymbol{\lambda})}{\int p(\mathbf{Y}^{(2)} | \mathbf{Y}^{(1)}, \boldsymbol{\lambda}) p(\mathbf{Y}^{(1)} | \boldsymbol{\lambda}) p(\boldsymbol{\lambda}) d\boldsymbol{\lambda}}. \quad (62)$$

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 (62) 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,

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

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 (2015) for this downweighted likelihood leads to the algorithm described in Section 5.1.4.

F. The influence of prior information about the supply elasticity.

Our baseline prior implies a 6% probability that the supply elasticity α_{qp} is below 0.0258.¹⁵ This could be misinterpreted to suggest that our prior imposes a big value for the elasticity. The reason this conclusion is wrong is that it is the variance of the prior (concentration of mass over any fixed interval), not the probability of exceeding some specified bound, that determines the influence of the prior. The Bayesian posterior distribution is a weighted average of the likelihood, with weights given by the prior density. If the prior density has a very large variance, the weights are approximately uniform over the range for which the likelihood has nonnegligible mass, and the posterior is essentially the same as the likelihood, with the prior exerting no influence on the posterior.

We illustrate this with a simple parametric example, in which the data mildly favor a value

¹⁵Note for $c = 0.1$ and $\sigma = 0.2$ that

$$\frac{[\tilde{\Phi}_3((0.0258 - c)/\sigma) - \tilde{\Phi}_3(-c/\sigma)]}{1 - \tilde{\Phi}_3(-c/\sigma)} = 0.062$$

where $\tilde{\Phi}_3(x)$ denotes the probability that a standard Student t variable with $\nu = 3$ degrees of freedom would be less than some value x .

of $\alpha = 0.05$. Our baseline prior results in a posterior that is virtually identical to that implied by the data (panel 3 in Figure F-1). By contrast, the Kilian-Murphy (KM) prior (insisting that $\alpha < 0.0258$) hugely distorts the data (panel 4). This is because our prior has a big variance and the KM prior has a tiny variance. The KM prior imposes a strong prior belief, whereas ours does not.

Details of the example that produced this figure are as follows. Suppose that the parameter of interest α is the variance of a $N(0, \alpha)$ distribution; we use this as a simple example of a parameter that has to be positive. Suppose we have observed a sample y_1, \dots, y_T of size $T = 10$ and that the average squared value of y_t in the observed sample is $s^2 = 0.05$. Thus the likelihood function is

$$f(y|\alpha) = \frac{1}{(2\pi\alpha)^{T/2}} \exp\left(-\frac{Ts^2}{2\alpha}\right) \quad (63)$$

and the MLE is $\hat{\alpha} = 0.05$. The likelihood is plotted as a function of α as the solid black curve in Figure F-1.¹⁶ Suppose first that our prior takes the form of a Student $t(c, \sigma, \nu)$ distribution truncated to be positive with $c = 0.1$ and $\nu = 3$. The first three panels of the graph correspond to three different truncated Student t priors plotted as dotted red curves, the first with a relatively small variance ($\sigma = 0.02$) the second with a somewhat bigger variance ($\sigma = 0.05$), and the third with a still bigger variance. In fact, the third distribution has exactly the same variance and exactly the same parameters as the prior in our baseline parameterization ($\sigma = 0.2$). The posterior associated with each prior is plotted in dashed blue.¹⁷ The prior with a small variance has significant impact on the posterior inference about α . The prior with a bigger variance has a much smaller impact, and the prior with the variance as in our baseline specification has zero influence on the posterior distribution in this example. This is true even though the latter prior only assigns a 12% probability to a value lower than 0.05. What matters for how much influence the prior has on the conclusion depends on the variance and bounds of the prior distribution. Our prior has a relatively large variance, and imposes no bounds, and in this example has essentially zero measurable impact on the posterior distribution.

By contrast, the implicit prior used by Kilian and Murphy is the uniform distribution over

¹⁶For ease of visual comparison with the prior and posterior, we divide (63) by the sum of the values over α between 0 and 1, so that the likelihood (like the prior and posterior that will also be plotted) integrates to unity with respect to α .

¹⁷If $p(\alpha)$ denotes the prior and $f(y|\alpha)$ the likelihood, this was evaluated at any point α by

$$p(\alpha|y) = \frac{p(\alpha)f(y|\alpha)}{\sum_{x=1}^N p(x/N)f(y|x/N)}$$

for $N = 1000$. Note this results in the numerically identical posterior $p(\alpha|y)$ for any value of the normalizing constant k used in the previous footnote, since numerator and denominator both get multiplied by k .

$(0, 0.0258)$ plotted as the dotted red line in the fourth panel. The variance of the KM prior is

$$\frac{(0.0258)^2}{12} = 0.000005547.$$

In addition to having a tiny variance, this distribution dogmatically rules out any possibility of $\alpha > 0.0258$, and so by force imposes this condition on the posterior distribution. As a result, the KM prior hugely distorts the posterior inference for this example, as seen in the fourth panel of Figure F-1.

Table 1. Prior distributions for model parameters **A**, **B** and **D**.

Parameter	Meaning	Location	Scale	Degrees of freedom	Skew	Sign restriction
Priors affecting contemporaneous coefficients A						
Student <i>t</i> distribution						
α_{qp}	Oil supply elasticity	0.1	0.2	3	--	$\alpha_{qp} > 0$
α_{yp}	Effect of p on economic activity	-0.05	0.1	3	--	$\alpha_{yp} < 0$
β_{qy}	Income elasticity of oil demand	0.7	0.2	3	--	$\beta_{qy} > 0$
β_{qp}	Oil demand elasticity	-0.1	0.2	3	--	$\beta_{qp} < 0$
ψ_1	Effect of q on oil inventories	0	0.5	3	--	none
ψ_3	Effect of p on oil inventories	0	0.5	3	--	none
h_2	Effect of economic activity shock on y	0.8	0.2	3	--	none
Beta distribution						
χ	Fraction of inventories observed	0.6	0.1	--	--	$0 \leq \chi \leq 1$
ρ	Importance of inventory measurement error	0.25χ	0.12χ	--	--	$0 \leq \rho \leq \chi$
Asymmetric <i>t</i> distribution						
h_1	Determinant of $\tilde{\mathbf{A}}$	0.6	1.6	3	2	none
Priors for structural variances D A						
Gamma distribution						
d_{ii}^{-1}	Reciprocal of variance	$1/(\mathbf{a}'_i \mathbf{S} \mathbf{a}_i)$	$1/(\sqrt{2} \mathbf{a}'_i \mathbf{S} \mathbf{a}_i)$	--	--	$d_{ii} > 0$
Priors for lagged structural coefficients B A,D						
Normal distribution						
b_{13}	Lagged supply response	0.1	eq (44)	--	--	none
b_{33}	Lagged demand response	-0.1	eq (44)	--	--	none
b_{ij}	All other	0	eq (44)	--	--	none

Notes to Table 1. For Student *t* and Normal distributions the location parameter refers to the mode; for Beta and Gamma distributions the location parameter is the mean and the scale parameter is the standard deviation.

Table 2. Prior and posterior probabilities that the impact of a specified structural shock on the indicated variable is positive.

Variable	<i>Oil supply shock</i>		<i>Economic activity shock</i>		<i>Oil consumption demand shock</i>		<i>Oil inventory demand shock</i>	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Prior	Posterior	Prior	Posterior	Prior	Posterior	Prior	Posterior
q	0.915	1.000	0.973	1.000	0.973	1.000	0.973	1.000
y	0.859	1.000	1.000	1.000	0.027	0.000	0.027	0.000
p	0.141	0.000	0.973	1.000	0.973	1.000	0.973	1.000
Δi	0.696	0.200	0.234	0.165	0.234	0.165	0.973	1.000

Table 3. Sensitivity of parameter inference when less weight is placed on various components of the prior.

Benchmark	Supply and demand elasticities	Measurement error	Pre-1975 data	Lagged structural coefficients	Variances of shocks	Replace RAC with WTI
(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>A. Short-run price elasticity of oil supply α_{qp}</i>						
0.15	0.11	0.15	0.14	0.14	0.15	0.13
(0.09, 0.22)	(0.06, 0.19)	(0.09, 0.24)	(0.09, 0.23)	(0.10, 0.22)	(0.10, 0.23)	(0.09, 0.21)
<i>B. Short-run price elasticity of oil demand β_{qp}</i>						
-0.35	-0.47	-0.35	-0.35	-0.35	-0.35	-0.31
(-0.51, -0.24)	(-0.78, -0.28)	(-0.53, -0.23)	(-0.50, -0.24)	(-0.48, -0.24)	(-0.50, -0.24)	(-0.45, -0.20)
<i>C. Effect of oil supply shock that raises real oil price by 10% on economic activity 12 months later</i>						
-0.50	-0.62	-0.50	-0.35	-0.52	-0.49	-0.55
(-0.91, -0.17)	(-1.20, -0.22)	(-0.92, -0.16)	(-0.74, -0.03)	(-0.92, -0.16)	(-0.89, -0.16)	(-0.91, -0.24)
<i>D. Effect of oil consumption demand shock that raises real oil price by 10% on economic activity 12 months later</i>						
0.13	0.05	0.14	0.21	0.04	0.14	0.02
(-0.14, 0.44)	(-0.21, 0.36)	(-0.14, 0.49)	(-0.05, 0.51)	(-0.25, 0.37)	(-0.13, 0.46)	(-0.22, 0.31)
<i>E. Effect of oil inventory demand shock that raises real oil price by 10% on economic activity 12 months later</i>						
-0.36	-0.46	-0.35	-0.14	-0.55	-0.35	-0.41
(-0.81, 0.07)	(-1.02, -0.02)	(-0.92, 0.12)	(-0.57, 0.28)	(-1.04, -0.09)	(-0.79, 0.08)	(-0.83, -0.02)

Notes to Table 3. Table reports posterior median (in bold) and 68% credibility regions (in parentheses) for indicated magnitudes. Baseline uses priors specified in Table 1. Alternatives put less weight on indicated component of the prior as detailed in the text.

Table 4. Effects on oil prices attributed to supply shocks 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)
June-Oct 90	74.83	34.47 (23.84, 46.62)	24.66 (13.09, 39.41)	35.27 (23.58, 49.10)	33.86 (22.61, 47.29)	33.79 (23.56, 45.71)	35.07 (23.71, 48.76)	72.92	35.90 (25.09, 48.66)
		<i>46.1%</i>	<i>33.0%</i>	<i>47.1%</i>	<i>45.3%</i>	<i>45.2%</i>	<i>46.9%</i>		<i>49.2%</i>
Jan 07-June 08	86.80	40.89 (26.89, 56.84)	29.24 (14.76, 47.98)	41.82 (26.67, 59.59)	38.82 (24.23, 56.56)	43.25 (29.02, 59.39)	41.59 (26.79, 59.50)	-83.12	39.25 (25.48, 55.77)
		<i>47.1%</i>	<i>33.7%</i>	<i>48.2%</i>	<i>44.7%</i>	<i>49.8%</i>	<i>47.9%</i>		<i>47.2%</i>
June 14-Jan 16	-129.88	-49.52 (-71.35, -31.31)	-33.44 (-58.70, -15.15)	-50.86 (-75.75, -30.90)	-50.91 (-74.85, -31.52)	-49.61 (-70.87, -31.84)	-50.58 (-75.23, -31.16)	-121.03	-45.89 (-67.78, -28.32)
		<i>38.1%</i>	<i>25.7%</i>	<i>37.3%</i>	<i>39.2%</i>	<i>38.2%</i>	<i>38.9%</i>		<i>37.9%</i>
Feb-Dec 16	54.01	16.61 (8.48, 27.15)	9.68 (2.63, 20.98)	17.19 (8.35, 29.32)	15.90 (7.45, 27.63)	16.92 (8.69, 27.64)	17.06 (8.44, 29.11)	51.80	17.39 (9.15, 28.73)
		<i>30.7%</i>	<i>17.9%</i>	<i>31.8%</i>	<i>29.4%</i>	<i>31.3%</i>	<i>31.6%</i>		<i>33.6%</i>

Notes to Table 4. Table reports posterior median (in bold), 68% credibility regions (in parentheses) and percent contribution accounted for by oil supply shocks (in italics). Baseline uses priors specified in Table 1. Alternatives put less weight on indicated component of the prior as detailed in the text.

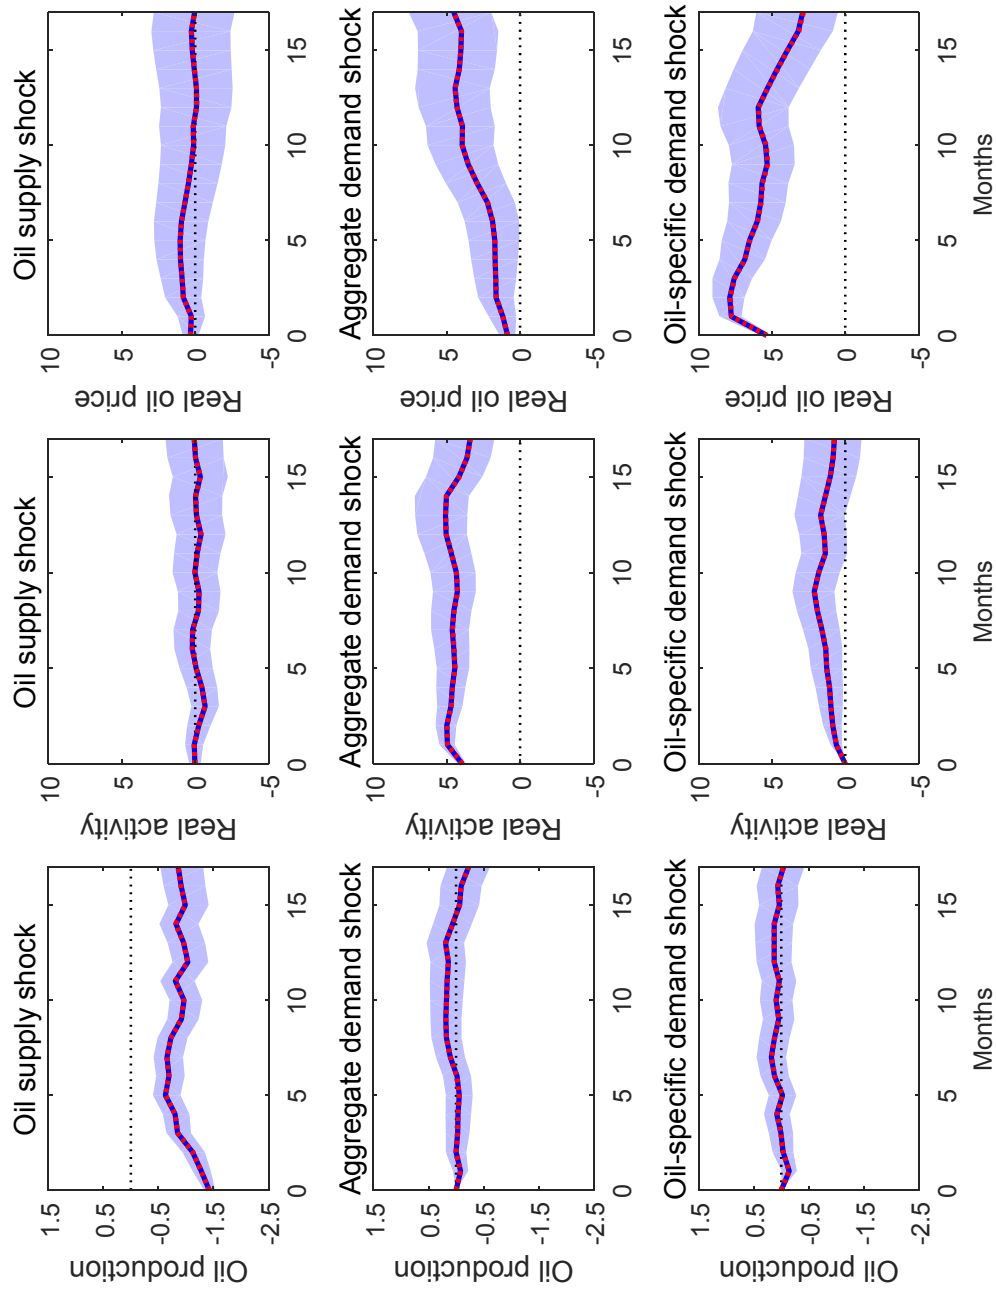


Figure 1. Impulse-response functions for 3-variable model under traditional Cholesky identification. Red dashed lines: point estimates arrived at using Kilian's (2009) original methodology; blue solid lines: median of Bayesian posterior distribution; shaded regions: 95% posterior credible set.




Figure 2. Prior (red lines) and posterior (blue histograms) distributions for the three unknown elements in **A** using traditional 3-variable Cholesky-type identification and posterior distribution for the implied short-run price elasticity of oil demand (lower right panel).




Figure 3. Prior (red lines) and posterior (blue histograms) distributions for the unknown elements in **A** and **H** in the Bayesian implementation of the 3-variable set-identified model of Kilian and Murphy (2012) and posterior distribution for the short-run price elasticity of oil demand (lower right panel).




Figure 4. Impulse-response functions for 3-variable model set identified using sign restrictions and bounds. Red dotted 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. 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 1.1.a (<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 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. Baseline prior (solid red curves) and posterior (blue histograms) distributions concerning the contemporaneous coefficients in **A** in baseline 4-variable model.




Figure 8. Effects of “speculative oil demand shock” for the Kilian and Murphy (2014) specification and data set using two different seeds for the random number generator. Left panel: effect on real activity. Right panel: effect on real price of oil. Red dotted lines: seed = 316, which was the original seed used by Kilian and Murphy and which reproduces panels (3,2) and (3,3) In Kilian and Murphy’s Figure 1. Blue solid lines: seed = 613.

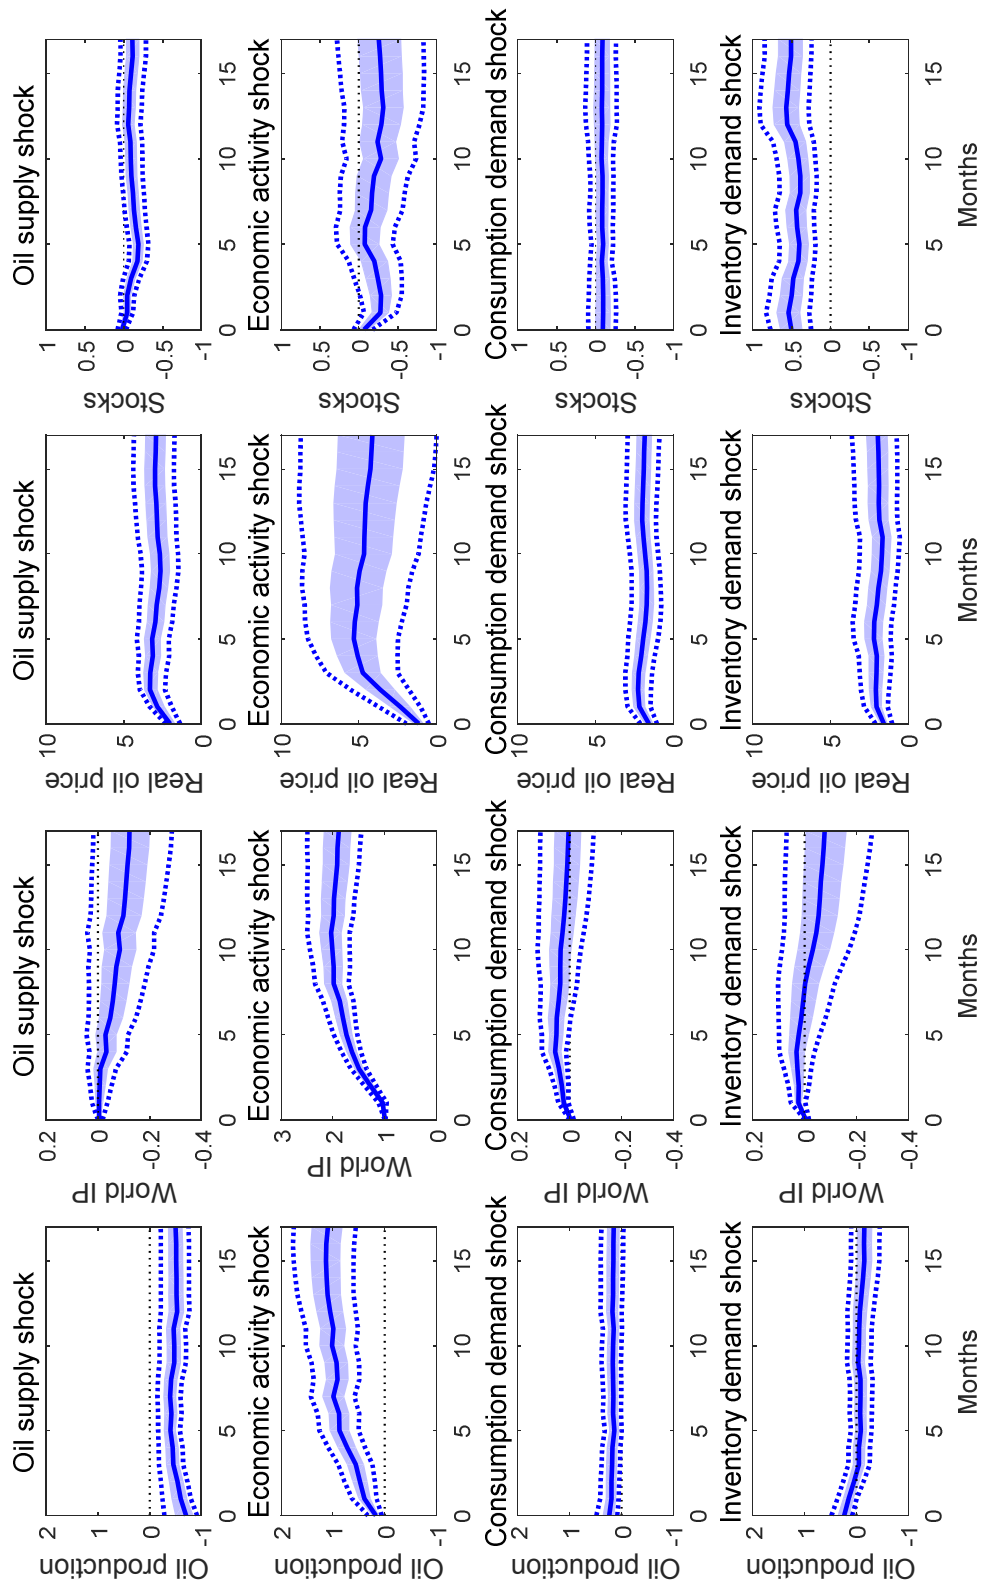


Figure 9. Impulse-response functions for baseline 4-variable model. Solid lines: Bayesian posterior median; shaded regions: 68% posterior credible sets; dotted lines: 95% posterior credible sets.




Figure 10. Actual changes in oil prices (red dashed lines) and historical contribution of separate structural shocks with 95% posterior credibility regions (blue and shaded) for baseline 4-variable model.

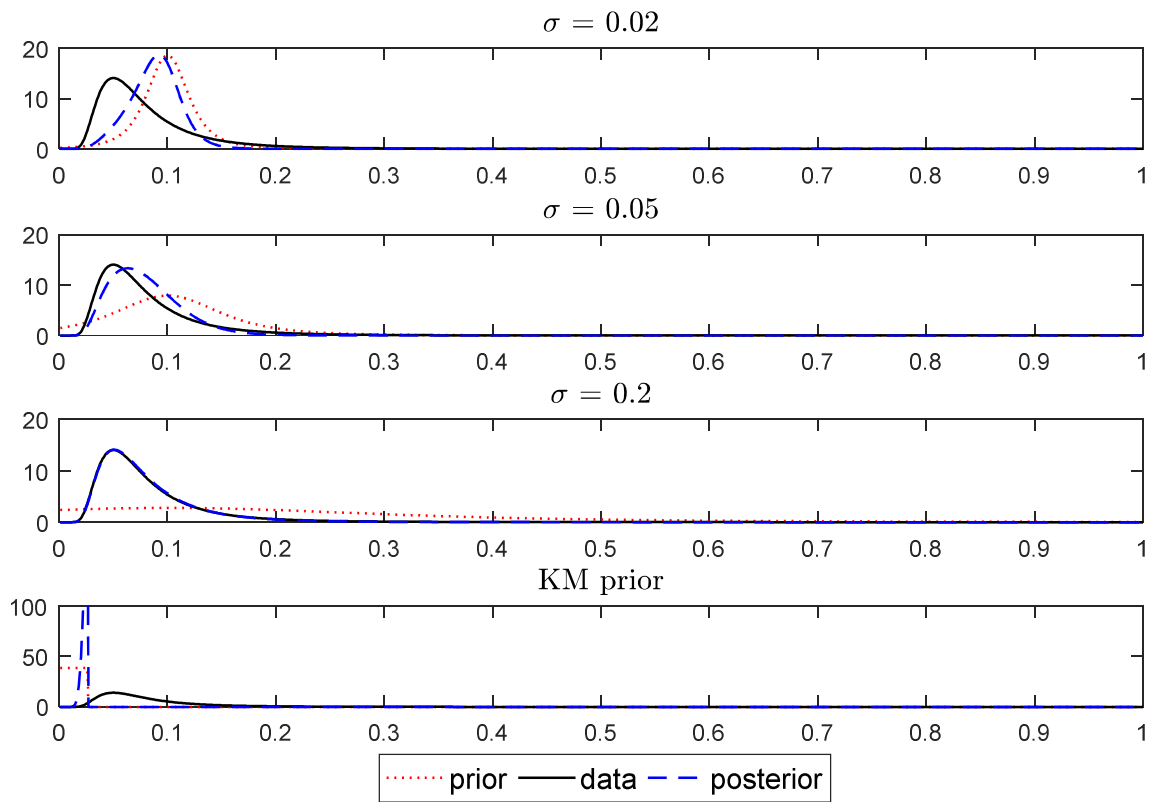


Figure F-1. Prior and posterior distributions for three different examples of truncated Student t priors and for Kilian-Murphy uniform prior.