

NBER WORKING PAPER SERIES

SHOCK RESTRICTED STRUCTURAL VECTOR-AUTOREGRESSIONS

Sydney C. Ludvigson
Sai Ma
Serena Ng

Working Paper 23225
<http://www.nber.org/papers/w23225>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2017, Revised July 2018

Ludvigson acknowledges financial support from the C.V. Starr Center for Applied Economics at NYU. Ng acknowledges support from the National Science Foundation under grant SES-0962431 and SES-1558623. We are grateful to Lutz Kilian, José Luis Montiel Olea, Stephen Terry, Jim Stock, Eric Swanson, and Ian Dew-Becker for many valuable discussions. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 Sydney C. Ludvigson, Sai Ma, and Serena Ng. 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.

Shock Restricted Structural Vector-Autoregressions
Sydney C. Ludvigson, Sai Ma, and Serena Ng
NBER Working Paper No. 23225
March 2017, Revised July 2018
JEL No. C01,C5,C51,E17

ABSTRACT

Identifying assumptions need to be imposed on dynamic models before they can be used to analyze the dynamic effects of economically interesting shocks. Often, the assumptions are only rich enough to identify a set of solutions. This paper considers two types of restrictions on the structural shocks that can help reduce the number of plausible solutions. The first is imposed on the sign and magnitude of the shocks during unusual episodes in history. The second restricts the correlation between the shocks and components of variables external to the model. These non-linear inequality constraints can be used in conjunction with zero and sign restrictions that are already widely used in the literature. The effectiveness of our constraints are illustrated using two applications of the oil market and Monte Carlo experiments calibrated to study the role of uncertainty shocks in economic fluctuations.

Sydney C. Ludvigson
Department of Economics
New York University
19 W. 4th Street, 6th Floor
New York, NY 10002
and NBER
sydney.ludvigson@nyu.edu

Serena Ng
Department of Economics
Columbia University
420 West 118th Street
New York, NY 10027
and NBER
Serena.Ng@columbia.edu

Sai Ma
Department of Economics,
New York University
19 W. 4th Street, 6th Floor
New York, NY 10012
sai.ma@nyu.edu

1 Introduction

A challenge in economic analysis is that the data represented by a vector autoregression (VAR) can be consistent with many causal structures differentiated by distinct economic models and primitive assumptions. Structural vector-autoregressive models (SVAR) provide a simple framework that enables researchers to perform counter-factual experiments without fully characterizing all primitives or micro-foundations that lead to the dynamic system. Concisely stated, a n -variable SVAR analysis consists of finding an $n \times n$ matrix \mathbf{B} that relates the reduced form errors $\boldsymbol{\eta}_t$ to the structural errors \mathbf{e}_t :

$$\boldsymbol{\eta}_t = \mathbf{B}\mathbf{e}_t.$$

The data provide $\boldsymbol{\eta}_t$, but the above relationship only provides $n(n+1)/2$ pieces of information about \mathbf{B} . Hence $n(n-1)/2$ restrictions are necessary for identification. Sims (1980) originally proposed a triangularized system to identify \mathbf{B} , but long-run economic restrictions, statistical restrictions based on heteroskedasticity of the VAR innovations, or additional information in the form of high frequency data, variables external to the SVAR, or narrative descriptions have also been employed to identify \mathbf{B} .¹ Point identification exists if these restrictions are enough to yield a unique solution. In some cases, point identification is achievable only by imposing restrictions that are difficult to defend. It may then be desirable to abandon the goal of point identification in favor of less restrictive economic assumptions that are supported by the data.

In this paper we are concerned with applications for which defensible economic restrictions may not permit point identification, but may nonetheless be enough to achieve a substantial constriction of the set of model parameters consistent with the data. Our focus specifically is on two new methods for winnowing the set of plausible solutions, both of which directly restrict the identified structural shocks. The first approach rests on restrictions that require the identified shocks to be consistent with economic reasoning in a small number of extraordinary events, or to accord with our historical understanding of events at particular points in the sample. We refer to these as “event timing constraints,” or simply *event constraints*. The second approach

¹For a comprehensive review of SVAR models, see Ramey (2016), Kilian and Lutkepohl (2016).

rests on non-linear restrictions pertaining to the comovement between components of external variables and the SVAR shocks. We refer to these as “component correlation constraints,” or *correlation constraints* for short. We show in this paper that the event and component correlation constraints can be used separately, and possibly in conjunction with other identification schemes that already exist in the literature. The effectiveness of these constraints in constricting the set of plausible solutions is also investigated.

The motivation for the event constraints is as follows. Every identification scheme necessarily yields a set of primitive shocks \mathbf{e} . While the dynamic causal effects and the forecast error variance attributable to the primitive shocks are often intensively analyzed, a typical SVAR analysis pays much less attention to the estimated shocks themselves, even though the stated goal of the exercise is to identify \mathbf{e} . A credible identification scheme should produce estimates of \mathbf{e} with features that accord with our ex-post understanding of historical events, at least during episodes of special interest. For example, a scheme that identifies a large positive output shock in the 1982 recession would be dismissed because the existence of a shock would be hard to defend given the historical account of the events at the time. Such special events turn out to be valuable for identification because, although two feasible structural models \mathbf{B} and $\tilde{\mathbf{B}}$ will generate shocks \mathbf{e}_t and $\tilde{\mathbf{e}}_t$ with equivalent first and second moments, the \mathbf{e}_t and $\tilde{\mathbf{e}}_t$ are not necessarily the same at any given t . In other words, two series with equivalent properties “on average” can still have distinguishable features in certain subsamples.

The idea of using specific events to identify shocks is not new. For example, the narrative approach to summarizing historical information has been widely used to construct shock series. Most often, this is done by introducing a dummy variable that takes on the value of one when a primitive shock to some variable is thought to have occurred based on a reading of historical events. Oil price shocks have been constructed from disruptive political events, tax shocks from specific fiscal policy announcements, and monetary policy shocks from a reading of FOMC meetings. In such narrative analyses, the constructed shock series is typically used as though it were exogenous and accurately measured, and subsequently used in an SVAR with additional restrictions imposed. But as noted in Ramey (2016), the narrative construction does not ensure

that the shocks are truly exogenous or are correctly measured. Mertens and Ravn (2014) suggest that failure to account for measurement error in the tax shocks constructed by the narrative approach can explain why some SVAR estimates of tax multipliers can be close to zero. Their approach is to use the narrative tax changes as an external instrument, to generate additional moments that complement those generated from the covariance structure. But exogeneity of the tax shocks is still assumed.

The event constraints used in this paper differ from the narrative approach to constructing events in a number of ways. Our event constraints are data driven and not narrative based. We use features of the shocks during selected episodes to determine whether a possible solution is admissible. This is tantamount to creating dummy variables from the timing of specific events, and then putting restrictions on their correlation with the identified shocks. Note that the same SVAR is used to identify all shocks simultaneously; it is not a two-step procedure that identifies some shocks ahead of others. A challenge with exactly identified SVARs is that there are few tools for model validation. In our framework, a recursive ordering can be recovered if such a structure is consistent with the constraints. Such flexibility comes at the expense of foregoing point identification in favor of set identification.

Our second type of economic restriction, what we call component correlation constraints, relates to the external instrumental variable (IV) approach studied by Stock (2008) and Mertens and Ravn (2013) that exploits external variables to identify one shock in a SVAR. For the external IV idea to work, each external variable must be a valid instrument in the sense of satisfying both a relevance and an exogeneity condition. Our approach differs in several ways. We identify not just one, but all shocks in the system, and we do not require the external variables to be exogenous, or at least not known to be exogenous a priori. To be clear that our external variables are not valid instruments, we refer to our external variables as *external proxies* and label them \mathbf{S} . Such proxies must be presumed to have valuable information about the parameters of interest. ‘Valuable’ is defined in terms of a lower bound on the unit-free correlation between relevant components of \mathbf{S} and the identified shocks, akin to the instrument relevance condition.

In what follows, we first use a SVAR for the oil market to illustrate how event and correlation constraints can be separately used with zero and sign restrictions to shrink the set of plausible solutions. The constraints thus have more general applicability than what was considered in Ludvigson, Ma, and Ng (2015) (LMN hereafter). In that paper, both types of constraints were used jointly to study the role of uncertainty shocks in economic fluctuations. We then use the oil application to design simulation experiments. This is used to assess the sampling errors that could be expected of these constraints.

2 Econometric Framework

Let \mathbf{X}_t denote a $n \times 1$ vector time series. We suppose that \mathbf{X}_t has a reduced-form vector autoregression and an infinite moving average representation given, respectively, by:

$$\mathbf{A}(L)\mathbf{X}_t = \boldsymbol{\eta}_t. \quad (1)$$

$$\mathbf{X}_t = \boldsymbol{\Psi}(L)\boldsymbol{\eta}_t \quad (2)$$

$$\boldsymbol{\eta}_t \sim (0, \boldsymbol{\Omega}), \quad \boldsymbol{\Omega} = \mathbb{E}(\boldsymbol{\eta}_t\boldsymbol{\eta}_t') = \mathbf{P}\mathbf{P}'$$

where $\mathbf{A}(L) = \mathbf{I}_n - \sum_{j=1}^p \mathbf{A}_j L^j$, $\boldsymbol{\Psi}(L) = \mathbf{I}_n + \boldsymbol{\Psi}_1 L + \boldsymbol{\Psi}_2 L^2 + \dots$ is a polynomial in the lag operator L of infinite order, and $\boldsymbol{\Psi}_s$ is the $(n \times n)$ matrix of coefficients for the s th lag of $\boldsymbol{\Psi}(L)$. The reduced form innovations $\boldsymbol{\eta}_t$ are related to the structural-form vector autoregression (SVAR) shocks \mathbf{e}_t by an invertible $n \times n$ matrix \mathbf{H} :

$$\boldsymbol{\eta}_t = \mathbf{H}\boldsymbol{\Sigma}\mathbf{e}_t \equiv \mathbf{B}\mathbf{e}_t \quad (3)$$

$$\mathbf{e}_t \sim (0, \mathbf{I}_K), \quad \boldsymbol{\Sigma} = \begin{pmatrix} \sigma_{11} & 0 & \cdot & 0 \\ 0 & \sigma_{22} & 0 & 0 \\ 0 & \cdot & \cdot & 0 \\ 0 & 0 & \cdot & \sigma_{KK} \end{pmatrix}, \quad \sigma_{jj} \geq 0 \forall j. \quad (4)$$

where $\mathbf{B} \equiv \mathbf{H}\boldsymbol{\Sigma}$. The structural shocks \mathbf{e}_t are mean zero with unit variance, and are serially and mutually uncorrelated. The objective is to study the dynamic effects of the structural shocks, or the impulse responses functions (IRF), defined by $\boldsymbol{\Theta}(L) = \boldsymbol{\Psi}(L)\mathbf{B}$,

By definition, $\boldsymbol{\Omega} = \mathbf{P}\mathbf{P}'$ where \mathbf{P} is the unique lower-triangular Cholesky factor with non-negative diagonal elements. Let \mathcal{O}_n be the set of $n \times n$ orthonormal matrices. Now any

$\mathbf{B} = \mathbf{P}\mathbf{Q}$ is consistent with the reduced form covariance matrix $\mathbf{\Omega} = \mathbf{B}\mathbf{B}'$ provided that $\mathbf{Q} = (q_1 \ q_2 \ \dots \ q_n) \in \mathbb{O}_n$. The set of observationally equivalent \mathbf{B} s, given $\mathbf{\Omega}$, can thus be defined as $\mathcal{B} = \{\mathbf{B} = \mathbf{P}\mathbf{Q} : \mathbf{Q} \in \mathbb{O}_n\}$. At this point, the only innocuous restriction we can impose is

$$\text{diag}(\mathbf{H}) = 1, \quad \text{or} \quad \text{diag}(\mathbf{B}) \geq 0, \quad (5)$$

where the restriction on the sign of \mathbf{B} follows from combining the unit effect normalization on \mathbf{H} with the restriction $\sigma_{jj} \geq 0$. We may then interpret a unit change in structural shock j as a standard deviation increase in variable j , and express the set of observationally equivalent \mathbf{B} corresponding to $\mathbf{\Omega}$ as

$$\hat{\mathcal{B}} = \{\mathbf{B} = \mathbf{P}\mathbf{Q} : \mathbf{Q} \in \mathbb{O}_n, \text{diag}(\mathbf{B}) \geq 0\}.$$

But economic theory, intuition, and our understanding of historical events often allow us to discard some \mathbf{B} in $\hat{\mathcal{B}}$. As in Rubio Ramírez, Waggoner, and Zha (2010), Moon, Schorfheide, and Granziera (2013) and related work in this literature, let $\bar{g}_Z(\mathbf{B})$ denote the collection of zero restrictions imposed on the model. Point identification obtains if there is a unique \mathbf{Q} that satisfies $\bar{g}_Z(\mathbf{B}) = \mathbf{0}$. If $\mathbf{Q} = \mathbf{I}_n$, then the solution is the recursive structure consistent with the ordering of the variables used to construct \mathbf{P} . Let f_{zi} be the number of zero restrictions on q_i and $f_z = \sum_{i=1}^n f_{zi}$. For models with variables ordered such that $f_{z1} \geq f_{z2} \dots \geq f_{zn} \geq 0$, Rubio Ramírez, Waggoner, and Zha (2010) show that a necessary and sufficient condition for point identification is $f_{zi} = n - i$ for all $i = 1, \dots, n$. A model is under-identified by zero restrictions if $f_{zi} \leq n - i$ with strict inequality for at least one i .

A model is partially identified when there is more than one \mathbf{Q} that satisfy $\bar{g}_Z(\mathbf{B}) = \mathbf{0}$. In such cases, researchers have used various additional restrictions to dismiss solutions in $\hat{\mathcal{B}}$ leading to a smaller set $\hat{\mathcal{B}}$ that satisfies the additional identifying restrictions. A notable example is sign restrictions, or more generally inequality restrictions, of the form $\bar{g}_S(\mathbf{B}) \geq 0$ (e.g., Uhlig (2005)). Existing theoretical and empirical work tends to place these constraints on the impulse response functions and/or \mathbf{B} . In terms of $\mathbf{X}_t = \mathbf{\Psi}(L)\mathbf{B}\mathbf{e}_t$, the restrictions have focused on $\mathbf{\Psi}(L)\mathbf{B}$. Our restrictions involve the identified shocks \mathbf{e}_t either on their own in the

form of event constraints $\bar{g}_E(\mathbf{e}_t(\mathbf{B}); \bar{\boldsymbol{\tau}}, \bar{\mathbf{k}})$, or combined with external variables in the form of component correlation constraints $\bar{g}_C(\mathbf{e}(\mathbf{B}); \bar{\boldsymbol{\lambda}}, \mathbf{S})$. These constraints are explained in the next two subsections.

2.1 Event Constraints

Event constraints put bounds \bar{k} on the sign and magnitude of $\mathbf{e}_t = \mathbf{B}^{-1}\boldsymbol{\eta}_t$ during selected episodes collected into τ^* . These shocks are useful for identification because, from $\mathbf{e}_t = \mathbf{B}^{-1}\boldsymbol{\eta}_t = \mathbf{Q}'\mathbf{P}^{-1}\boldsymbol{\eta}_t$, we see that for $\tilde{\mathbf{Q}} \neq \mathbf{Q}$,

$$\tilde{\mathbf{e}}_t = \tilde{\mathbf{Q}}'\mathbf{P}^{-1}\boldsymbol{\eta}_t = \tilde{\mathbf{Q}}\mathbf{e}_t \neq \mathbf{e}_t$$

at any given t . This implies that constraints involving the shocks at specific time periods in the sample could be used to constrict the number of solutions in $\hat{\mathcal{B}}$. To illustrate the point, consider the $n = 2$ case:

$$\begin{pmatrix} \eta_{1t} \\ \eta_{2t} \end{pmatrix} = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} \begin{pmatrix} e_{1t} \\ e_{2t} \end{pmatrix}.$$

so that solving for e_{1t} gives

$$e_{1t} = |\mathbf{B}|^{-1}(B_{22}\eta_{1t} - B_{12}\eta_{2t}),$$

where $|\mathbf{B}| = B_{11}B_{22} - B_{12}B_{21}$ is the determinant of \mathbf{B} . The values of η_{1t} and η_{2t} are given since we have data for t in the span $[\tau_1, \tau_2]$. Hence, a restriction on the behavior of e_{1t_1} at specific time t_1 is a non-linear restriction on \mathbf{B} , or equivalently, on \mathbf{Q} . With non-Gaussian reduced-form errors, one can also see that the third and higher order moments of \mathbf{e}_{1t} are not invariant to \mathbf{B} , hence \mathbf{Q} . This is in spite of the fact that the first and second moments of \mathbf{e}_t are invariant to \mathbf{Q} . There is thus information in \mathbf{e}_t that can be used to identify \mathbf{B} ,

Imposing restrictions on \mathbf{e}_t at specific time periods τ^* in SVAR was an idea first considered in LMN. Recent work by Antolin-Diaz and Rubio Ramírez (2016) also suggest to use restrictions on the sign and relative importance of the shocks to help identification. That the properties of the shocks have not been much exploited in identification is somewhat surprising since economic theory and reasoning, narrative interpretations of history, and/or statistical analyses of the data often provide guidance as to what shocks have occurred when, and which ones should be

systematically related to variables we can observe. We denote these event constraints with the vector

$$\bar{g}_E(\mathbf{e}_t(\mathbf{B}); \bar{\boldsymbol{\tau}}, \bar{\mathbf{k}}) = (\bar{g}_{E1}(\mathbf{e}(\mathbf{B}); \bar{\tau}_1, \bar{k}_1), \dots, \bar{g}_{En}(\mathbf{e}(\mathbf{B}); \bar{\tau}_n, \bar{k}_n))' \geq 0,$$

where \mathbf{X} enters because $\boldsymbol{\eta}_t = \mathbf{A}(L)\mathbf{X}_t$. Each $\bar{g}_{Ei}(\mathbf{e}(\mathbf{B}); \bar{\tau}_i, \bar{k}_i)$ represents a vector of constraints of magnitude \bar{k}_i on e_{it} for all $t \in \tau_i^*$. The motivation is that if a \mathbf{Q} implies a shock series that is difficult to defend in certain episodes, it can be removed from $\hat{\mathcal{B}}$. Such constraints could be imposed on extraordinary events such as the major recessions, wars, and natural disasters that have been well-documented. For example, if the first shock (say monetary policy) is presumed to be strongly contractionary in $\tau_i^* = (1979:10, 1979:11, 1979:12)$, then one could formulate

$$\bar{g}_E = \begin{pmatrix} -\sum_{t=1}^T 1_{t=,1979:10} \cdot e_{1t} \\ -\sum_{t=1}^T 1_{t=,1979:11} \cdot e_{1t} \\ -\sum_{t=1}^T 1_{t=,1979:12} \cdot e_{1t} \end{pmatrix} - \begin{pmatrix} \bar{k}_1 \\ \bar{k}_1 \\ \bar{k}_1 \end{pmatrix} \geq 0$$

to dismiss solutions that imply highly expansionary monetary policy shocks in these episodes. The parameter \bar{k}_1 could be two (standard deviations) or some other lower bound that reflects how contractionary these shocks are thought to be. The k -th row of \bar{g}_{E1} represents an inequality with $1_{t \in \tau_{i,k}^*}$ as instrument. In essence, $\bar{g}_E(\mathbf{e}_t(\mathbf{B}); \bar{\boldsymbol{\tau}}, \bar{\mathbf{k}})$ defines conditions based on the timing, sign, and magnitude of the events to help identification. Of course, if \bar{k}_i is too big, or if the timing of the event τ_i^* is inaccurate, the solutions will be meaningless even if they exist. On the other hand, if shocks are systematically found at particular episodes when no restrictions are explicitly imposed, we can be more confident of their occurrence.

2.2 Correlation Constraints

Correlation constraints put bounds \bar{c} on the correlation between the identified shocks and certain components of variables external to the VAR, which we denote by \mathbf{S} . Similar to the event constraints, constraints involving the correlation between $\mathbf{e}_t = \mathbf{B}^{-1}\boldsymbol{\eta}_t$ and observables \mathbf{S} , can be used to constrict the number of solutions in $\hat{\mathcal{B}}$. Thus let

$$\bar{g}_C(\mathbf{e}(\mathbf{B}); \bar{\boldsymbol{\lambda}}, \mathbf{S}) \geq 0,$$

be the collection of component correlation constraints that put possibly nonlinear restrictions across the parameters. An example helps understand the motivation of these constraints.

Consider a two variable model

$$\mathbf{A}(L) \begin{pmatrix} X_{1t} \\ X_{2t} \end{pmatrix} = \begin{pmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{pmatrix} \begin{pmatrix} e_{1t} \\ e_{2t} \end{pmatrix}, \quad \begin{pmatrix} e_1 \\ e_2 \end{pmatrix} \sim \mathcal{N}(0, I_2),$$

where I_2 is a 2×2 identity matrix. The covariance structure $\mathbf{\Omega} = \mathbf{B}\mathbf{B}'$ provides three unique pieces of information, so one more restriction would be needed for point identification. If an instrumental variable Z_t exists such that (i) $\mathbb{E}[Z_t e_{1t}] = 0$ (exogeneity) and (ii) $\mathbb{E}[Z_t e_{2t}] \neq 0$ (relevance), then a unique solution for \mathbf{B} can be obtained.

Suppose that, instead of a valid instrument Z_t , we have an external variable S_t that is not assured to be exogenous. Being endogenous means that S_t is contemporaneously correlated with at least one structural shock. This suggests we could represent S_t by

$$S_t = \gamma e_{1t} + \Gamma e_{2t} + \sigma_S e_{St} \tag{6}$$

where e_{St} is a S -specific shock uncorrelated with e_{1t}, e_{2t} by assumption. We may want to discard solutions in $\hat{\mathbf{B}}$ for which the absolute correlation between S_t and e_{2t} is too small. The quantity

$$c(\mathbf{B}) = \frac{\Gamma}{\sqrt{\Gamma^2 + \gamma^2 + \sigma_S^2}},$$

measures the correlation between the component S_t and e_2 . Requiring that $c(\mathbf{B}) > \bar{c}$ is the same as $\frac{\Gamma^2}{\Gamma^2 + \gamma^2 + \sigma_S^2} > \bar{c}^2$, which is a non-linear constraint on the parameters of the S_t equation, which are themselves functions of the shocks and data on S_t . That c is between zero and one facilitates the parameterization of \bar{c} .

In conventional IV estimation, instrument exogeneity is a maintained assumption. Point identification is achieved by assuming that the instruments have a zero correlation with some shocks and a non-zero correlation with others. By contrast, our approach makes no such exogeneity assumption. We only assert that the events and external variables be driven at least in part by one or more of the shocks, thereby allowing us to narrow the set of solutions but not achieve point identification. Of course, S_t itself is a valid exogenous instrument if $\gamma = 0$. But

when validity of the exogeneity assumption is questionable, then S_t is at best *plausibly exogenous* in the terminology of Conley, Hansen, and Rossi (2012). These authors consider estimation of the equation for X_{1t} with endogenous regressor X_{2t} when instrument exogeneity is not known to hold exactly, but the parameter is point identified when the exogeneity assumption is valid. They put bounds on the effect due to the invalid instrument, or what we refer to as S_t , on X_{1t} . In contrast, we analyze X_{1t} and X_{2t} jointly, and we put more structure on S_t so that a lower bound can be placed between its subcomponent that is relevant as an instrument and the shocks. Like Conley, Hansen, and Rossi (2012), such a bound will not, in general, be enough to achieve point identification. But it could dismiss solutions that do not achieve this bound, akin to dismissing weak instruments.

2.3 Overview

The event and correlation constraints can be used individually, jointly, and possibly in conjunction with zero restrictions, short- or long-run restrictions, and/or sign restrictions. Suppose estimates of \mathbf{B} are required to satisfy all restrictions, i.e., the covariance structure restrictions, the sign restrictions, the event constraints, and correlation constraints. Then the *identified* solution set is defined by

$$\begin{aligned} \bar{\mathcal{B}}(\mathbf{B}; \bar{\mathbf{k}}, \bar{\boldsymbol{\tau}}, \mathbf{S}) &= \{\mathbf{B} = \hat{\mathbf{P}}\mathbf{Q} : \mathbf{Q} \in \mathbb{O}_n, \text{diag}(\mathbf{B}) > 0; \\ &\quad \bar{g}_Z(\mathbf{B}) = 0, \bar{g}_E(\mathbf{B}; \bar{\boldsymbol{\tau}}, \bar{\mathbf{k}}) \geq 0, \bar{g}_C(\mathbf{B}; \mathbf{S}) \geq 0\}. \end{aligned}$$

where $\bar{g}_Z(\mathbf{B})$ is the collection of zero restrictions, $\bar{g}_E(\mathbf{e}_t(\mathbf{B}); \bar{\boldsymbol{\tau}}, \bar{\mathbf{k}})$ is the set of event constraints and $\bar{g}_C(\mathbf{B}; \mathbf{S})$ is the collection of component correlation constraints. To simplify notation, we simply write $\bar{\mathcal{B}}(\mathbf{B}; \bar{\mathbf{k}}, \bar{\boldsymbol{\tau}}, \bar{\boldsymbol{\lambda}}, \mathbf{S})$ as $\bar{\mathcal{B}}$. A particular solution can be in both $\hat{\mathcal{B}}$ and $\bar{\mathcal{B}}$ only if all the event and correlation restrictions are satisfied. Though $\bar{\mathcal{B}}$ is still a set, it should be smaller than $\hat{\mathcal{B}}$, which is based on the covariance restrictions alone. Though no one solution in $\bar{\mathcal{B}}$ is any more likely than another, we sometimes use what will be referred to as the ‘maxG’ solution as

reference point:

$$\mathbf{B}^{\max G} \equiv \arg \max_{\mathbf{B} \in \bar{\mathcal{B}}} \sqrt{\bar{g}(\mathbf{B})' \bar{g}(\mathbf{B})}, \quad \text{where} \quad \bar{g}(\mathbf{B}; \bar{\mathbf{k}}, \bar{\boldsymbol{\tau}}, \bar{\boldsymbol{\lambda}}, \mathbf{S}) = \begin{pmatrix} \bar{g}_Z(\mathbf{B})' \\ \bar{g}_E(\mathbf{B}; \bar{\boldsymbol{\tau}}, \bar{\mathbf{k}})' \\ \bar{g}_C(\mathbf{B}; \bar{\boldsymbol{\lambda}})' \end{pmatrix}'. \quad (7)$$

This is the solution at which the value of the inequalities are jointly maximized. It should, however, be noted that selecting one particular solution for further analysis does not mean that the solution is more likely to explain the data, as discussed in Kilian and Murphy (2012).

In the illustrations to follow, $\bar{g}_Z(\mathbf{B})$ is a vector of $n(n+1)/2$ restrictions implied by the covariance structure of the model, but these are not enough to identify n^2 parameters. The first of two applications in Section 3 does not involve correlation constraints of type $\bar{g}_C(\mathbf{B}; \mathbf{S})$, while the second includes all types of winnowing constraints. In all cases, the possible solutions in $\hat{\mathcal{B}}$ are obtained by initializing \mathbf{B} to be the lower Cholesky factorization of $\boldsymbol{\Omega}$ for an arbitrary ordering of the variables, and then rotating it by $K = 1.5$ million random orthogonal matrices \mathbf{Q} . More precisely, each rotation begins by drawing an $n \times n$ matrix \mathbf{G} of NID(0,1) random variables. Then \mathbf{Q} is taken to be the orthonormal matrix in the \mathbf{QR} decomposition of $\mathbf{G} = \mathbf{QR}$ and $\mathbf{Q}\mathbf{Q}' = \mathbf{I}_n$.

3 An Empirical Analysis of Oil Market

As an example of the identifying potential of the previously proposed shock-restrictions, this section re-examines shocks to the oil market using event and correlation constraints. The SVAR system builds on the oil market system studied in Kilian and Murphy (2012) (KM hereafter). They consider an SVAR with three variables estimated from 1973:01-2008:09:

$$\mathbf{X}_t = \begin{pmatrix} X_{1t} \\ X_{2t} \\ X_{2t} \end{pmatrix} = \begin{pmatrix} \Delta prod_t \\ rea_t \\ rpo_t \end{pmatrix}, \quad (8)$$

where $\Delta prod_t$ is the percentage change in global crude oil production, rea_t is the global demand of industrial commodities variable constructed in Kilian (2009), and rpo_t is the real oil price. KM refer to rea_t simply as an ‘‘aggregate demand’’ shock.

The three structural shocks of interest are collected into the vector \mathbf{e}_t ,

$$\mathbf{e}_t = \begin{pmatrix} e_{1t} \\ e_{2t} \\ e_{3t} \end{pmatrix} = \begin{pmatrix} e_{\Delta prod,t} \\ e_{rea,t} \\ e_{rpo,t} \end{pmatrix}.$$

The first is a shock to the production of crude oil (oil supply shock), the second is a shock to aggregate commodity demand, and the third is a oil market specific demand shock. Kilian (2009) imposes sign restrictions, which we denote $\bar{g}_S(\mathbf{B})$, on the impact responses as follows:

X	Shock		
	Oil supply disruption	Aggregate demand	Oil demand
$\Delta prod_t$	−	+	+
rea_t	−	+	−
rpo_t	+	+	+

The table shows the presumed impact response of each variable named in the row to an impulse in each shock named in the column. Note that an oil supply disruption is by construction a negative supply shock. KM impose further restrictions, denoted $\bar{g}_{KM}(\mathbf{B})$, that limit the elasticity of supply to demand shocks, which amounts to imposing an upper bound on the ratios of impact impulse responses $\frac{B_{13}}{B_{33}}$ and $\frac{B_{12}}{B_{32}}$. The upper bound for both is set to be 0.0258. We refer to the Kilian sign restrictions in combination with the KM elasticity bounds as the KKM restrictions. The next two subsections re-examine the SVAR using some new identifying constraints.

3.1 Application 1: Event Constraints on Oil Production

Much has been written about the correlation between oil price on the one hand, and geopolitical as well economic events on the other. See, for example, Hamilton (2013) and Baumeister and Kilian (2016) for recent reviews. However, their causal relations and more precisely the relative importance of the sources of fluctuations in the oil market is still a matter of debate. The latent shocks that we recover necessarily depend on the identifying assumptions used. Instead of a recursive structure or an elasticity bound, we now consider an application that combines sign restrictions with event constraints to constrict the set of possible solutions.

Figure 1 plots the change in the real oil price over time. The three largest spikes in the unconditional oil price change occur in (i) 1974:01, following the start of the OPEC embargo in

October 1973 (which lasted until March of 1974), (ii) 1986:02, following the collapse of OPEC, and (iii) 1990:08, the month of the Kuwait invasion by the U.S. If a reading of events in these months suggests that the large oil price changes are partly attributable to oil supply shocks, we would expect a spike of the appropriate sign in the structural e_1 shock during these episodes. Such an assumption leads to the following event constraints on oil supply shocks:

Oil event constraints $\bar{g}_E(\mathbf{B}; \bar{\tau}, \bar{\mathbf{k}})$

- i $\bar{k}_1 - e_{1\bar{\tau}_1} \geq 0$ where $\bar{\tau}_1$ is the period 1974:01 during the OPEC Embargo.
- ii $e_{1\bar{\tau}_2} - \bar{k}_2 \geq 0$ where $\bar{\tau}_2$ is the period 1986:02 following the the collapse of OPEC.
- iii $\bar{k}_1 - e_{1\bar{\tau}_3} \geq 0$ where $\bar{\tau}_3$ is the period 1990:08 of the Kuwait invasion.

Note that there is nothing in the above that explicitly precludes the presence of the demand shocks e_2 and e_3 from playing an important role in these episodes, or that restricts the relative importance of the three shocks in any particular episode. The event constraints merely require that supply shocks played some role in the price spikes, where the magnitude and sign of that role are determined by the bounds \bar{k}_1 and \bar{k}_2 . The constraints dismiss parameter estimates that suggest overly favorable supply conditions during the post-OPEC embargo event, or the month of the Kuwait invasion, and too unfavorable conditions during the OPEC collapse. To implement these restrictions, we need to set bounds $\bar{k} \equiv (\bar{k}_1, \bar{k}_2)'$ for the oil supply shocks. In the baseline case, we set $\bar{k}_1 = -3$ and $\bar{k}_2 = 2.5$. This is guided by the fact that the unconditional change in the real price of oil was at least 4 standard deviations higher than its mean in the periods corresponding to events (i) and (iii), and 4 standard deviation lower than its mean in the period corresponding to event (ii). As a sensitivity check, we also consider bounds that are less restrictive by setting $(\bar{k}_1, \bar{k}_2) = (-2.5, 2)$. In total, the $g(\bar{\tau}, \bar{\mathbf{k}}) = \{\bar{g}_S, \bar{g}_E(\cdot)\}$ constraints for this application combine the Kilian sign restrictions with the event constraints.

We extend the data used in KM to the sample 1973:01 to 2016:06. Our baseline case uses six lags in the VAR. After losing observations to lags and differencing, the estimation sample is 1973:08-2016:06. The first price spike occurs in 1974:01, shortly after the onset of the OPEC

embargo that occurred in 1973:10. We therefore estimate the SVAR with lags $p = 6$ so that the large price change episode is in the sample. The KM analysis starts in 1973:01 and uses $p = 24$ lags to capture the long swings in the oil market. We also consider a 24 lag version, which shortens the sample to 1975:02-2016:06. Estimation over this shorter sample eliminates the first event entirely, thereby removing an identifying restriction.

To assess how the constraints affect the identified impulse response functions, we first impose sign restrictions alone as in Kilian (2009). Among 1.5 million rotations, 4,878 solutions satisfy the sign restrictions. We then impose the KKM constraints that combine the sign restrictions with the elasticity bounds of KM. Only 34 solutions satisfy the KKM constraints, implying that the KM constraint alone severely shrinks $\hat{\mathcal{B}}$. Not surprisingly, there is no solution that survives the sign restrictions, the elasticity bounds, and our event constraints. However, there are 2,143 solutions that satisfy our $g(\bar{\tau}, \bar{\mathbf{k}})$ constraints with $\bar{k}_1 = -3$ and $\bar{k}_2 = 2.5$, and 3,147 solutions when $\bar{k}_1 = -2.5$ and $\bar{k}_2 = 2$.

Figure 2 shows the responses of the real price of oil to an oil supply shock, an aggregate commodity demand shock, and an oil market specific demand shock. All figures show responses to one standard deviation shocks in the direction that raise the price of oil. The left panel displays the results under the KKM constraints, and the right panel shows the results under the $g(\bar{\tau}, \bar{\mathbf{k}})$ constraints. Note first that the response of the oil price to the two demand shocks is rather robust to the set of restrictions imposed. Hence we focus on the response to the oil supply shock. The left panel shows that the KKM constraints yield identified responses of the oil price to an oil supply shock that constitute a small subset of those formed from the Kilian sign restrictions alone. The identified responses in this small subset are numerically small (top panel), indicating that the solutions generating the large responses of oil price to a supply shock have been eliminated by the KM constraint.

The right panel of Figure 2 shows the responses produced by our $g(\bar{\tau}, \bar{\mathbf{k}})$ constraints which combine the Kilian sign restriction with the event constraints. This is shown for the base case with $(\bar{k}_1, \bar{k}_2) = (-3, 2.5)$, and also for the less restrictive constraint of $(\bar{k}_1, \bar{k}_2) = (-2.5, 2)$. As seen from the right panel of Figure 2, the responses identified by the tighter bounds are a subset

of those identified by the weaker bound, which are in turn a subset of the responses constructed from the sign restrictions alone. In contrast to the KKM constraints, the $g(\bar{\tau}, \bar{\mathbf{k}})$ restrictions preserve solutions that produce larger oil price responses to an oil supply shock. These results are robust to using 24 lags and losing the first event, as reported in the Appendix. The results are also robust to using the shorter KM sample, and are available on request.

Interestingly, among all solutions that survive our winnowing constraints, the smallest elasticity $\frac{B_{13}}{B_{33}}$ is 0.1792, which is bigger than the KM bound of 0.0258. The lowest elasticity $\frac{B_{12}}{B_{32}}$ is 0.00004. Amongst the solutions that survive the KKM constraints, the smallest $e_{1\bar{\tau}_1}(\mathbf{B})$ is 0.2777 and the largest $e_{1\bar{\tau}_2}(\mathbf{B})$ is 0.6950. But both of these solutions are rejected by our event constraints, even when we use the weaker bound of $(\bar{k}_1, \bar{k}_2) = (-2.5, 2)$. Taken together, the estimates imply that, if we restrict the values for the elasticity of oil production to oil *demand* shocks, as in KM, the survived solutions imply that oil *supply* shocks play little or no role in oil price fluctuations. On the other hand, solutions that satisfy the event constraints imply a greater elasticity of oil supply to oil demand shocks than that permitted by KM.

Since a stated objective of an SVAR analysis is to identify the structural shocks, the properties of the shocks are of interest. The e_{1t} series identified by the $g(\bar{\tau}, \bar{\mathbf{k}})$ constraints exhibits non-Gaussian features. Averaged across solutions, the coefficient of skewness and kurtosis are -0.7736 and 6.8513 , respectively. By contrast, the e_{1t} series implied by the KKM constraints have an average skewness of -1.4603 and a kurtosis of 11.0722 . The stronger departure from Gaussianity arises because the KKM constraints accept solutions that imply larger shocks.²

To have a clearer picture of what shocks are identified by $g(\bar{\tau}, \bar{\mathbf{k}})$, the left panel of Figure 3 plots the timing of “large shocks” over the sample; the solid vertical lines are three event dates given above. Large shocks are defined to be those in excess of two standard deviations above or below the mean. The figure reports the average magnitude of the large shocks so defined across all solutions in the winnowed set $\bar{\mathcal{B}}$. The figure compares the large shock episodes under KKM with those under $g(\bar{\tau}, \bar{\mathbf{k}})$. By design, the e_1 shocks generated by $g(\bar{\tau}, \bar{\mathbf{k}})$ displayed in

²Specifically, the KKM constraints identified large negative e_{1s} in 1975:10, 1977:01, 1978:01, 1980:10, 1983:02, 1986:09, 1990:08, and large positive e_{1s} in 1975:09, 1977:02 and 1987:07.

red should be less than -3 standard deviations in 1974:01 and 1990:08, which is indeed the case. While no additional large oil supply shocks are found, there are 19 occasions when e_{1t} exceeds two standard deviations. By way of comparison, the KKM constraints also identified two negative supply shocks around the events corresponding to our t_1 and t_3 dates, but the negative supply shock in the first event occurs slightly earlier, in 1973:11 rather than 1974:01. On the other hand, the KKM restrictions find no big positive supply shocks in the middle event; on the contrary, all big supply shocks surrounding these dates are negative rather than positive. Interestingly, the KKM constraints only identify a few big demand shocks of either type prior to 1990. On the other hand, the $g(\bar{\tau}, \bar{\mathbf{k}})$ constraints did not find the big negative supply shock associated with the Iranian Revolution around 1978 that the KKM constraints recovered.

We have imposed restrictions on the supply shock e_{1t} , and it is of interest to examine what they imply for the demand shocks e_{2t} and e_{3t} . The left panel of Figure 3 shows that our event constraints by no means preclude other big shocks from occurring at the same time. Indeed, the $g(\bar{\tau}, \bar{\mathbf{k}})$ constraints generate several large aggregate commodity demand shocks of both signs after 2010, several negative oil-specific demand shocks before 1980, and negative demand shocks of both types concurrently in 1986. While demand shocks appear to be important, large supply shocks outside of constrained event dates are numerous under both identification schemes. Regardless of which identification scheme is used, large negative supply shocks coexist with large positive demand shocks in both 1973/74 and 1990.

We have thus far illustrated our approach using restrictions on e_{1t} at three particular episodes. While supply shocks during these three episodes seem reasonable, some might see demand shocks as being relatively more important. In a separate exercise, we place the event constraints on the demand shocks e_{2t} or e_{3t} instead of e_{1t} , thereby forcing at least one them to be large in the earlier event dates characterized by large price changes. This approach is agnostic as to which demand shock moves. That is, the restrictions require that *at least one* of e_{2t} or e_{3t} be large and positive (≥ 3 standard deviations above the mean) in 1974:01 and 1990:08, but large and negative (≤ -2 standard deviations) in 1986:02. Sixty solutions survived these constraints. But as seen from the right hand panel of Figure 3, we still find supply shocks

during the oil embargo in 1973:11 and again in the Kuwait invasion month 1990:08 to be large and negative. This is similar to when the supply shock e_{1t} was explicitly constrained, though the negative supply shock in the first episode is now two months earlier than the price spike itself. However, e_{1t} does not spike up at all in any month of 1986. Together with the consistent finding of large negative demand shocks in 1986, the interpretation of the 1986 OPEC collapse episode is less clear. We return to this point below.

The above analysis focused on the identified sets of impulse responses. It is straight forward to likewise construct sets of variance decompositions. We briefly report a few of these statistics here. For the long-horizon (infinite-step-ahead) forecast error variance, the KKM restrictions imply that effectively none of the forecast error variance in the real price of oil is attributable to supply shocks: the fraction explained by supply shocks covers the range [0.02, 0.04]. That is, of all the solutions that survive the KKM restrictions, the lowest fraction of variance explained is 2% and the highest is 4%. By contrast, the aggregate commodity demand shock explains between 32% and 91% of the oil price variance across the solutions, while the oil demand shock explains between 7% and 65%. The corresponding numbers using the $g(\bar{\tau}, \bar{\mathbf{k}})$ constraints are [0.19, 0.68] for the supply shock, [0.19, 0.80] for the aggregate demand shock, and [0.00, 0.32] for the oil demand shock. Thus, supply shocks play a much larger relative role in oil price fluctuations under these identifying restrictions. Under either type of identifying restrictions, the relative importance of the demand shocks ranges widely across the solutions in the winnowed set. These ranges can be tightened with additional constraints on the forecast errors attributable to the identified shocks. Such constraints can also be interpreted as a type of correlation constraint.

The purpose of this application is to show that event constraints can be combined with sign restrictions to shrink the set of plausible solutions. Our intent is not to argue in favor of one event constraint over another, but rather to show that shocks identified by different \mathbf{Q} s have distinguishable features at specific points in the sample even though they may have equivalent first and second moments computed from the full sample. Such distinguishable features can substantially narrow the set of solutions deemed credible. The results of this application bear

this out.

3.2 Application 2: Shortfall in Oil Production as External S_t

Instrumental variables are typically chosen because economic reasoning or theory suggests that they ought to be correlated with the endogenous variable of interest. In the oil example above, some have argued that oil production shortfalls are at least partly caused by political events such as wars or embargoes (e.g., Hamilton (2013)). Hence if the production shortfalls are uncorrelated with the two demand shocks e_{2t} and e_{3t} , one might consider a measure of such shortfalls as an instrument for oil supply shocks, if data for such a variable is available. Kilian (2008) constructed such a series, which Olea, Stock, and Watson (2015) use as an external instrument in an SVAR designed to identify oil supply shocks.³

To use the shortfall series as an external instrument, the variable must be truly exogenous and the data measured without error. Relaxing the exogeneity assumption would yield inconsistent estimates. However, as discussed above, even if the shortfall series is not exogenous with respect to the two demand shocks, it may still be relevant for oil supply shocks and we can restrict the set of possible solutions by placing a lower bound on the presumed relevance. We therefore consider a second oil shock application that illustrates the use of the correlation constraints discussed above, in which the oil supply shortfall variable is our S_t . The data for this application are based on the replication files of Olea, Stock, and Watson (2015) and span the period 1973:02 to 2004:09, or 1973:09-2004:09 after losing observations to lags and differencing. Hence the sample is slightly different from the previous application.

The correlation constraint for this application is defined as follows:

Oil Correlation Constraint $\bar{g}_C(\mathbf{B}; \mathbf{S})$:

³This variable is an estimate of production short-falls caused by wars or civil disturbances in OPEC countries. As explained in Kilian (2008), the production short-fall variable is generated by first computing a counterfactual production level for the country in question that would have occurred if the war had not. The counterfactual production level is generated by extrapolating its pre-war production level based on the average growth rate of production in other countries that are subject to the same global macroeconomic conditions and economic incentives, but are not involved in the war. Kilian chooses the countries that belong in the benchmark group on a case-by-case basis drawing on historical accounts and industry sources.

- $|c(\mathbf{B})| > \bar{c}$, where $c(\mathbf{B}) = \text{corr}(S_t, e_{1t}(\mathbf{B}))$ is the sample correlation between S_t and the supply shock $e_{1t}(\mathbf{B})$.

As c is a correlation coefficient, \bar{c} must be between zero and one in absolute value. Olea, Stock, and Watson (2015) are concerned that the production shortfall variable could be a weak instrument. We do not treat S_t as a valid instrument, but nonetheless assume that it is relevant for supply shocks while taking the weak IV consideration into account. In the baseline implementation, we set a relatively unrestrictive bound of $\bar{c} = 0.03$. We assess the sensitivity to this bound.

The results are displayed in Figure 4. The benchmark is the identified set of impulse responses based on Kilian’s sign restrictions alone, $\bar{g}_S(\mathbf{B})$. These are shown on the left. This is compared to the set of responses after imposing the correlation constraint $\bar{g}_C(\mathbf{B}; \mathbf{S})$, combining them with sign restrictions, and also adding the oil event constraints on e_{1t} considered in the previous application to both \bar{g}_S and \bar{g}_C to arrive at $g(\bar{\tau}, \bar{\mathbf{k}}, S) = \{\bar{g}_S, \bar{g}_C(\cdot), \bar{g}_E(\cdot)\}$. These results are shown on the right panel. Because the sample size is different from the first application and correlation constraints are now imposed, the \bar{k} bounds used earlier renders the winnowed set empty. Hence we adjust the bounds on the event constraints to be less restrictive at $\bar{k}_1 = -2$ and $\bar{k}_2 = 2$, compared with $\bar{k}_1 = -3$ and $\bar{k}_2 = 2.5$ used earlier. Among 1.5 million rotations, 3,838 solutions satisfy the sign restrictions. While 687 solutions satisfy the sign and correlation constraints, only 30 solutions satisfy sign, event, and correlation constraints. Further investigation reveals that the main reason that only 30 solutions survive all three constraints is that few solutions deliver a positive oil supply shock in 1986:02.

The plots on the left show that sign restrictions alone leave a wide range of possible responses to each structural shock, with no clear response pattern associated with any one structural shock, as KM have previously emphasized. The responses based on sign and correlation constraints are narrower, while those based on $g(\bar{\tau}, \bar{\mathbf{k}}, S)$ that combine the three sets of restrictions are significantly narrower for all three shocks. That the identified impulse responses are narrower than using sign restrictions alone arises in part because external information in the form of correlation constraints contributes to identification, especially for the responses to the

aggregate demand shock e_{2t} . The event constraints are relatively more important for the supply shock. Taken together, this analysis suggests that oil supply shocks and aggregate commodity demand shocks have large and persistent effects on the price of oil, while oil specific demand shocks have much smaller effects.

To learn more about the effectiveness of the correlation constraints, Figure 5 depicts the large shocks based on sign and correlation constraints, but without imposing event constraints. A direct comparison with Figure 3 must be made with care because the sample for the second application ends in 2004:9. Keeping in mind this difference, we see that imposing \bar{g}_C yields large negative oil supply shocks in 1974:01 and 1990:08, the previously imposed event constraint dates and both episodes of dramatic price spikes upward. This shows that the identified set obtained by imposing the component correlation constraints using the external proxy variable S_t implies that large negative supply shocks contributed non-trivially to the price spikes during the oil embargo and Kuwait invasion, even if we don't impose event constraints in those episodes. However, a large (positive) supply shock is found in the OPEC collapse of 1986:02 only if we additionally impose an event constraint that seems incongruent with all but eight solutions. Meanwhile there were large negative shocks to aggregate demand e_2 in 1986:02, and oil specific demand in 1986:09. Taking the results from the two applications together, the sharp drop in oil price in 1986 is likely due not to one shock, but to a combination of shocks. This is unlike the oil price hikes during the oil embargo and the Kuwait invasion when evidence for important negative supply shocks is more compelling.

4 Repeated Sampling Simulation

In point-identified models, sampling uncertainty can be evaluated using frequentist confidence intervals or Bayesian credible regions, and they coincide asymptotically. Inference for set-identified SVARs is, however, more challenging because no consistent point estimate is available. As pointed out in Moon and Schorfheide (2012), the credible regions of Bayesian identified impulses responses will be distinctly different from the frequentist confidence sets, with the implication that Bayesian error bands cannot be interpreted as approximate frequentist error

bands. Our analysis is frequentist, and while the two applications presented above illustrate how the dynamic responses vary across estimated models, where each model is evaluated at a solution in $\bar{\mathcal{B}}(\mathbf{B}; \bar{\mathbf{k}}, \bar{\boldsymbol{\tau}}, \mathbf{S})$, we still need a way to assess the robustness of our procedure, especially since it is new to the literature.

Unfortunately, few methods are available to evaluate the sampling uncertainty of set identified SVARs from a frequentist perspective, and these tend to be specific to the imposition of particular identifying restrictions. Moon, Schorfheide, and Granziera (2013) suggest a projections based method within a moment-inequality setup, but it is designed to study SVARs that only impose restrictions on one set of impulse response functions. Furthermore, the method is computationally intense, requiring a simulation of critical value for each rotation matrix. Gafarov, Meier, and Olea (2015) suggest to collect parameters of the reduced form model in a $1 - \alpha$ Wald ellipsoid but the approach is conservative. For the method to get an exact coverage of $1 - \alpha$, the radius of the Wald-ellipsoid needs to be carefully calibrated. As discussed in Kilian and Lutkepohl (2016), even with these adjustments, existing frequentist confidence sets for set-identified models still tend to be too wide to be informative. It is fair to say that there exists no generally agreed upon method for conducting inference in set-identified SVARs.

We use a bootstrap/Monte Carlo experiment to assess the robustness of our inequality restrictions when S_t is a variable external to the three variable SVAR.

Let M be the number of solutions in the identified set $\bar{\mathcal{B}}(\mathbf{B}; \bar{\mathbf{k}}, \bar{\boldsymbol{\tau}}, \bar{\boldsymbol{\lambda}}, \mathbf{S})$ and let m index an arbitrary solution in the set. Let R be the number of replications in a repeated sampling experiment. To generate samples of the structural shocks from this solution in a way that ensures the events that appear in historical data also occur in our simulated samples, we draw randomly with replacement from the sample estimates of the shocks \mathbf{e}_t^m for the m th solution, with the exception that we fix the values for these shocks in each replication in the periods $\bar{\tau}_1$, $\bar{\tau}_2$ and $\bar{\tau}_3$ to be the observed ones, where $\bar{\tau}_1$ is the period 1974:01 of OPEC embargo and $\bar{\tau}_2 = 1986:02$ following the collapse of OPEC and $\bar{\tau}_3 = 1990:08$ of the Kuwait invasion. Each draw of the shocks \mathbf{e}_t^m is combined with the \mathbf{B} parameters of the m th solution, \mathbf{B}^m , and the reduced form estimates of the VAR parameters $\mathbf{A}(L) = \mathbf{I} - \mathbf{A}_1L - \dots - \mathbf{A}_pL^p$, to generate

$R = 1,000$ samples of size T of $\boldsymbol{\eta}_t = \mathbf{B}^m \mathbf{e}_t^m$ and \mathbf{X}_t using the SVAR $\mathbf{X}_t = \sum_{j=1}^p \mathbf{A}_j \mathbf{X}_{t-j} + \mathbf{B}^m \mathbf{e}_t^m$. In each of these new samples $r = 1, \dots, R$ we regress \mathbf{X}_t on p lags of itself and generate a new values of $\mathbf{B}^{mr} = \mathbf{P}^{mr} \mathbf{Q}$ and the dynamic responses to shock j summarized by the IRF:

$$\frac{\partial \mathbf{X}_{t+s}}{\partial e_{jt}} = \boldsymbol{\Psi}_s^{mr} \mathbf{b}^{mrj}, \quad (9)$$

where \mathbf{B}^{mr} is the value of \mathbf{B} in the r th replication generated by drawing from the shocks of the m th solution, \mathbf{b}^{mrj} is the j th column of \mathbf{B}^{mr} and the coefficient matrixes $\boldsymbol{\Psi}_s^{mr}$ are given by $\boldsymbol{\Psi}^{mr}(L) = \boldsymbol{\Psi}_0^{mr} + \boldsymbol{\Psi}_1^{mr} L + \boldsymbol{\Psi}_2^{mr} L^2 + \dots = \mathbf{A}^{mr}(L)^{-1}$.

To generate samples of S from this solution in a way that ensures that the correlations with the oil supply shocks that appear in our historical data also appear in our simulated samples, we generate idiosyncratic oil supply shocks e_{St} from

$$e_{St}^m = \underbrace{S_t - \hat{\rho}_s S_{t-1}}_{\hat{u}_t} - d_1 e_{1t}^m - d_2 e_{2t}^m - d_3 e_{3t}^m, \quad (10)$$

where $\hat{\rho}_s$ is the estimated first-order autocorrelation coefficient for S_t from historical data, and the $d = (d_1, d_2, d_3)'$ parameters in (10) are calibrated to target the observed correlations $c_1^m = \text{corr}(u_t, e_{1t}^m)$ for the m th solution in historical data. Given the d parameters, observations on \mathbf{e}_t^m and \hat{u}_t , we observe e_{St}^m on the left-hand-side of (10). We generate 1,000 samples of S_t by drawing with replacement from these e_{St} , in the same manner described above for \mathbf{e}_t^m and recursively iterating on (10) using the first observations on S_1 in our historical sample as initial values where S_t is the return on the CRSP value-weighted index. Note that the auto-regressive residual for S_t , upon which our correlation constraints are based, is a function of the four shocks: $S_t - \hat{\rho}_s S_{t-1} = e_{St}^m + d_1 e_{1t}^m + d_2 e_{2t}^m + d_3 e_{3t}^m$.

For each of these $R = 1,000$ replications, we construct an identified set of solutions $\bar{\mathbf{B}}(\cdot)$. In each replication r , $K = 1.5$ million possible solutions for \mathbf{B}^{mr} are generated by initializing \mathbf{B}^{mr} to be the lower Cholesky factorization of $\boldsymbol{\Omega}^{mr}$ for an arbitrary ordering of the variables. These are then rotated by 1.5 million random orthogonal matrices \mathbf{Q} .

We construct confidence sets for the set of solutions in repeated samples as follows. In each replication of each solution, $K = 1.5$ million rotation matrices are entertained, but only

$K_{mr} \leq K$ rotations of \mathbf{Q} will generate solutions that are admitted into the identified set for that replication, $\bar{\mathcal{B}}^{mr}(\cdot)$. Let $\Theta_{i,j,s}^{m,r,k}$ be the s -period ahead response of the i th variable to a standard deviation change in shock j at the k -th rotation of replication r of the m th solution. Let $\underline{\Theta}_{i,j,s}^{m,r} = \min_{k \in [1, K_{mr}]} \Theta_{i,j,s}^{m,r,k}$ and $\bar{\Theta}_{i,j,s}^{m,r} = \max_{k \in [1, K_{mr}]} \Theta_{i,j,s}^{m,r,k}$. Each $(\underline{\Theta}_{i,j,s}^{m,r}, \bar{\Theta}_{i,j,s}^{m,r})$ pair represents the extreme (highest and lowest) dynamic responses in replication r of solution m . From the quantiles of the set $\{\underline{\Theta}_{i,j,s}^{m,r}\}_{m=1, r=1}^{M,R}$ that includes all replications for all solutions we can obtain the $\alpha/2$ critical point $\underline{\Theta}_{i,j,s}(\alpha/2)$. Similarly, from the quantiles of $\{\bar{\Theta}_{i,j,s}^{m,r}\}_{m=1, r=1}^{M,R}$, we have the $1 - \alpha/2$ critical point $\bar{\Theta}_{i,j,s}(1 - \alpha/2)$. Eliminating the lowest and highest $\alpha/2$ percent of the samples gives a $(1 - \alpha)\%$ percentile-based confidence interval defined by

$$CI_{\alpha,g} = \left[\underline{\Theta}_{i,j,s}(\alpha/2), \bar{\Theta}_{i,j,s}(1 - \alpha/2) \right].$$

$CI_{\alpha,g}$ denotes the confidence intervals for sets of solutions that satisfy all constraints, including the event and correlation constraints: $\bar{g}_Z(\mathbf{B}) = 0$, $\bar{g}_E(\mathbf{B}; \bar{\tau}, \bar{\mathbf{k}}) \geq 0$, $\bar{g}_C(\mathbf{B}; \mathbf{S}) \geq 0$. We use $CI_{\alpha,gZ}$ to denote the confidence intervals for sets of solutions that satisfy only the reduced form covariance restrictions $\bar{g}_Z(\mathbf{B}) = 0$.

The confidence intervals $CI_{\alpha,g}$ (light shaded area) and $CI_{\alpha,gZ}$ (white area) for the identified set are displayed in Figure (6) along with the identified set itself (dark shaded area). The results show that the confidence intervals $CI_{\alpha,g}$, formed from estimations that impose the event and correlation constraints are noticeably narrower than $CI_{\alpha,gZ}$ formed from estimations that impose only covariance restrictions.

5 Other Tests

A final result of interest is whether this particular SVAR is compatible with a recursive structure. Recursivity is a convenient assumption and is often used in the empirical literature on the effects of uncertainty shocks. The identifying restrictions employed here impose no such structure, and it is straight forward to check whether the estimated values of B are consistent with it. Figure ?? plots the distribution over 5,000 replications of \hat{B}_{YM} , \hat{B}_{MY} , \hat{B}_{YF} , and \hat{B}_{MF} based on the “max-C” solution. None of the distributions are centered around zero. The implication is that

the recursive structure is inconsistent with any ordering of the variables.

6 Conclusion

Identifying assumptions need to be imposed in order to give impulse responses generated by autoregressions an economically meaningful interpretation. But our information in the form of commonly used zero and sign restrictions are often not rich enough to conclude that the data are consistent with a clear causal pattern among the variables. This paper explores the properties of two new types of identifying restrictions as moment inequalities that can help constrain the number of plausible solutions. The first restricts the sign and magnitude of identified shocks to accord with an historical understanding of events at particular points in the sample, while the second restricts the correlations between the identified shocks and components of external variables. We use applications and simulations to show how these constraints can be used. The issue of how to best conduct frequentist inference in set-identified SVARs remains an important topic for future research.

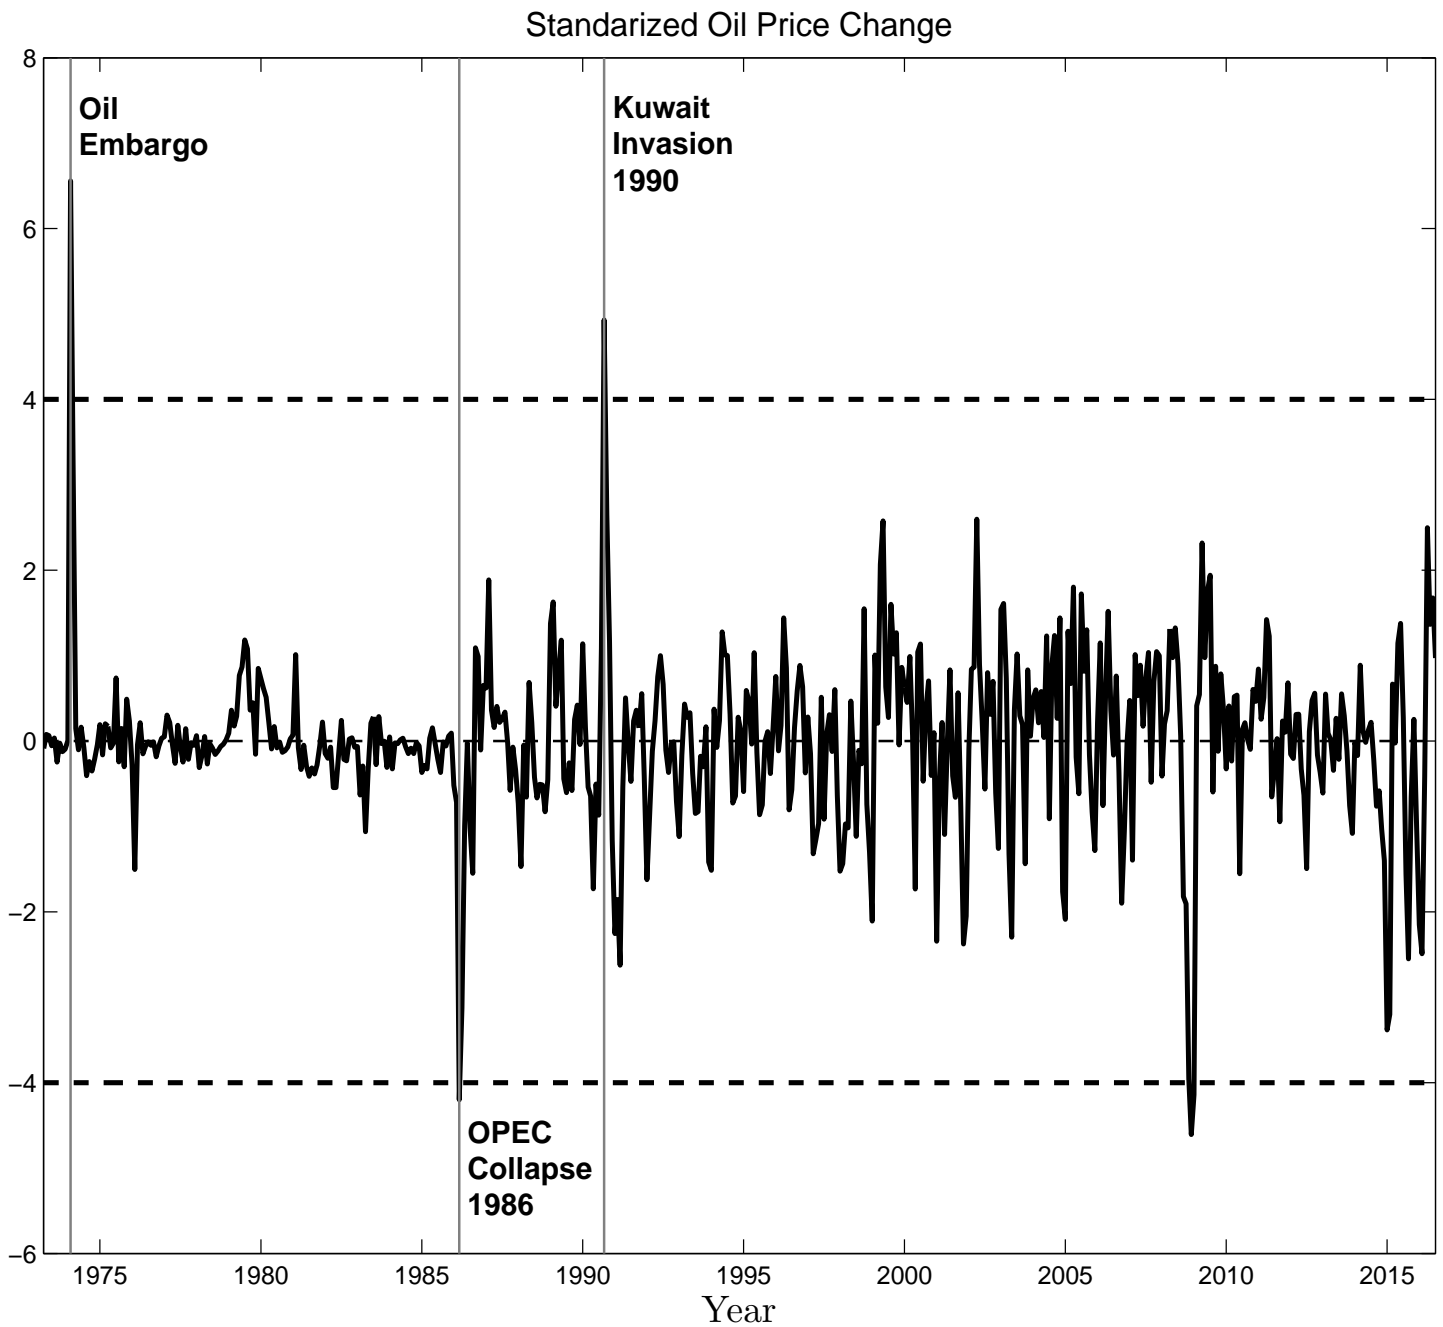
References

- ANTOLIN-DIAZ, J., AND J. RUBIO RAMÍREZ (2016): “Narrative Sign Restrictions for SVARs,” Unpublished paper, Emory University.
- BAUMEISTER, C., AND L. KILIAN (2016): “Forty years of oil price fluctuations: Why the price of oil may still surprise us,” *The Journal of Economic Perspectives*, 30(1), 139–160.
- CONLEY, T. G., C. B. HANSEN, AND P. E. ROSSI (2012): “Plausibly Exogenous,” *The Review of Economics and Statistics*, 94(1), 260–272.
- GAFAROV, B., M. MEIER, AND J. M. OLEA (2015): “Projection Inference for Set-Identified SVARs,” Unpublished paper, New York University.
- HAMILTON, J. D. (2013): “Historical Oil Shocks,” in *Routledge Handbook of Major Events in Economic History*, ed. by R. E. Parker, and R. Whaples, pp. 49–56. Routledge Taylor and Francis Group, New York.
- KILIAN, L. (2008): “Exogenous oil supply shocks: how big are they and how much do they matter for the US economy?,” *The Review of Economics and Statistics*, 90(2), 216–240.
- (2009): “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, 99(3), 1053–1069.
- KILIAN, L., AND H. LUTKEPOHL (2016): *Structural Vector Autoregressive Analysis*. Cambridge University Press.
- KILIAN, L., AND D. 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.
- LUDVIGSON, S. C., S. MA, AND S. NG (2015): “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?,” <http://www.econ.nyu.edu/user/ludvigsons/ucc.pdf>.

- MERTENS, K., AND M. O. RAVN (2013): “The dynamic effects of personal and corporate income tax changes in the United States,” *The American Economic Review*, 103(4), 1212–1247.
- (2014): “A reconciliation of svar and narrative estimates of tax multipliers,” *Journal of Monetary Economics*, 68, S1–S19.
- MOON, H. R., AND F. SCHORFHEIDE (2012): “Bayesian and frequentist inference in partially identified models,” *Econometrica*, 80(2), 755–782.
- MOON, H. R., F. SCHORFHEIDE, AND E. GRANZIERA (2013): “Inference for VARs Identified with Sign Restrictions,” Unpublished Paper.
- OLEA, J. L. M., J. H. STOCK, AND M. D. WATSON (2015): “Inference in Structural VARs with External Instruments,” Unpublished manuscript, Harvard University.
- RAMEY, V. A. (2016): “Macroeconomic shocks and their propagation,” NBER Working paper No. 21978.
- RUBIO RAMÍREZ, J. F., D. F. WAGGONER, AND T. ZHA (2010): “Structural vector autoregressions: Theory of identification and algorithms for inference,” *The Review of Economic Studies*, 77(2), 665–696.
- SIMS, C. A. (1980): “Macroeconomics and Reality,” *Econometrica*, pp. 1–48.
- STOCK, J. H. (2008): “What’s New in Econometrics: Time Series,” *NBER Summer Institute*.
- UHLIG, H. (2005): “What are the effects of monetary policy on output? Results from an agnostic identification procedure,” *Journal of Monetary Economics*, 52(2), 381–419.

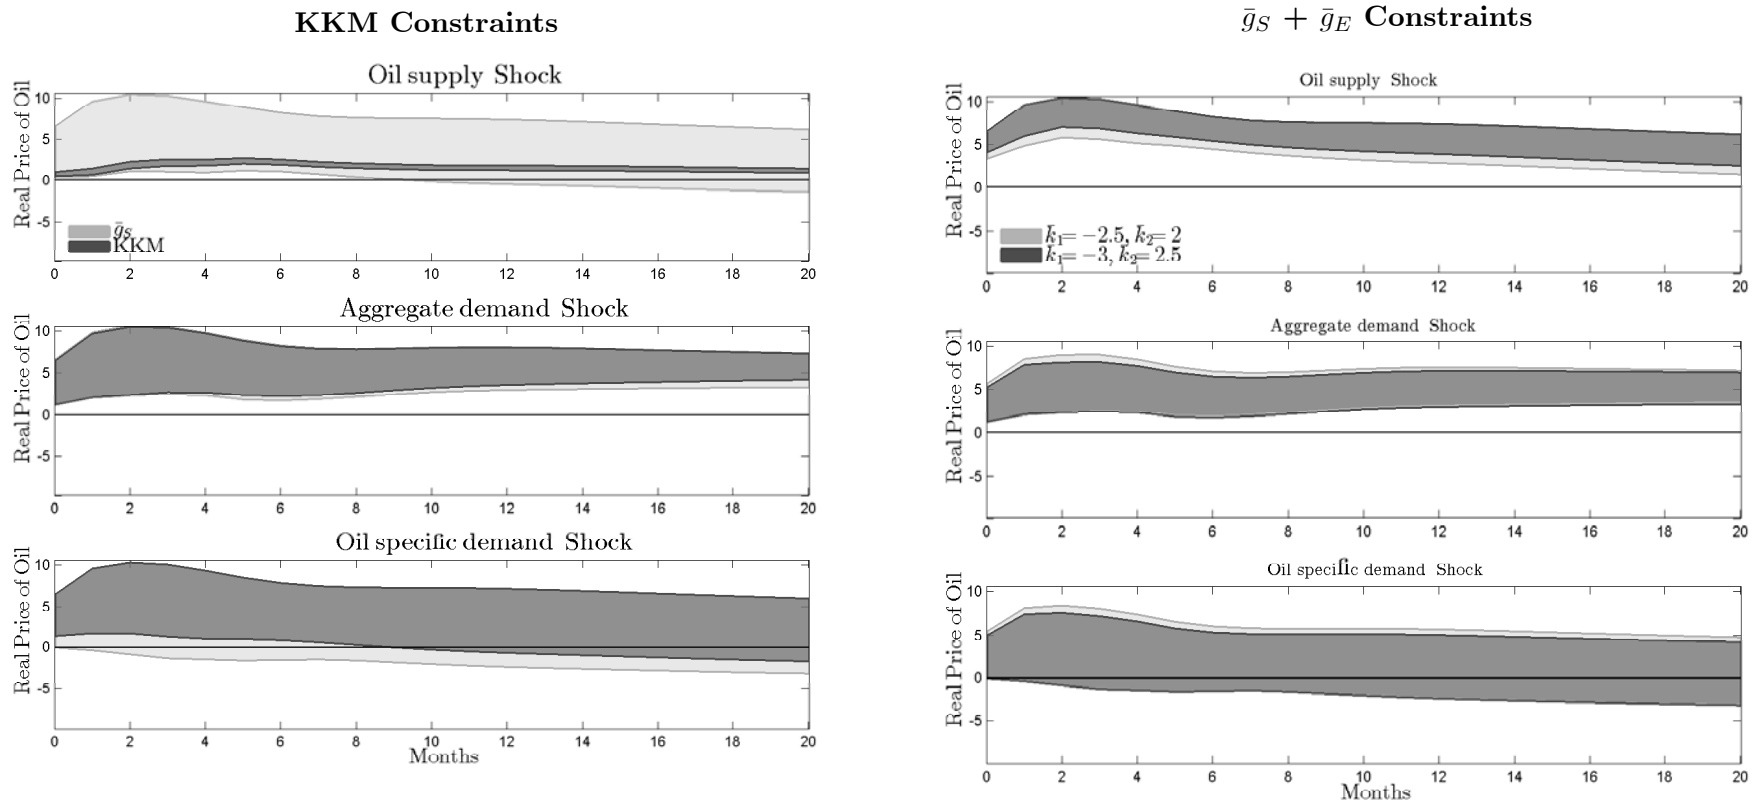
Figures and Tables

Figure 1: Oil Price Change



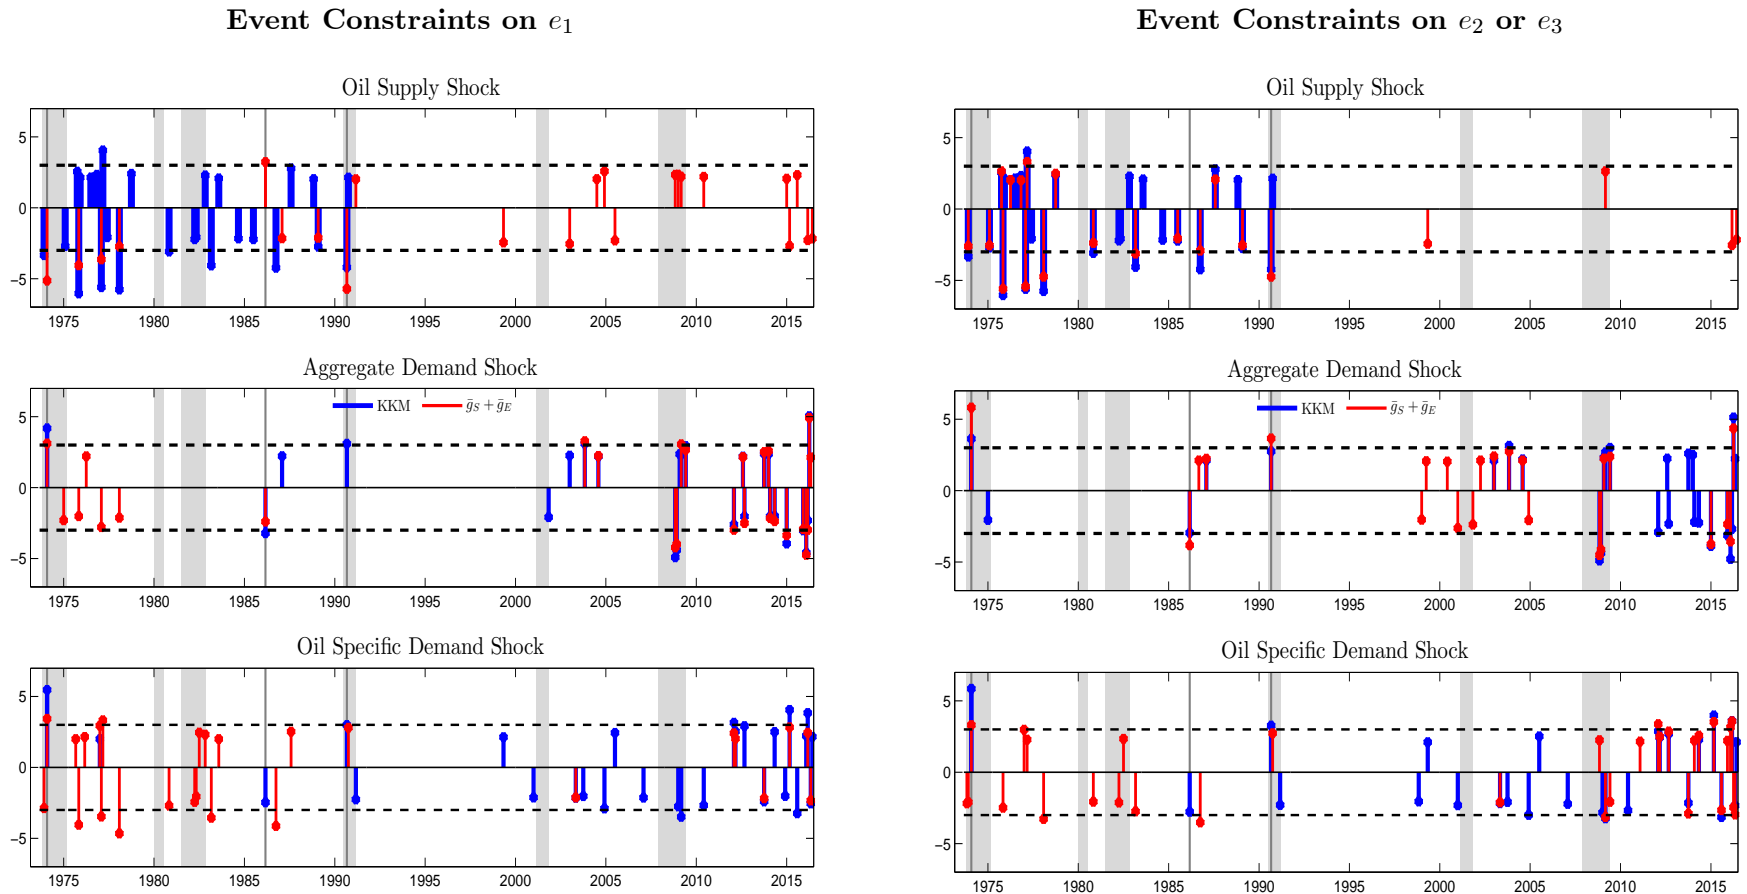
The standardized change in the real oil price change is reported. The sample spans the period 1973:01 to 2016:06.

Figure 2: Oil Shock Application 1



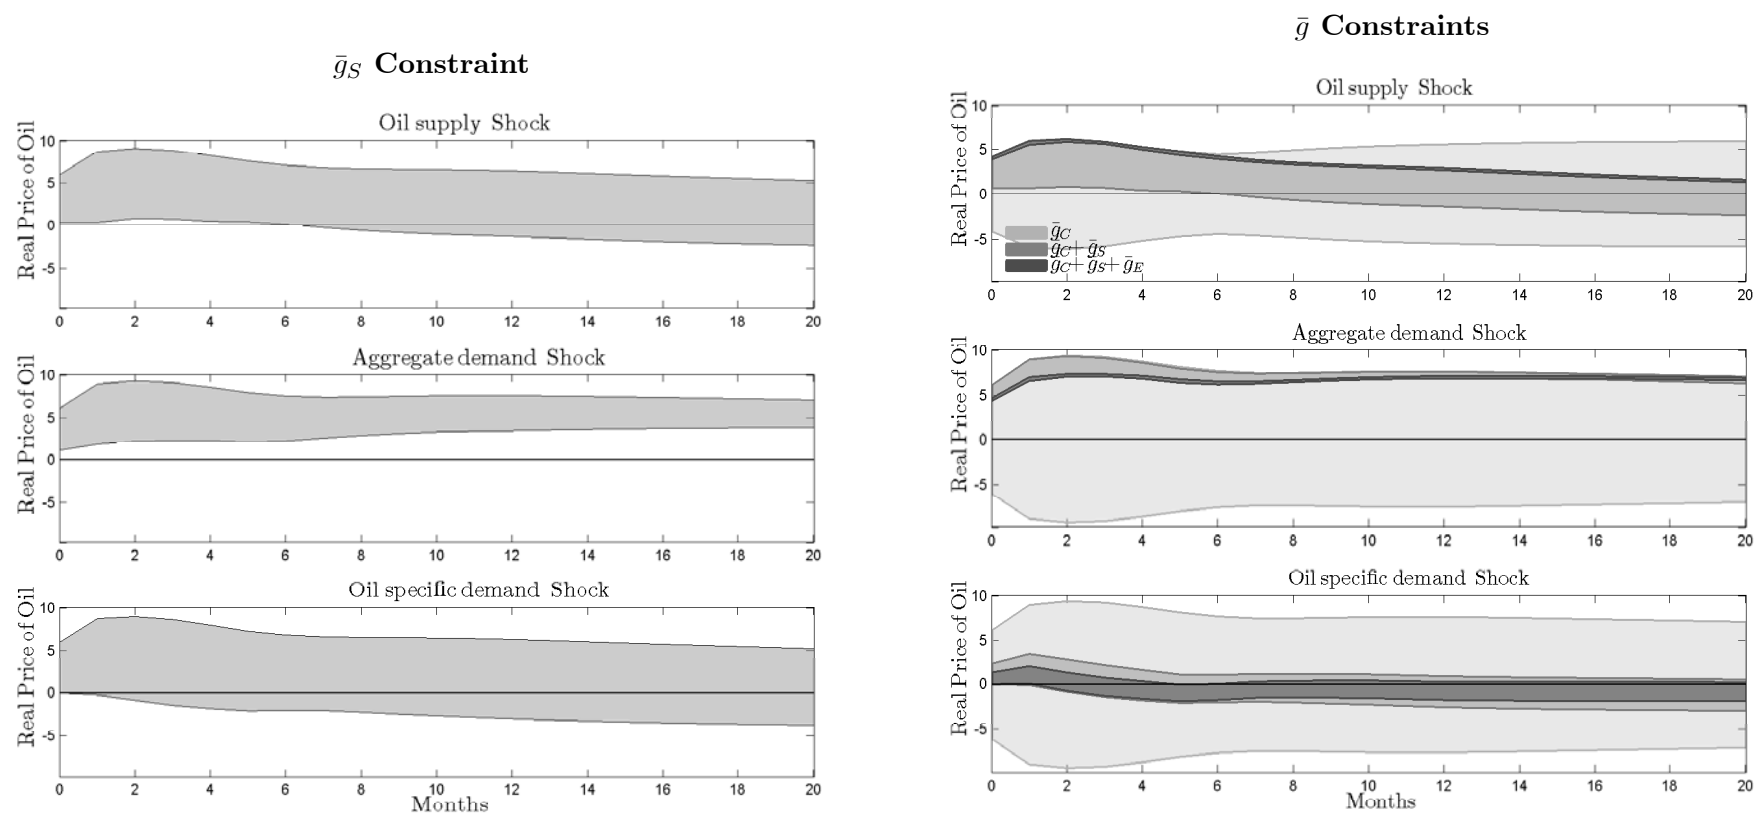
Shaded areas stack the impulse responses for all admissible solutions based on the type of restriction listed in the subplot heading. The sample spans the period 1973:01-2016:06.

Figure 3: Large Shocks in Oil Application 1



The figure exhibits shocks that are at least 2 standard deviations above/below the unconditional mean for e . The horizontal line corresponds to 3 standard deviations shocks. The solid vertical lines are the event dates 1974:01, 1986:02, and 1990:08 discussed in the text. The light grey shaded areas give the dates of NBER recessions. The sample spans the period 1973:02 to 2004:09.


Figure 4: Oil Shock Application 2




Shaded areas stack the impulse responses for all admissible solutions based on the type of restriction listed in the subplot heading. \bar{g}_C constraint is based on $\bar{c}=0.03$. The sample spans the period 1973:02-2004:09.

Figure 5: Large Shocks in Oil Application 2


Oil Supply Shock



Aggregate Demand Shock

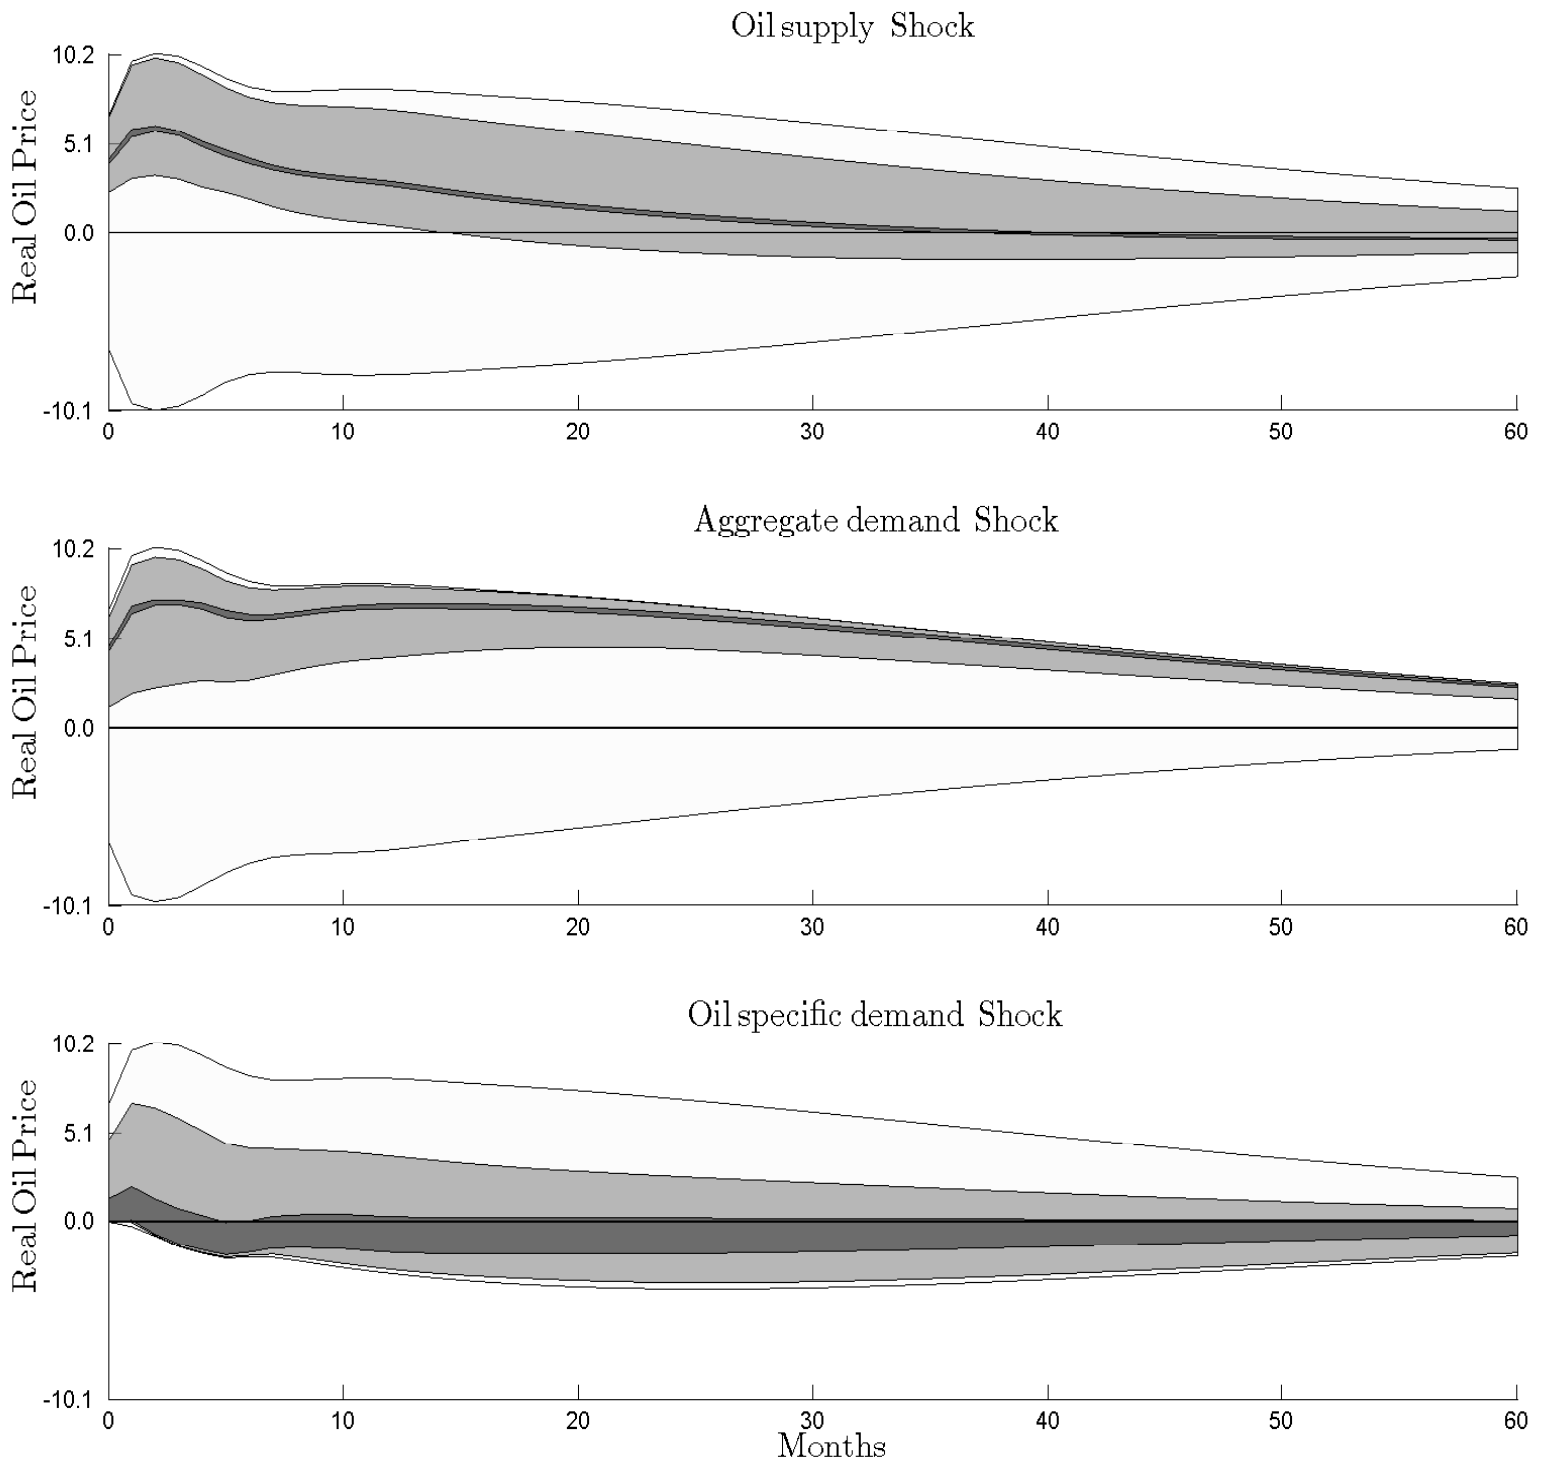


Oil Specific Demand Shock



The figure exhibits shocks that are at least 2 standard deviations above/below the unconditional mean for e . The horizontal line corresponds to 3 standard deviations shocks. The solid vertical lines are the event dates 1974:01, 1986:02, and 1990:08 discussed in the text. The light grey shaded areas give the dates of NBER recessions. The sample spans the period 1973:02 to 2004:09.

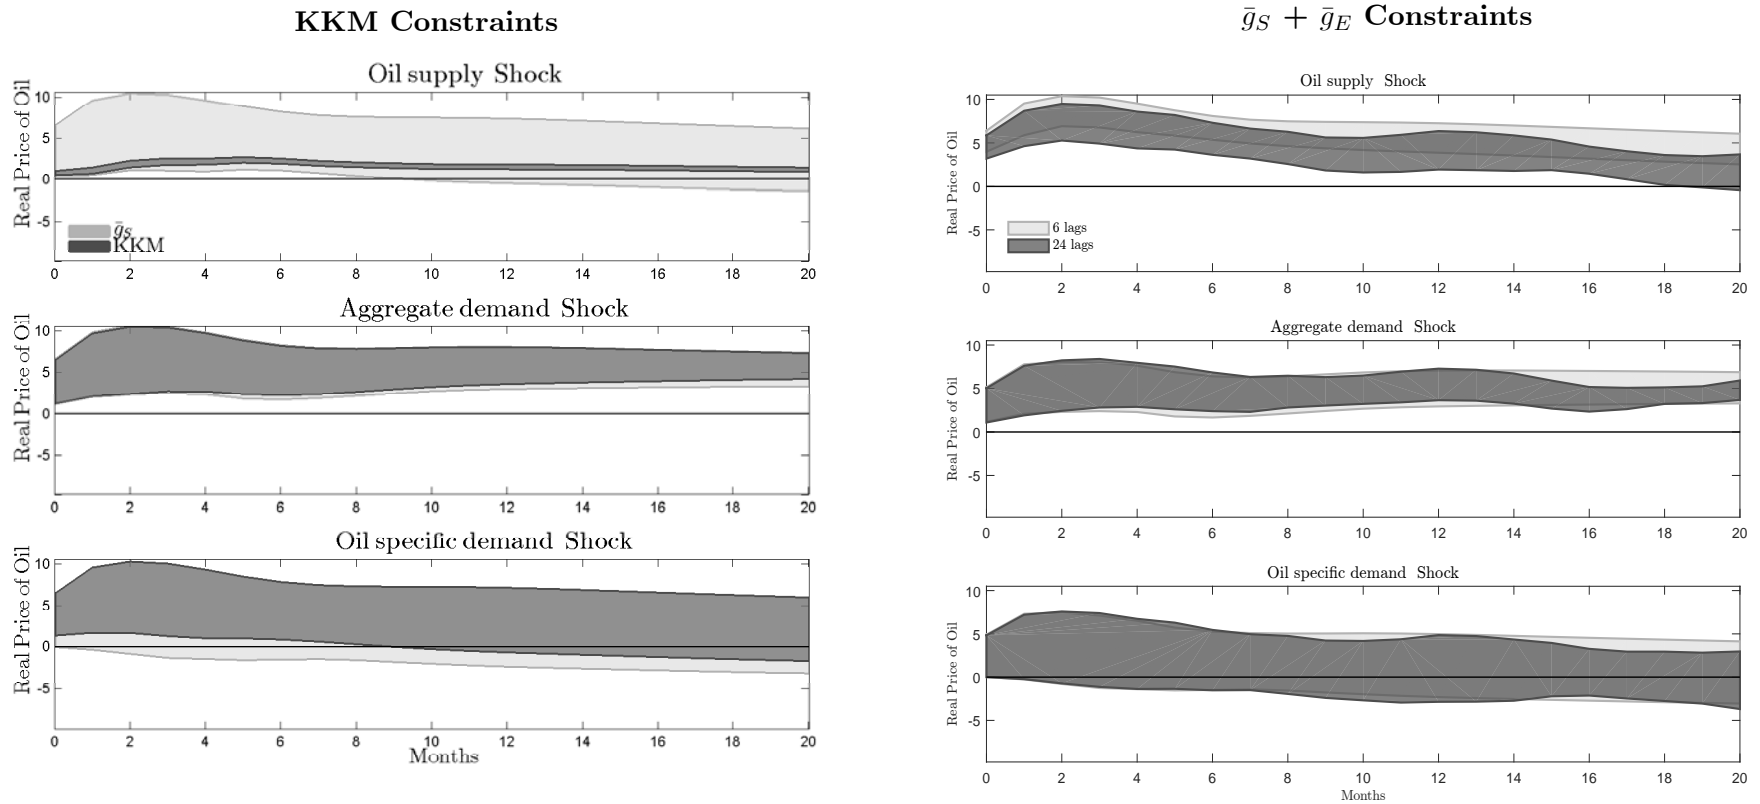
Figure 6: Sampling Simulation



The shaded area reports the 90 percent confidence (CI) interval across 1000 replications. The dark bands are historical set of IRF in the identified set. $CI_{10,g}$ is the CI with all constraints imposed. $CI_{10,gz}$ is the CI with only reduced form covariance restrictions imposed. The 1.5 millions random rotations are used for each replication.

Appendix Figures and Tables

Figure A1: Oil Shock Application 1: Longer Lags



Shaded areas stack the impulse responses for all admissible solutions based on the type of restriction listed in the subplot heading. The sample spans the period 1973:01-2016:06.