

# WHY AGNOSTIC SIGN RESTRICTIONS ARE NOT ENOUGH: UNDERSTANDING THE DYNAMICS OF OIL MARKET VAR MODELS

---

**Lutz Kilian**  
University of Michigan

**Daniel P. Murphy**  
University of Michigan

## Abstract

Sign restrictions on the responses generated by structural vector autoregressive models have been proposed as an alternative approach to the use of exclusion restrictions on the impact multiplier matrix. In recent years such models have been increasingly used to identify demand and supply shocks in the market for crude oil. We demonstrate that sign restrictions alone are insufficient to infer the responses of the real price of oil to such shocks. Moreover, the conventional assumption that all admissible models are equally likely is routinely violated in oil market models, calling into question the use of posterior median responses to characterize the responses to structural shocks. When combining sign restrictions with additional empirically plausible bounds on the magnitude of the short-run oil supply elasticity and on the impact response of real activity, however, it is possible to reduce the set of admissible model solutions to a small number of qualitatively similar estimates. The resulting model estimates are broadly consistent with earlier results regarding the relative importance of demand and supply shocks for the real price of oil based on structural vector autoregressive (VAR) models identified by exclusion restrictions, but imply very different dynamics from the posterior median responses in VAR models based on sign restrictions only. (JEL: C68, E31, E32, Q43)

---

## 1. Introduction

There has been increasing interest in recent years in identifying the demand and supply shocks underlying the evolution of the real price of oil. Shifts in the composition of oil demand and oil supply shocks over time tend to cause temporal instabilities in conventional regressions of macroeconomic outcomes on the price of oil. They also invalidate traditional causal interpretations of responses to oil price shocks and overturn the standard view of how policy makers should respond to oil price shocks (see, e.g., Kilian (2008a) for a review).

Much of the literature on identifying oil demand and oil supply shocks has relied on structural vector autoregressive (VAR) models. The first generation of oil market VAR models building on Kilian (2009a) was based on exclusion restrictions imposed on the impact multiplier matrix. These identifying assumptions are often economically

---

*The editor in charge of this article was Fabio Canova.*

Acknowledgments: We thank Christiane Baumeister, Luca Benati, Fabio Canova, Ana María Herrera, and three referees for helpful comments.

E-mail: lkilian@umich.edu (Kilian); dpmurphy@umich.edu (Murphy)

interpretable in terms of the slopes of short-run supply and demand curves. Kilian proposed a decomposition of shocks to the real price of crude oil into three components: shocks to the supply of oil, shocks to global demand for all industrial commodities including crude oil, and demand shocks that are specific to the crude oil market. More recently, several authors have aimed to relax some of the identifying assumptions in Kilian (2009a) with the help of sign restrictions on the implied responses of the real price of oil, of global crude oil production and of global real activity to oil demand and oil supply shocks. Such models are less restrictive than Kilian (2009a) in some dimensions, but more restrictive in other dimensions. Examples include Baumeister and Peersman (2012, forthcoming), Peersman and Van Robays (2009), and Baumeister, Peersman, and Van Robays (2010). For a closely related approach see Lippi and Nobili (forthcoming). Unlike the first generation of exactly identified oil market VAR models, these structural VAR models are only partially identified—in the sense discussed in Moon and Schorfheide (2009)—and allow for a wide range of different estimates.

Our article makes four contributions to this literature. First, we demonstrate that imposing sign restrictions alone is not sufficient to resolve the question of the relative importance of oil demand and oil supply shocks. For example, the results from purely sign-identified VAR models are equally consistent with a large response of the real price of oil to oil supply shocks combined with a small response to oil demand shocks (the traditional view in the literature until 2003) and with a small response to oil supply shocks, yet a large response to oil demand shocks (the view more recently espoused in Kilian (2009a) and related work).

Second, we show that conventional approaches of aggregating the set of admissible solutions into a single summary statistic are inadequate. This criticism applies not to the use of sign restrictions for identification *per se*, but rather to the use of commonly employed summary statistics for the results of such models. In particular, we show that the practice of inferring an impulse response function from the pointwise medians of the posterior distribution of the admissible responses can be misleading. This problem occurs not only for the reasons already articulated in Fry and Pagan (2005, 2011), but also because the underlying assumption that all admissible models are equally likely, which is implicitly used in constructing the posterior distribution of the responses, is economically implausible in oil market VAR models. For example, some structural models that are admissible based on the pure sign restriction approach imply a large instantaneous jump in global oil production in response to positive oil demand shocks. That response is inconsistent with the consensus view in the literature that the short-run elasticity of oil supply is low. Including such implausible models in the construction of median responses distorts the results. We demonstrate that these distortions can be substantial in practice. In settings consistent with the consensus view, the responses that satisfy all identifying restrictions may be three times as large as the posterior median response or one-third of the magnitude of the posterior median response. In addition, the responses that satisfy all identifying restrictions often are well outside the customary 68% impulse posterior response error bands. This finding casts doubt on empirical results reported in the literature based on medians and other quantiles of the posterior distribution of VAR impulse responses obtained using sign restrictions

only. In particular, we show that the latter approach tends to overestimate the response of the real price of oil to oil supply shocks and to underestimate the response to oil-market-specific demand shocks.

Our third contribution is to show how fairly agnostic sign restrictions may be strengthened with the help of additional economically motivated inequality restrictions.<sup>1</sup> This approach allows us to avoid the use of posterior median responses altogether. First, we show that it is possible to reduce the set of admissible model solutions by combining standard sign restrictions with an empirically plausible bound on the magnitude of the oil supply elasticity on impact. This additional restriction alone suffices to assess the relative importance of oil supply and oil demand shocks for the real price of oil. We show that one can rule out large responses of the real price of oil to oil supply shocks for any reasonable bound on the oil supply elasticity. Second, we show that the relative importance of different demand shocks for the real price of oil depends on the impact response of real activity to oil-market-specific demand shocks. There is reason to believe that the impact effect of oil-market-specific demand shocks on global real activity is small. We show that imposing a bound on this impact response has little effect on the response of the real price of oil to oil supply shocks, but favors models with larger responses to oil-market-specific demand shocks at the expense of smaller oil price responses to global aggregate demand shocks.

Fourth, we show that a sign-identified model with both additional identifying restrictions imposed supports the substantive insights provided in Kilian (2009a) regarding the relative importance of different oil demand and oil supply shocks.

The new methodology we develop in this article is applicable to a range of problems in which VAR models of markets are used to identify demand and supply shocks. We explicitly rely on the market structure of the model. For example, bounds on impact price elasticities can only be imposed in models that identify exogenous shifts in demand and supply curves. There are many other applications of sign-identified VAR models in which our approach to limiting the range of admissible responses cannot be applied. In the latter case, often alternative approaches are available such as narrowing down the set of admissible responses based on robust implications of macroeconomic theory (see, e.g., Canova and De Nicolo 2002; Canova and Paustian 2010). Finally, in some situations none of these approaches will be applicable or the range of admissible models remains wide even after imposing all available identifying information. This does not mean that we should revert to using median response functions. For one thing, the concerns over the use of median responses articulated in Fry and Pagan (2011) remain valid even in this case. Unless the pointwise posterior medians of all impulse response coefficients implied by the VAR model correspond to the same structural model, which is unlikely, the median response functions will have no structural interpretation. Instead, a more informative approach will be to search for the admissible model most favorable to each of the competing economic interpretations,

---

1. The refinements of the sign restriction approach we propose in this article complement related ideas for narrowing down the set of admissible structural models in Uhlig (2005), Canova and De Nicolo (2002), and Canova and Paustian (2010).

as in Faust (1998), and to discuss the range of economic interpretations consistent with the data. In all three cases, sign-identified VAR models have the potential of revealing information in the data that may be difficult to capture by other econometric methods.

The remainder of the article is organized as follows. Section 2 describes the baseline VAR model involving only sign restrictions and motivates the modified approach of this article. Section 3 focuses on the substantive implications of our analysis for the analysis of the global market for crude oil, contrasts our findings with those in Kilian (2009a), and illustrates the bias in conventional posterior median responses. We conclude in Section 4.

**2. VAR Methodology**

*2.1 Imposing VAR Sign Restrictions*

Consider the reduced-form VAR model  $A(L)y_t = e_t$ , where  $y_t$  is the  $N$ -dimensional vector of variables,  $A(L)$  is a finite-order autoregressive lag polynomial, and  $e_t$  is the vector of white-noise reduced-form innovations with variance–covariance matrix  $\Sigma_e$ . Let  $\varepsilon_t$  denote the corresponding structural VAR model innovations. The construction of structural impulse response functions requires an estimate of the  $N \times N$  matrix  $\tilde{B}$  in  $e_t = \tilde{B}\varepsilon_t$ .<sup>2</sup> Let  $\Sigma_e = P\Lambda P'$  and  $B = P\Lambda^{0.5}$  such that  $B$  satisfies  $\Sigma_e = BB'$ . Then  $\tilde{B} = BD$  also satisfies  $\tilde{B}\tilde{B}' = \Sigma_e$  for any orthogonal  $N \times N$  matrix  $D$ . One can examine a wide range of possibilities for  $\tilde{B}$  by repeatedly drawing at random from the set  $\mathbf{D}$  of orthogonal matrices  $D$ . Following Rubio-Ramirez, Waggoner, and Zha (2010) we construct the set  $\tilde{\mathbf{B}}$  of admissible models by drawing from the set  $\mathbf{D}$  of rotation matrices and discarding candidate solutions for  $\tilde{B}$  that do not satisfy a set of a priori restrictions on the implied impulse response functions. The procedure consists of the following steps:

- (1) Draw an  $N \times N$  matrix  $K$  of  $NID(0, 1)$  random variables. Derive the  $QR$  decomposition of  $K$  such that  $K = Q \cdot R$  and  $QQ' = I_N$ .
- (2) Let  $D = Q'$ . Compute impulse responses using the orthogonalization  $\tilde{B} = BD$ . If all implied impulse response functions satisfy the identifying restrictions, retain  $D$ . Otherwise discard  $D$ .
- (3) Repeat the first two steps a large number of times, recording each  $D$  that satisfies the restrictions and storing the corresponding impulse response functions.

The resulting set  $\tilde{\mathbf{B}}$  comprises the set of admissible structural VAR models.

---

2. For a review of the construction of these structural impulse responses the reader is referred to Fry and Pagan (2005, 2011), for example.

## 2.2. *Interpreting the Set of Admissible Structural VAR Models*

A fundamental problem in interpreting VAR models identified based on sign restrictions is that there is no point estimate of the structural impulse response functions. Unlike conventional structural VAR models based on short-run restrictions, sign-identified VAR models are only set identified. They do not imply a unique solution, but a set of solutions that are all equally consistent with the identifying assumptions. Without further information, it is impossible to know which of these candidate impulse response functions depicts the true system dynamics. In fact, it is difficult merely to represent the estimation results.

Plotting all admissible response functions typically results in a plot that obscures the response patterns associated with any one structural model. It may seem that this problem could be addressed by plotting a small number of randomly selected response functions, but this alternative approach entails the risk of overlooking response functions with different qualitative patterns than those selected at random. Another attempt to resolve this problem is the common practice of reporting the posterior median response at each horizon for all responses in the identified set.<sup>3</sup> Applied researchers often treat the vector of pointwise medians as though it were a point estimate of the response function (even if they do not refer to it as a point estimate) and use it as the basis for forecast error variance decompositions. As pointed out by Fry and Pagan (2005, 2011), this practice is misleading because the response functions constructed from the posterior median responses for each horizon typically is a composite of different structural response functions. There is no reason to believe that any one structural model among the set of admissible models will generate the median response at each horizon for all response functions. In addition, as we demonstrate in Section 3 of this article, the premise underlying the construction of median responses that all admissible models are equally likely is typically violated. This fact would invalidate median responses even if they corresponded to responses from the same structural model.

In this article we illustrate in the context of oil market VAR models how these problems of interpretation can be minimized, if not avoided altogether. Since the introduction of VAR models based on sign restrictions, several researchers have made proposals to facilitate the interpretation of a set of admissible structural impulse response functions. Broadly speaking, there are two approaches. One approach has been to narrow down the set of admissible responses by imposing additional restrictions. For example, Canova and De Nicolo (2002) and Canova and Paustian (2010) propose to reduce the number of admissible solutions by imposing additional structure in the form of sign restrictions on dynamic cross-correlations. They motivate

---

3. See, for example, Baumeister and Peersman (forthcoming), Lippi and Nobili (forthcoming), and Baumeister, Peersman, and Van Robays (2010) in the context of oil market models and Peersman (2005), Dedola and Neri (2007), Canova and Paustian (2010), and Scholl and Uhlig (2008) in other contexts. An alternative approach is to report probabilities of events based on the fraction of candidate impulse responses that yield, for example, a positive response of a non-restricted variable to a given shock. That approach suffers from the same conceptual limitations as the median response, however.

these restrictions based on properties of DSGE models. In related work, Uhlig (2005, p. 402) discusses how to strengthen sign restrictions based on imposing zero restrictions on selected impact responses. An alternative approach involves the use of a penalty function to narrow down the set of admissible models to a singleton. The specification of the penalty function also utilizes additional identifying information. For example, Francis, Owyang, and Roush (2007) identify a technology shock as that shock which satisfies sign restrictions and maximizes the forecast-error variance (MFEV) share in labor productivity at a finite horizon. Faust (1998) appeals to an analogous argument regarding the effects of monetary policy shocks on real output.

None of these techniques is directly applicable to oil market VAR models. For example, it is not possible to implement the approach of Canova and De Nicolo (2002). While there have been significant improvements in the design of DSGE models of the relationship between the real price of oil and the domestic economy in recent years, to date these models have focused on a subset of the relevant oil demand and oil supply shocks in the interest of keeping the model tractable (see, e.g., Bodenstein, Erceg, and Guerrieri 2011; Nakov and Pescatori 2010). This makes it inadvisable to use the sign of conditional correlations implied by these models in narrowing down the set of admissible structural VAR models. Moreover, as Bodenstein et al. (2011) and Lippi and Nobili (forthcoming) demonstrate, the implications of DSGE models will be sensitive to the assumed elasticities of substitution. Since we have little knowledge of the magnitude of these elasticities, it seems prudent not to impose such sign restrictions. Likewise, the minimum forecast error variance approach does not easily generalize to the analysis of oil markets because there is no natural choice for the penalty function.

Instead, in this article, we employ a new method of identifying structural VAR models building on the sign restriction approach. Although this approach is distinct from earlier approaches, it shares with the work of Canova and De Nicolo (2002) and Canova and Paustian (2010) the feature that we successively narrow down the range of admissible structural models by imposing additional economically plausible identifying assumptions. Our restrictions are soft in that they involve bounds rather than strict equality or exclusion restrictions. Starting with the usual sign restrictions, we proceed in three steps. We first impose additional identifying assumptions in the form of bounds on the short-run oil supply elasticity. It is widely accepted that this elasticity is low on impact (see, e.g., Hamilton 2009a,b; Kilian 2009a). We propose a plausible upper bound for that elasticity based on historical data. This bound helps us greatly reduce the number of admissible solutions and allows us to resolve some of the ambiguities in the qualitative results implied by purely sign-identified VAR models. While the specific bound we derive is suggestive only, we show that our results are robust for any reasonable upper bound on the oil supply elasticity.

In the second step, we impose economically plausible bounds on the impact response of global real activity to oil-market-specific demand shocks. As discussed in Kilian (2009a,b), Alquist and Kilian (2010), and Kilian and Murphy (2010), these shocks primarily reflect revisions to expectations about future demand and supply conditions. Such shocks affect current real activity only indirectly through their contemporaneous effect on the real price of oil. That effect is likely to be small within the month, motivating a bound on the impact response. Rather than pinning

down this bound with any precision, we illustrate the sensitivity of the results to the magnitude of this bound. The latter is shown to affect only the relative importance of different types of oil demand shocks for the real price of oil, but not the overall importance of oil supply shocks.

Our approach to identification suffices to narrow down the set of structural models to a handful of models with essentially identical dynamics, allowing us to choose one of these models as the final model without loss of generality. Like existing approaches, our approach is not intended as a general solution (and indeed may not be feasible in all applications). It is specifically designed to shed light on the identification of oil demand and oil supply shocks. Nevertheless, the methodological insights provided by this example will apply to other empirical studies as well, as discussed in the conclusion.

### 3. Modeling the Global Crude Oil Market

Building on Kilian (2009a), we consider a fully structural oil market VAR model of the form

$$B_0 y_t = \alpha + \sum_{i=1}^{24} B_i y_{t-i} + \varepsilon_t, \quad (1)$$

where  $\varepsilon_t$  is a vector of orthogonal structural innovations and  $y_t$  consists of the percent change in global crude oil production, an index of real economic activity representing the global business cycle, and the log of the real price of oil from 1973.1–2008.9.<sup>4</sup> The vector  $\varepsilon_t$  consists of a shock to the world production of crude oil (“oil supply shock”), a shock to the demand for crude oil and other industrial commodities associated with the global business cycle (“aggregate demand shock”), and a shock to demand for oil that is specific to the oil market (“oil-market-specific demand shock”). The latter shock is designed to capture innovations to the demand for crude oil that are orthogonal to aggregate demand shocks such as speculative oil demand shocks.

#### 3.1. VAR Identification by Exclusion Restrictions

Kilian (2009a) imposes identifying restrictions on the slope of the short-run oil supply curve conditional on past data. Notably, Kilian postulates that the short-run oil supply

---

4. The price of crude oil is based on US refiners’ acquisition cost for imported crude oil, as reported in the *Monthly Energy Review* of the Energy Information Administration. The series has been extrapolated backwards as in Barsky and Kilian (2002) and deflated by the US consumer price index. Unlike Kilian (2009a) we do not annualize the growth rate of oil production. This is merely a scaling convention and does not affect the substantive results. The global oil production data from the same source are measured in millions of barrels of oil and have been expressed in percent changes. The index of real activity has been obtained by cumulating average rates of increase in dry cargo ocean shipping freight rates, deflating the nominal index by the US CPI and linearly detrending the resulting series. The index is expressed in percent deviations from trend. For a full discussion of the data sources and construction of the data see Kilian (2009a).

TABLE 1. Sign restrictions on impact responses in VAR model.

	Oil supply disruption	Aggregate demand shock	Oil-specific demand shock
Oil production	–	+	+
Real activity	–	+	–
Real price of oil	+	+	+

curve is vertical, implying that global oil production does not respond to oil demand shocks instantaneously, but only with a delay of a month. This assumption can be motivated based on costs to adjusting production and is consistent with anecdotal evidence on OPEC production decisions. Note that the model permits the presence of exogenous oil supply shocks in the form of shifts of the short-run oil supply curve. In addition, the model imposes a delay restriction on real activity. While allowing for unrestricted lagged feedback from fluctuations in the real price of oil to global industrial commodity markets, Kilian rules out instantaneous feedback within the month. That restriction is motivated by the fact that there is no evidence of instantaneous feedback from fluctuations in the real price of oil to the dry cargo ocean freight rates that form the basis of Kilian's measure of global real activity. These exclusion restrictions jointly imply a recursive structural model of the form

$$e_t \equiv \begin{pmatrix} e_t^{\Delta prod} \\ e_t^{rea} \\ e_t^{rpo} \end{pmatrix} = \begin{bmatrix} a_{11} & 0 & 0 \\ a_{21} & a_{22} & 0 \\ a_{31} & a_{32} & a_{33} \end{bmatrix} \begin{pmatrix} \varepsilon_t^{oil\ supply\ shock} \\ \varepsilon_t^{aggregate\ demand\ shock} \\ \varepsilon_t^{oil-market-specific\ demand\ shock} \end{pmatrix}.$$

Although economically plausible, this model is not without limitations. For example, the assumption of a short-run oil supply elasticity of zero may be a good approximation, but is unlikely to be literally correct.<sup>5</sup> This fact has prompted Baumeister and Peersman (forthcoming), for example, to explore the use of sign restrictions that do not require the researcher to take a stand on the short-run price elasticity of oil supply. Below we examine the robustness of the results in Kilian (2009a) to such alternative methods of identification. Sign restrictions on the VAR response functions arise naturally in the context of structural models of the oil market.

### 3.2. VAR Identification Based on Sign Restrictions Only

Table 1 shows our baseline sign restrictions. The sign restrictions we impose on the impact responses are fully consistent with the restrictions used in Baumeister and Peersman (forthcoming) and related work. We postulate that a positive aggregate

5. In a slightly different context, Davis and Kilian (2011) and Kilian (2010) show how this type of assumption may be relaxed, building on work by Abraham and Haltiwanger (1995), if we are willing to take a stand on the value of the elasticity. As observed by Canova and Paustian (2010), however, exact magnitude restrictions are unlikely to hold in the data.



demand shock will tend to raise oil production, stimulate real activity and increase the real price of oil on impact. A positive oil-market-specific demand shock on impact will raise the real price of oil and stimulate oil production, but lower real activity. An unexpected oil supply disruption will by construction lower oil production on impact. It also will lower real activity, while increasing the real price of oil. The set of restrictions imposed in Table 1 implies a unique response pattern for each structural shock.

Unlike Baumeister and Peersman (forthcoming) we do not impose sign restrictions on any of the response functions beyond the impact period. The reason is that we are more agnostic about some of the dynamic responses of crude oil production and global real activity and do not wish to rule out that general equilibrium effects may cause a sign reversal. Nevertheless, we find that our estimates even in the absence of further restrictions usually satisfy the additional sign restrictions imposed by Baumeister and Peersman (forthcoming), and, if not, tend to be close to zero.<sup>6</sup>

To implement the baseline specification we start with 1.5 million random draws of the rotation matrix. Of these, 30,860 draws satisfy the criteria in Table 1. The implied response functions are displayed in Figure 1(a). For expository purposes we focus on the responses of the real price of oil. In this and the subsequent figures, all responses have been normalized such that an innovation will tend to raise the price of oil. Figure 1(a) illustrates that sign restrictions are consistent with a range of response functions of very different amplitude and shape.

It may seem that we could summarize these response functions by reporting median responses. As Figure 2 illustrates, this approach would be misleading because there is reason to believe that not all admissible structural models are equally likely. Figure 2 shows selected response functions for two of the structural models underlying Figure 1. For each structural model, we show the responses of oil production (left column) and of the real price of oil (right column) to each of the three structural shocks. Figure 2 illustrates that imposing sign restrictions alone is not sufficient to resolve the question of the relative importance of different oil demand and oil supply shocks. Model 1 implies a large response of the real price of oil to oil supply shocks and a small response to oil-market-specific demand shocks, consistent with the traditional view in the literature until 2003. In contrast, Model 2 implies a moderate price response to oil supply shocks, yet a large response to oil-market-specific demand shocks, consistent with the view more recently espoused in Kilian (2009a) and related work. Both models are equally consistent with the sign restrictions we imposed, but mutually contradictory.

The common practice of reporting median results as a way of aggregating these opposing views relies on the premise that the two structural models shown in Figure 2 are equally likely. That presumption is not supported upon closer examination. Model 1 implies that a positive oil-market-specific demand shock is associated with a large jump in global crude oil production within the month. This jump strains credulity, all the more so, as the same shock only implies a small increase in the real price of oil. One can cast this result in terms of the implied short-run price elasticity of oil supply.

---

6. Also note that our reduced-form VAR specification follows the original analysis in Kilian (2009a) rather than that in Baumeister and Peersman (forthcoming).

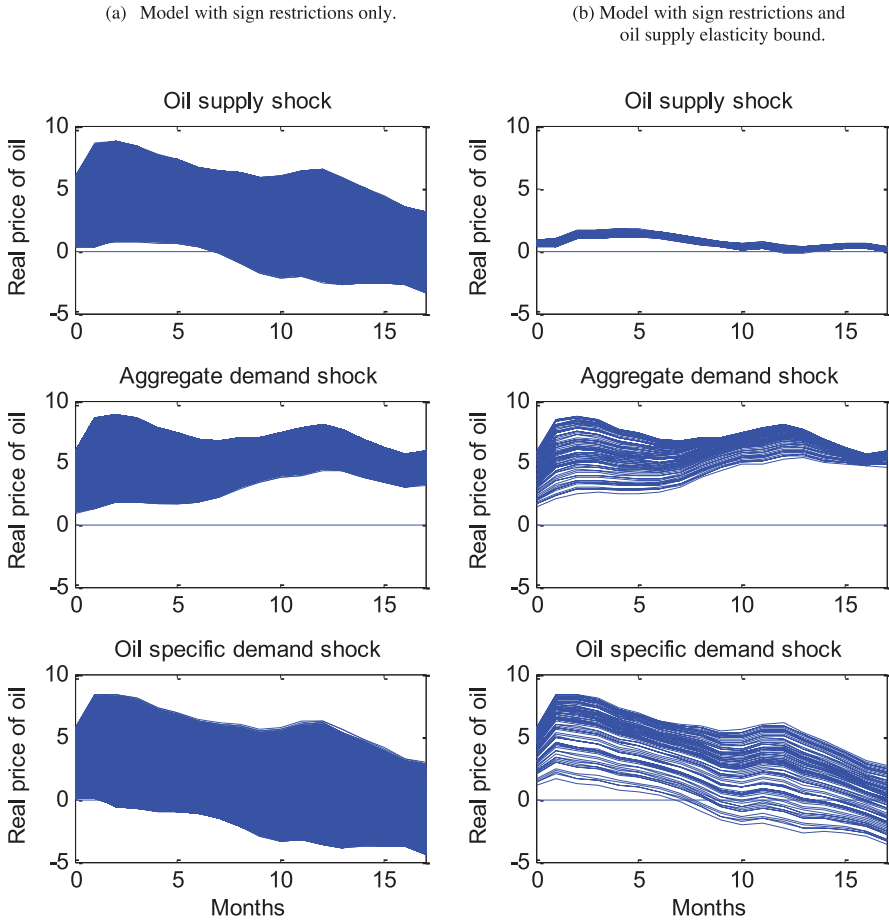


FIGURE 1. Set-identified responses to one-standard-deviation structural shocks. (a) Model with sign restrictions only. (b) Model with sign restrictions and oil supply elasticity bound.

Recall that  $a_{13}$  is the impact response of global oil production to an oil-market-specific demand shock and  $a_{33}$  is the impact response of the real price of oil to an oil-market-specific demand shock. Because our variables are expressed in logs, we can interpret  $a_{13}/a_{33}$  as the impact elasticity of the supply of oil with respect to the real price of oil. That elasticity based on Model 1 is 1.89. Given the prevailing view in the literature that the short-run oil supply elasticity is very low and perhaps zero, an implied elasticity of 1.89 suggests that Model 1 is not credible and that the traditional view in the literature until 2003 regarding the effects of oil supply shocks on the real price of oil cannot be right.

In contrast, Model 2 implies a jump on impact in the real price of oil in response to an oil-market-specific demand shock with little adjustment in the level of crude oil production. The implied oil supply elasticity of 0.01 is much more reasonable. It also is much more in line with the oil supply elasticity implied by the other demand shock

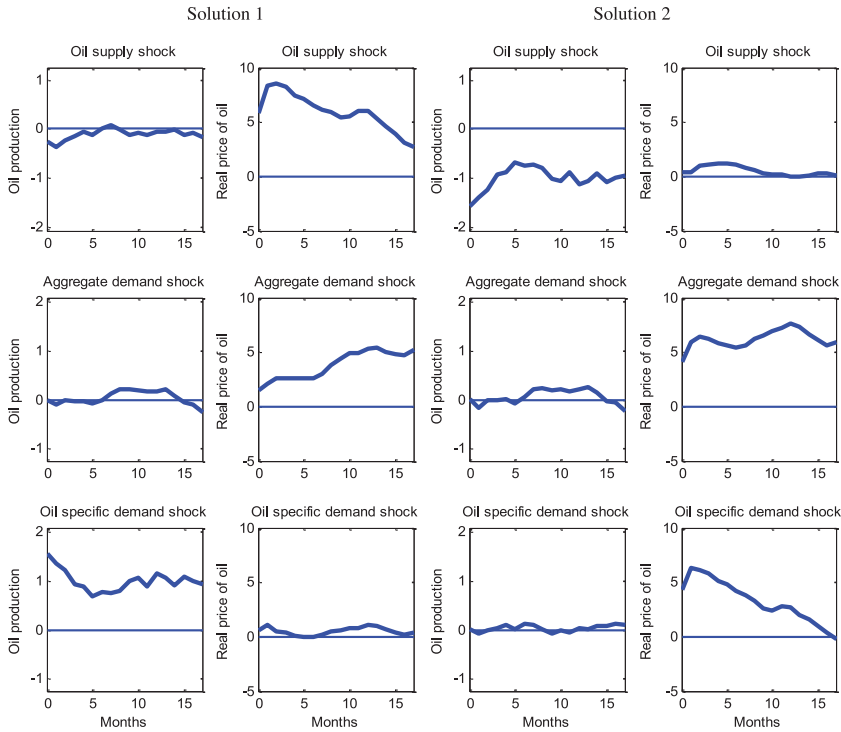


FIGURE 2. Alternative responses to one-standard deviation structural shocks.

than in Model 1. Thus, we must abandon the premise that all admissible models are equally likely. We conclude that, in this example, rather than reporting the median across Models 1 and 2, we should discard Model 1 and focus on Model 2 only. More generally, this line of reasoning suggests that we should be able to narrow down the range of admissible structural models by imposing additional constraints on the magnitude of the oil supply elasticity, as discussed in Section 3.3.

### 3.3. VAR Identification Based on Sign Restrictions Combined with Elasticity Bounds

The consensus in the literature is that the short-run price elasticity of oil supply is close to zero. For example, Hamilton (2009b, p. 25) observes that “in the absence of significant excess production capacity, the short-run price elasticity of oil supply is very low.” In practice, it often will take years for significant production increases.<sup>7</sup>

7. In related work, Kilian (2009a) makes the case that—even in the presence of spare capacity—the response of oil production to price signals will be negligible within the month because changing oil production is costly. As a result, oil producers tend to set production targets based on expected demand growth. Oil producers will respond to unanticipated oil price increases only if that increase is expected to persist. Detecting persistent shifts in expected demand growth is akin to the problem of detecting trend

This motivates the imposition of an upper bound on the ratios  $a_{13}/a_{33}$  and  $a_{12}/a_{32}$ . Our starting point is an upper bound of 0.0258. This bound can be motivated based on historical experience, as discussed in the Online Appendix, or may be viewed as merely a working hypothesis for expositional purposes.

Figure 1(b) shows the impulse responses functions generated by the admissible models after imposing this inequality constraint in addition to the sign restrictions of Table 1. The additional identifying assumption greatly reduces the number of admissible models from 30,860 to 80, given 1.5 million initial draws. Even without additional information, these results are highly informative. Figure 1b shows that all structural models that imply large price responses to oil supply shocks have been eliminated in favor of models that imply large price responses to oil-market-specific demand shocks.<sup>8</sup>

An obvious concern in developing additional identifying restrictions is that these restrictions must not involve circular reasoning. The key controversy we are interested in resolving in this article is over the magnitude of response of the real price of oil to oil supply shocks. We want the data to tell us what that magnitude is. Imposing an upper or a lower bound on the magnitude of that response, therefore, would simply amount to imposing the answer, defeating the purpose of the exercise. It is important to stress that our bound on the impact price elasticity of oil supply does not involve such circular reasoning. First, this ratio bound relates to responses to oil demand shocks rather than responses to oil supply shocks. Thus, we obviously have not bounded any of the responses to oil supply shocks. Second, by construction a bound on ratios of impact responses does not bound the level of the impact responses (or for that matter of the responses at higher horizons). Thus, we leave the magnitude of all responses in the VAR including the price response to oil supply shocks to be determined by the data, protecting ourselves from circular reasoning.

### *3.4. How Robust are the Results to the Magnitude of the Elasticity Bound?*

Although the results in Figure 1(b) are instructive, our numerical bound on the oil supply elasticity based on historical evidence is suggestive only. A useful robustness check is to report results for oil supply elasticity bounds of twice and three times the bound in the baseline model. Such elasticity values are still consistent with the consensus in the literature that the short-run oil supply curve is steep without imposing that the curve is near vertical. It can be shown that the peak effect of an oil supply shock on the real price of oil, as we vary the supply elasticity bound, increases from 1.73 in the baseline model to 2.5, as we double the elasticity, and to 3.19, as we triple

---

breaks in economic data. It requires a sufficient time span of data, making it inadvisable for oil producers to respond to month-to-month variation in fundamentals. Indeed, OPEC production decisions have been remarkably sluggish.

8. The elasticity bound we impose may equivalently be interpreted from a Bayesian point of view as a truncated uniform prior on the value of the impact oil supply elasticity. When prior economic information is less diffuse, structural models may instead be weighted according to other prior distributions (see, e.g., Gambetti, Pappa, and Canova 2008).

TABLE 2. Oil supply elasticity bound required for oil supply shocks to explain a given percentage of the forecast error variance of the real price of oil.

Horizon	10%	20%	30%	40%
6	0.07	0.11	0.16	0.19
12	0.08	0.12	0.17	0.22
18	0.09	0.14	0.19	0.26

Note: Based on the model with sign restrictions only.

it. Thus, even with three times the original elasticity bound, the upper bound on the price response is much lower than the highest peak response of 8.71 obtained without any elasticity restrictions (see Figure 1(a)). Interestingly, changing the elasticity in this manner has little effect on the other oil price responses, except that the set of admissible models becomes denser. This additional evidence shows that our results regarding the relative importance of oil supply shocks and oil demand shocks for the real price of oil are not sensitive to reasonable changes in the oil supply elasticity bound and do not hinge on the specific bound we imposed in the baseline model.

A different way of posing the same question is to ask what we must believe about the slope of the oil supply curve (as implied by the oil supply elasticity bound) for the oil supply shocks to explain at least a certain percentage of the variability of the real price of oil. This alternative exercise is in the spirit of the analysis in Faust (1998). Table 2 shows that for oil supply shocks to explain 10% of the variability of the real price of oil would require an oil supply elasticity between 0.07 and 0.09, depending on the horizon. Even a still moderate share of 30% would require an oil supply elasticity between 0.16 and 0.19. These results demonstrate that oil supply shocks cannot have large effects on the real price of oil, given the consensus view that the oil supply curve is very steep. Moreover, the conventional view in the literature that the oil supply elasticity is near zero is incompatible with the equally popular view that oil supply shocks have large effects on the real price of oil.

### ***3.5. Identification based on Sign Restrictions, the Elasticity Bound and a Bound on the Response of Real Activity to Oil-Market-Specific Demand Shocks***

The preceding analysis suffices to answer the central question of what the relative role of oil supply and oil demand shocks is in determining the real price of oil, but leaves considerable uncertainty about the relative importance of different oil demand shocks for the real price of oil. Although we cannot resolve this question with the same degree of confidence as the question of how quantitatively important oil supply shocks are, we can bring some additional identifying information to bear. It can be shown that structural models in Figure 1(b) with large negative  $a_{23}$  values are associated with larger price responses to aggregate demand shocks and smaller price responses to oil-market-specific demand shocks (while implying similar responses to oil supply shocks). Although there is no reason to expect  $a_{23}$  to be literally zero, as postulated

TABLE 3. Percentage of the forecast error variance of the real price of oil explained by each shock as a function of alternative lower bounds on  $a_{23}$ .

Horizon:	Oil supply shock			Aggregate demand shock			Oil-specific demand shock		
	6	12	18	6	12	18	6	12	18
-0.5	2.9	2.4	1.8	10.5	21.3	32.3	86.6	76.4	66.0
-1	2.2	1.8	1.3	18.7	30.8	42.2	79.1	67.4	56.5
-1.5	2.8	2.2	1.7	28.0	40.7	51.9	69.3	57.1	46.5
-2	2.9	2.4	1.7	39.3	52.1	62.4	57.8	45.6	35.9
-2.5	2.0	1.6	1.2	53.0	64.8	73.4	45.0	33.6	25.4
-3	2.1	1.7	1.3	65.8	75.9	82.2	32.1	22.4	16.5

Note: Based on the model with sign restrictions and the baseline supply elasticity bound.

in Kilian (2009a), we would expect the value of  $a_{23}$  to be at least close to zero. The reason is that oil-market-specific demand shocks tend to be associated with revisions in expectations about future oil supply and oil demand conditions. Such shocks affect real activity within the same month only through their effect on the real price of oil. This reasoning casts doubt on models that imply large negative impact responses of real activity to oil-market-specific demand shocks.<sup>9</sup>

Table 3 shows the percentage of the forecast error variance of the real price of oil explained by each of the three shocks as a function of hypothetical lower bounds on the value of  $a_{23}$ . The relative importance of oil supply shocks is essentially unaffected by the bound on  $a_{23}$ . At the same time, the relative importance of different oil demand shocks appears remarkably robust to the bound on  $a_{23}$ . Even for bounds on  $a_{23}$  as low as  $-1.5$  more than half of the forecast error variance of the real price of oil at horizons of one year is explained by oil-market-specific demand shocks. Only for even wider bounds on  $a_{23}$  aggregate demand shocks become quantitatively more important than oil-market-specific demand shocks.

If we are willing to discard as economically implausible all models in Figure 1(b) that do not satisfy  $-1.5 < a_{23} < 0$ , consistent with the notion that  $a_{23}$  is near zero, but not equal to zero, we are left with 17 admissible models that are nearly equivalent in terms of their implied response functions. In Section 3.6, we focus on one of these models and ask how sensitive the impulse response estimates and historical decompositions in Kilian (2009a) are to the use of this alternative set of identifying assumptions. For expository purposes we choose the model with the largest response of the real price to oil supply shocks in Figure 1(b).

9. In support of this exclusion restriction on  $a_{23}$ , Kilian (2009a) provided evidence that the dry cargo ocean shipping freight rates, on which the global real activity index is based, do not covary with the real price of oil at high frequency. This result is plausible because the bunker fuel used to run ships is a residual product in the refining process and not as sensitive to fluctuations in the real price of oil as the prices of gasoline and other higher-level refined products.

TABLE 4. Alternative estimates of  $B_0^{-1}$ .

Recursively Identified Structural Model <sup>a</sup>			Sign Restrictions Combined with All Additional Restrictions		
1.572	0	0	1.567	0.023	0.128
-0.007	4.357	0	0.007	4.252	-0.947
-0.349	0.955	5.924	-0.835	2.211	5.526

Note: The first column of the identification matrix has signs opposite of the impact responses displayed in the graphs because those graphs depict the effects of an unanticipated oil supply disruption, following the convention adopted in Kilian (2009a).

<sup>a</sup> Based on Kilian (2009a).

### 3.6. Comparison with the Recursively Identified Model

It can be shown that the response functions implied by this model, with few exceptions, are within the confidence bands computed for the recursively identified structural model of Kilian (2009a). One exception is that the response of the real price of oil to the aggregate demand shock for the first few months is slightly higher than the upper bound of the confidence band. The other exception is that the response of real activity to the oil-market-specific demand shock for the first few months lies below the lower bound of the confidence band. These differences in the impulse response results can be traced to differences in the estimates of the impact multiplier matrix  $B_0^{-1}$ . Table 4 contrasts the estimate of that matrix in Kilian (2009a) with our estimate. Three results stand out. First, although neither approach to identification nests the other, most unrestricted parameter values are broadly similar in the two models. One exception is the larger value of  $a_{32}$  which implies a larger impact response of the real price of oil to aggregate demand shocks. Second, there is little evidence against the assumption of a vertical supply curve embodied in the two exclusion restrictions of the first row. Our alternative model implies values of  $a_{12}$  and  $a_{13}$  that are quite close to zero. In fact, the implied impact oil supply elasticities are about 0.01 and 0.02, respectively. Third, a somewhat larger discrepancy arises for  $a_{23}$ . The estimated value of  $-0.947$  is clearly different from zero, which explains the larger impact response of real activity to oil-market specific demand shocks in the nonrecursive model.

Notwithstanding these differences, the recursive model and the nonrecursive model provide broadly similar accounts of the evolution of the real price of oil. Historical decompositions measure the extent to which the evolution of the real price of oil can be traced to each of the three demand and supply shocks. Figure 3 superimposes the historical decomposition implied by the nonrecursive structural model developed in this article on that implied by the recursively identified model in Kilian (2009a). The relative contribution of the oil supply shock is invariant to the method of identification. Either way, the explanatory power of oil supply shocks is low. The nonrecursive structural model implies somewhat larger explanatory power for aggregate demand shocks during major oil price increases and somewhat lower explanatory power during

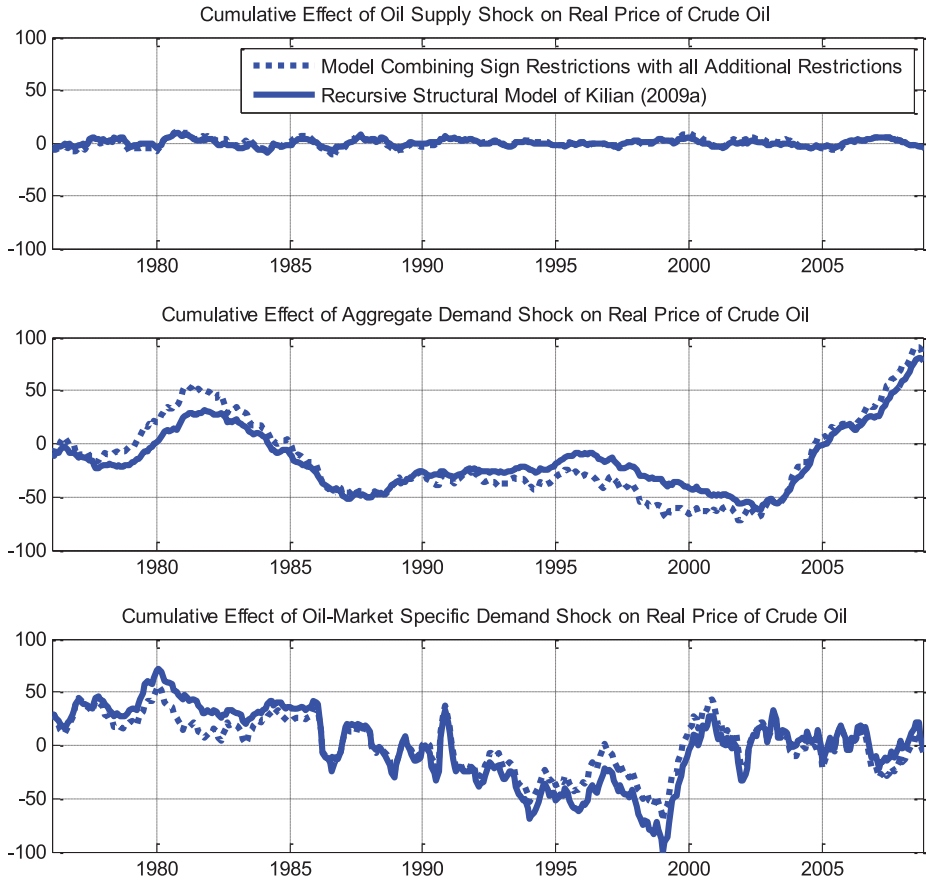


FIGURE 3. Historical decompositions of the real price of oil: 1976.1–2008.9: Model combining sign restrictions with all additional restrictions imposed.

other periods. The added explanatory power comes at the expense of that of the oil-market-specific demand shock in the last panel.

Despite this difference in emphasis, both models imply that a substantial part of the 1979/80 oil price shock can be attributed to shocks to the demand for all industrial commodities including crude oil associated with the global business cycle, with the remainder being largely attributed to oil-market-specific demand shocks. Both models suggest that virtually all the surge in the real price of oil after 2002 can be attributed to global aggregate demand shocks. Finally, both models imply that the oil price shock of 1990/91 was due almost entirely to oil-market-specific demand shocks, and both models imply that oil supply shocks never played a dominant role in driving the real price of oil. We conclude that the results in Kilian (2009a) are remarkably robust to relaxing all exclusion restrictions.<sup>10</sup>

10. This result is in striking contrast to the finding of Canova and De Nicolo (2002) who concluded that the exclusion restrictions imposed by standard semi-structural VAR models of monetary policy are



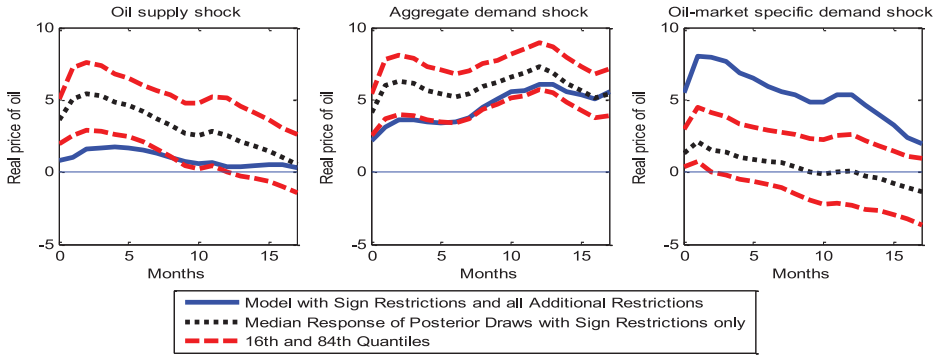


FIGURE 4. Responses of the real price of oil to one-standard deviation structural shocks: Quantile responses from posterior of model with sign restrictions only and response from model with sign restrictions and all additional restrictions imposed. Notes: The median response refers to the median of the posterior distribution of impulse responses at each horizon obtained for the model with sign restrictions only. The results are based on 200 draws from the posterior of the reduced form parameters with 20,000 rotations each. The response functions for the model combining sign restrictions with the additional identifying restrictions proposed in this article are based on the same model as in Figure 3.

### 3.7. Comparing the Responses in the Fully Specified VAR Model to the Median Posterior Response for the Model Based on Sign Restrictions Only

Next we illustrate the extent to which our answer would have differed, had we inferred the system dynamics from the median posterior responses of the model based on sign restrictions only, as is common practice in the literature. Figure 4 illustrates that the posterior median response may grossly overstate or understate the magnitude of the actual response.<sup>11</sup> For example, the response of the real price of oil to an oil-market-specific demand shock in the model that imposes all identifying restrictions can easily be three times as large as the posterior median response, whereas the corresponding response to an oil supply shock may be one-third of the magnitude of the posterior median response. The discrepancy in results of course reflects the inclusion of economically implausible models in the construction of the posterior medians. This example illustrates that the posterior median response cannot credibly serve as a robustness check to recursively identified models; nor is it likely to be a credible representation of the mapping from reduced form errors to structural shocks.<sup>12</sup>

not consistent with estimates from sign-restricted VAR models. Our findings illustrate that it is possible for alternative identifying restrictions to yield similar results. A likely explanation of this result is that the correlation between the reduced-form errors in our VAR model is very low, indicating that the variance-covariance matrix of the regression errors is nearly diagonal. This was not the case in Canova and De Nicolò's application.

11. These results are based on 200 draws from the posterior distribution of the reduced form parameters. For each posterior draw, we draw 20,000 rotations of the impact multiplier matrix.

12. The same issues would arise if we focused on the response function corresponding to the median response on impact.

Allowing for uncertainty in the form of pointwise 68% posterior error bands does little to mitigate this problem. Figure 4 shows that the response to oil-market-specific demand shocks obtained under our identifying assumptions is well outside the posterior error bands at all horizons. The other two response functions are outside the pointwise posterior error bands at least at some horizons.<sup>13</sup>

### 3.8. Inference on the Impulse Responses

Pointwise 68% and 95% posterior error bands for the impulse response estimates may be constructed following the procedure outlined in Uhlig (2005) with suitable changes in the algorithm to reflect the additional identifying restrictions introduced in this article. These bands indicate regions of high posterior probability and are not specific to the model we selected earlier for expository purposes.<sup>14</sup> Figure 5 shows that the responses of oil production are rather precisely estimated compared with the responses of global real activity and of the real price of oil. The representative model whose responses we singled out for expository purposes earlier typically lies near the center of the posterior distribution and is always contained within the 68% error band. The posterior distribution assigns little probability mass to large responses of the real price of oil to oil supply shocks, yet assigns considerable probability mass to large responses to oil demand shocks, suggesting that our main substantive conclusions about the relative importance of oil demand and oil supply shocks are robust.

## 4. Conclusions

We presented an improved approach to the identification for VAR models based on sign restrictions. Our methodology is specifically designed for VAR models of markets and was illustrated in the context of oil market VAR models. Our identifying restrictions are substantially more agnostic about some aspects of the structural impact responses than the model of Kilian (2009a) and dispense with all exclusion restrictions, yet they allow identification of the structural model on the basis of other economic information not commonly imposed in the construction of the response functions.

At the same time, our approach is considerably less agnostic than oil-market models identified based on sign restrictions only in that we use additional identifying

---

13. In related work, Fry and Pagan (2005, 2011) made the suggestion that we focus on the structural response function that minimizes the distance to the median responses at each horizon and for all impulse response functions. Figure 4 illustrates that their proposal will not work in general. While their proposal resolves the problem that each coefficient of the response function potentially comes from a different structural model, it does not address the problem that the response functions underlying the median often are economically not equally plausible. Thus, the resulting estimates are not economically meaningful in the absence of further identifying information.

14. The analysis is based on 200 draws from the posterior of the reduced-form parameters with 200,000 rotations each. It is important to keep in mind that these error bands do not correspond to classical confidence bands even asymptotically. As Moon et al. (2009) have recently shown, classical impulse response confidence bands for sign-identified VAR models tend to be considerably wider than posterior error bands.

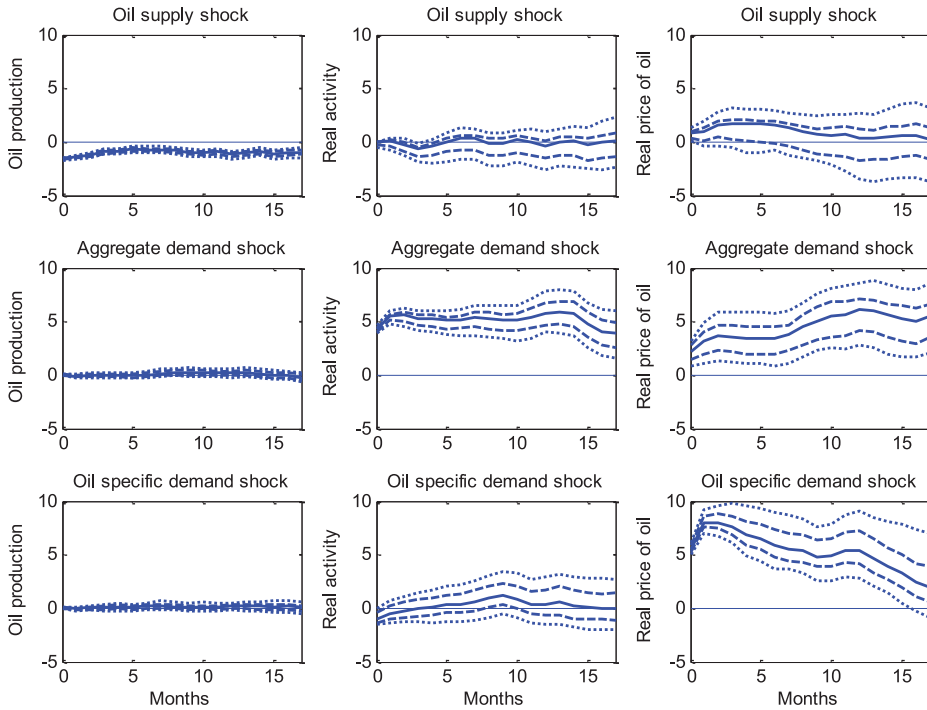


FIGURE 5. Responses to one-standard deviation structural shocks; model combining sign restrictions with all additional restrictions imposed representative response estimate with 68% and 95% posterior error bands. Notes: Error bands constructed as outlined in Uhlig (2005) with suitable changes to reflect the additional identifying restrictions we proposed in this article. Results based on 200 draws from the posterior distribution of the reduced-form parameters with 200,000 rotations each.

information. We view this fact as an advantage rather than a drawback. Our results support the observation in Fry and Pagan (2005) that remaining agnostic is not without cost. First, as we demonstrated, it can make it impossible to answer the questions of economic interest such as what the relative importance of oil demand and oil supply shocks is for the evolution of the real price of oil. Second, we showed that attempting to answer those questions anyway on the basis of insufficient information, for example by focusing on posterior median responses under the presumption that all admissible models are equally likely, can lead researchers seriously astray. Clearly, the extent of the biases in the median responses that we documented hinges on the credibility of the conventional view in the literature that the impact price elasticity of oil supply is near zero. Even for someone skeptical of this conventional view, however, our example illustrates how sensitive impulse response results from sign-identified VAR models can be to the introduction of additional identifying information.

Our analysis highlights the importance of imposing additional identifying information when such information exists. While ours is not the first study to observe that imposing additional restrictions could be useful in principle, unlike earlier studies such as Canova and Paustian (2010), we explicitly demonstrated that imposing such

restrictions actually changes the empirical results fundamentally. We showed that the response estimates obtained after imposing additional identifying restrictions can be qualitatively and quantitatively different from the median responses implied by more agnostic models. Finally, whereas Canova and De Nicolo (2002) concluded that the exclusion restrictions imposed by standard semi-structural VAR models of monetary policy are not consistent with estimates from sign-restricted VAR models, we found that in the context of oil market VAR models not much is lost by imposing the exclusion restrictions of Kilian (2009a) on the impact multiplier matrix, as long as we are interested in the response of the real price of oil only. A robust finding within the class of models we considered is that the fluctuations in the real price of oil are mainly driven by oil demand shocks with oil supply shocks playing a minor role, consistent with the substantive conclusions in Kilian (2009a). In particular, the increase in the real price of oil since 2003 is well explained by the cumulative effects of shocks to the global aggregate demand for all industrial commodities. In contrast, the oil price surge in 1979/80 reflected a mix of oil-market-specific demand shocks and global aggregate demand shocks and that of 1990/91 primarily oil-market specific demand shocks. In contrast, oil market VAR models based on sign restrictions alone have a tendency to overstate the response of the real price of oil to oil supply shocks and to understate its response to oil-market-specific shocks. This finding is important for the ongoing debate in the literature about the relative importance of oil demand and oil supply shocks for the real price of oil (see, e.g., Hamilton 2009a).

Although our discussion focused on the market for crude oil, our results highlight that more generally great care is required in the interpretation of the estimates of VAR models estimated using sign restrictions only. Our approach in this paper was designed to enhance the usefulness of sign restrictions in the VAR analysis of oil markets. With the help of additional inequality restrictions we studied the range of responses consistent with alternative identifying assumptions. This allowed us to explore how far certain bounds can be relaxed without affecting the substantive conclusions. For example, we showed that the price response to an oil supply shock in sign-identified VAR models is small for any reasonable bound on the short-run oil supply elasticity, even if we relax all other identifying restrictions. This means that, conditional on the sign restrictions, the conventional view in the literature that the oil supply elasticity is near zero is incompatible with the equally popular view that oil supply shocks have large effects on the real price of oil. In contrast, the relative importance of different oil demand shocks requires the researcher to take a stand on other features of the structural model. By illustrating these trade-offs we provided a framework that helps researchers to understand the role of alternative identifying restrictions and to convey the sensitivity of their results to alternative identifying assumptions to the reader.

The tools we developed to address these potential problems of interpretation also are applicable to the study of other markets. A case in point is Abraham and Haltiwanger's (1995) analysis of labor market responses to demand and supply shocks. Their analysis—like ours—focuses on the magnitude of demand and supply elasticities and lends itself to conducting a sensitivity analysis along the lines proposed in our paper. Macroeconomic studies of the effects of demand and supply shocks at the

aggregate level are another example. Rather than impose exclusion restrictions that reflect extreme assumptions about the slopes of aggregate demand or aggregate supply curves, our approach could be used to assess the sensitivity of the results based on elasticity bounds. It is important to stress that our approach to identification is useful beyond answering the question of what the relative importance of supply and demand shocks is. For example, Kilian and Murphy (2010) have used a similar approach (applied to a different VAR model) to estimate the speculative component of the price of oil and to test models of speculation in commodity markets. Similar methods have the potential of helping with the identification of shocks to market expectations in a variety of contexts.

We illustrated in the context of the oil market example that commonly used summary statistics for the results of sign-identified VAR models such as median response functions can be misleading. Our analysis complemented the concerns over the use of pointwise posterior median responses already expressed in Fry and Pagan (2005, 2011). We showed that, after imposing additional identifying information, in some cases it may not be necessary to compute summary statistics such as posterior medians because the range of admissible models is sufficiently narrow and the admissible response functions are so similar that qualitatively (or even quantitatively) similar conclusions are implied by all admissible models. There is of course no guarantee that this will happen in general. If the set of admissible models remains large, the most useful exercise will be to search for the admissible model most favorable to each of the competing economic interpretations, as in Faust (1998), and to discuss the range of economic interpretations consistent with the data. A similar strategy may be pursued in cases when no additional identifying restrictions beyond sign restrictions on the responses are available.

## Supporting Information

Additional Supporting Information may be found in the online version of this article:

**Appendix S1.** Constructing a Bound on the Impact Price Elasticity of Oil Supply Elasticity (pdf file)

**Appendix S2.** Data sets and software for replication of the analysis in the article (zip file)

Please note: Blackwell Publishing are not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

## References

- Abraham, Katharine G. and John C. Haltiwanger (1995). "Real Wages and the Business Cycle." *Journal of Economic Literature*, 33, 1215–1264.
- Alquist, Ron and Lutz Kilian (2010). "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, 25, 539–573.

- Barsky, Robert B. and Lutz Kilian (2002). "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative." In *NBER Macroeconomics Annual 2001*, edited by B. Bernanke and K. Rogoff, pp. 137–183.
- Baumeister, Christiane and Gert Peersman (2012). "Time-Varying Effects of Oil Supply Shocks on the U.S. Economy." Working paper, Ghent University.
- Baumeister, Christiane and Gert Peersman (forthcoming). "The Role of Time-Varying Price Elasticities in Accounting for Volatility Changes in the Crude Oil Market." *Journal of Applied Econometrics*.
- Baumeister, Christiane, Gert Peersman, and Ine Van Robays (2010). "The Economic Consequences of Oil Shocks: Differences Across Countries and Time." In *Inflation in an Era of Relative Price Shocks*, edited by R. Fry, C. Jones, and C. Kent. Proceedings of 2009 Conference, Sydney, pp. 91–128.
- Bodenstein, Martin, Christopher J. Erceg, and Luca Guerrieri (2011). "Oil Shocks and U.S. External Adjustment." *Journal of International Economics*, 83, 168–184.
- Canova, Fabio and Gianni De Nicolo (2002). "Monetary Disturbances Matter for Business Cycle Fluctuations in the G-7." *Journal of Monetary Economics*, 49, 1131–1159.
- Canova, Fabio and Matthias Paustian (2010). "Measurement with Some Theory: Using Sign Restrictions to Evaluate Business Cycle Models." Working paper, Pompeu Fabra.
- Davis, Lucas W. and Lutz Kilian (2011). "Estimating the Effect of a Gasoline Tax on Carbon Emissions." *Journal of Applied Econometrics*, 16, 1187–1214.
- Dedola, Luca and Stefano Neri (2007). "What Does a Technology Shock Do? A VAR Analysis with Model-Based Sign Restrictions." *Journal of Monetary Economics*, 54, 512–549.
- Faust, Jon (1998). "On the Robustness of Identified VAR Conclusions about Money." *Carnegie-Rochester Conference Series on Public Policy* 49, 207–244.
- Francis, Neville, Michael Owyang and Roush, Jennifer E. (2007). "A Flexible Finite-Horizon Identification of Technology Shocks." Working paper, Federal Reserve Bank of St. Louis.
- Fry, Renee and Adrian Pagan (2005). "Some Issues in Using VARs for Macroeconometric Research." Working Paper No. 19, Centre for Applied Macroeconomic Analysis, Australian National University.
- Fry, Renee and Adrian Pagan (2011). "Sign Restrictions in Structural Vector Autoregressions: A Critical Review." *Journal of Economic Literature*, 49, 938–960.
- Gambetti, Luca, Evi P. Pappa, and Fabio Canova (2008). "The Structural Dynamics of U.S. Output and Inflation: What Explains the Changes?" *Journal of Money, Credit and Banking*, 40, 369–388.
- Hamilton, James D. (2009a). "Causes and Consequences of the Oil Shock of 2007–08." *Brookings Papers on Economic Activity*, 1, 215–261.
- Hamilton, James D. (2009b). "Understanding Crude Oil Prices." *Energy Journal*, 30, 179–206.
- Kilian, Lutz (2008a). "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature*, 46, 871–909.
- Kilian, Lutz (2008b). "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" *Review of Economics and Statistics*, 90, 216–240.
- Kilian, Lutz (2009a). "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 (2009b). "Comment on 'Causes and Consequences of the Oil Shock of 2007–08' by James D. Hamilton." *Brookings Papers on Economic Activity*, 1, 267–278.
- Kilian, Lutz (2010). "Explaining Fluctuations in Gasoline Prices: A Joint Model of the Global Crude Oil Market and the U.S. Retail Gasoline Market." *Energy Journal*, 31, 87–104.
- Kilian, Lutz and Daniel Murphy (2010). "The Role of Inventories in Identifying Oil Demand and Oil Supply Shocks." Working paper, University of Michigan.
- Lippi, Francesco and Andrea Nobili (forthcoming). "Oil and the Macroeconomy: A Quantitative Structural Analysis." *Journal of the European Economic Association*.
- Moon, H. Roger and Frank Schorfheide (2009). "Bayesian and Frequentist Inference in Partially Identified Models." Working paper, University of Pennsylvania.
- Moon, H. Roger, Frank Schorfheide, Eleonora Granziera, and Miyhe Lee (2009). "Inference for VARs Identified with Sign Restrictions." Working paper, University of Southern California.

- Nakov, Anton and Andrea Pescatori (2010). "Monetary Policy Tradeoffs with a Dominant Oil Producer." *Journal of Money, Credit and Banking*, 42, 1–32.
- Peersman, Gert (2005). "What Caused the Early Millennium Slowdown? Evidence Based on Vector Autoregressions." *Journal of Applied Econometrics*, 20, 185–207.
- Peersman, Gert and Ine Van Robays (2009). "Oil and the Euro Area Economy." *Economic Policy*, 24, 603–651.
- Rubio-Ramirez, Juan F., Dan Waggoner, and Tao Zha. (2010). "Structural Vector Autoregressions: Theory for Identification and Algorithms for Inference." *Review of Economic Studies*, 77, 665–696.
- Scholl, A. and Harald Uhlig (2008). "New Evidence on the Puzzles: Results from Agnostic Identification on Monetary Policy and Exchange Rate." *Journal of International Economics*, 76, 1–13.
- Uhlig, Harald (2005). "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics*, 52, 381–419.