NO. 19 • DECEMBER 2012

DALLASFED

Staff Papers

How the Global Perspective Can Help Us Identify Structural Shocks

by Alexander Chudik
and Michael Fidora

Staff Papers *is published by the Federal Reserve Bank of Dallas. The views expressed are those of the authors and should not be attributed to the Federal Reserve Bank of Dallas or the Federal Reserve System. Articles may be reprinted if the source is credited and the Federal Reserve Bank of Dallas is provided a copy of the publication or a URL of the website with the material. For permission to reprint or post an article, email the Public Affairs Department at dal.webmaster@dal.frb.org.*

Staff Papers *is available free of charge by writing the Public Affairs Department, Federal Reserve Bank of Dallas, P.O. Box 655906, Dallas, TX 75265-5906; by fax at 214-922-5268; or by phone at 214-922-5254. This publication is available on the Dallas Fed website, www.dallasfed.org.*

No. 19, December 2012

StaffPAPERS Federal Reserve Bank of Dallas

How the Global Perspective Can Help Us Identify Structural Shocks

Alexander Chudik
Senior Research Economist
Federal Reserve Bank of Dallas

and

Michael Fidora
Economist
European Central Bank

Abstract

This paper argues that global perspective can help us with the identification of structural shocks by utilizing the information on the signs of the responses of individual countries (cross section units). We demonstrate the main idea by means of Monte Carlo experiments and present an empirical application where we look at the effects of oil supply shocks on output and on global exchange rate constellation. Using a large-scale GVAR model of oil prices and the global economy, we find supply shocks tend to have a stronger impact on emerging economies' real output as compared with mature economies, have a negative impact on real growth in oil-exporting economies as well, and tend to cause an appreciation (depreciation) of oil exporters' (oil importers') real exchange rates but also lead to an appreciation of the U.S. dollar. This paper also illustrates some pitfalls with the existing measures to summarize the available information on structural shocks identified using sign restrictions when the dimension of the model is large (as it is in the case of global models).

JEL codes: C32, F41, F47
Keywords: Identification of shocks, sign restrictions, global VAR, oil shocks

We are grateful to Valerie Grossman and Tyler Atkinson for excellent research assistance. We are also grateful to Jennifer Afflerbach, Gianni Amisano, Marcel Fratzscher, Renée Fry, Anil Kumar, Marco Lombardi, Hashem Pesaran, Roland Straub and Mark Wynne as well as seminar participants at the European Central Bank for their comments and suggestions. All errors are our own responsibility. The views expressed in this article are those of the author and do not necessarily reflect the views of the Federal Reserve Bank of Dallas, the Federal Reserve System, the European Central Bank or the Eurosystem.

Vector autoregressive models (VARs) have become an indispensable tool in macroeconometric modelling given their ability to describe *econometric* reduced-form relationships without the need to a priori impose any economic theory. When it comes to unveiling the underlying structural *economic* relationships, however, some method of identification is needed.

In practice, identification has often been based on economic theory rather than purely statistical considerations. Common identification approaches typically include a recursive ordering of variables (Cholesky decomposition), structural identification by imposing zero-restrictions in the system of linear equations (Blanchard and Watson 1986; Bernanke 1986) or a decomposition in temporary and permanent components (Blanchard and Quah 1989). An alternative way to identify shocks that has over the recent past become increasingly popular is to impose restrictions on the *signs* of structural impulse responses for a given number of periods after the shock. This identification approach has been developed inter alia by Faust (1998), Canova and Pina (1999), Canova and de Nicoló (2002), Uhlig (2005) and Mountford and Uhlig (2009). The basic intuition is that structural shocks can be identified by checking whether the signs of the corresponding impulse responses (IR) are in line with economic theory.

Sign restrictions do not pin down a unique structural model, and it is therefore not surprising that any sign restriction identification procedure is bound to be imperfect. Nevertheless, one would expect that, with an increasing number of sign restrictions imposed, one should obtain a better understanding of the structural shock in question. We argue that global or cross-section dimension offers an intuitive and straightforward way of imposing a large number of sign restrictions to identify shocks that are global in nature—i.e., shocks that affect many cross-section units at the same time. However, no matter how large the dimension of the model is and how many restrictions are imposed, there will typically be an infinite number of structural models that satisfy the sign restrictions, and therefore it is necessary to find a way to summarize the available information.

How to deal with the multiplicity of models turns out to be crucial in the models with a large number of variables. The mainstream literature reports the median of the impulse responses, which is sometimes interpreted as a "consensus" view of the magnitudes of the responses, and quantiles are used to give an impression of the distribution of IR. Fry and Pagan (2007) propose a different way of summarizing available models by selecting the IR that is closest to the median IR in some measure. Both approaches have been criticized by Inoue and Kilian (2012), who instead suggest deriving posterior distribution of the structural impulse response functions and looking at the posterior mode (as the most likely response) and the highest posterior density credible sets.

We argue that the global dimension offers a new way of looking at this problem. To focus on the underlying identification problem in a transparent and easy way, we initially abstract completely from the estimation uncertainty of the reduced-form parameters and look at the summary measure as an estimation problem. We expect that with an increasing cross-section dimension and an increasing number of restrictions imposed, the summary measure gets arbitrarily close to the true structural impulse response. We call this feature consistency. Surprising enough, neither the median IR nor the solution proposed by Fry and Pagan (2007) satisfy this intuitive consistency feature, and we show by means of Monte Carlo experiments that the distance between the structural shock and these two measures increases with the increasing number of restrictions imposed, contrary to what one would expect. The suggestion by Inoue and Kilian (2012) to pick a posterior mode does not necessarily solve our problem either in cases when the mode is not unique. We suggest an alternative way of summarizing the available models, based on a "scaled median" estimate of the structural impulse response, which performs well in our Monte Carlo experiments and gets closer to the true structural shocks with an increase in the number of restrictions imposed.

The paper then implements this approach in a large-scale global VAR (GVAR) model that allows us to explore the global dimension by adding a large number of sign restrictions.[1] Our application of the sign restrictions approach in the context of a GVAR model of oil prices and the global economy shows that oil supply shocks have a stronger impact on emerging economies' real output as compared with mature economies, have a negative impact on real growth in oil-exporting economies as well, and tend to cause an appreciation (depreciation) of oil exporters' (oil importers') real exchange rates but also lead to an appreciation of the U.S. dollar. One possible explanation would be the recycling of oil exporters' increased revenues in U.S. financial markets.

The remainder of the paper is as follows. Section 1 illustrates the main issues by means of Monte Carlo experiments in a very simple setup. Section 2 implements the sign restriction approach applying our alternative summary measure of impulse responses in the context of a GVAR model of the world economy, which given its global dimension allows for imposing a large number of sign restrictions for identifying the

[1] The GVAR approach has been proposed by Pesaran, Schuermann, and Weiner (2004).

effect of oil shocks on the global economy. The final section offers some concluding remarks. The Technical Appendix outlines the identification problem, briefly reviews the sign restriction approach, and discusses the problems of summarizing multiple models when the dimension (i.e., the number of variables) is large.

1. MONTE CARLO EXPERIMENTS

Monte Carlo Design

In this section, we illustrate the main idea by means of Monte Carlo experiments. In particular, we show that introducing a global dimension can turn weak information strong: As the number of variables—and hence the leeway to introduce sign restrictions—increases, our estimate, which we denote as $\widehat{\mathbf{r}}_1^{scaled}$, does converge to the "true" structural impulse response (SIR), while the traditionally used measures do not.

Let the data-generating process be $\mathbf{x}_t = \mathbf{u}_t$, where errors in the $n \times 1$ vector of the reduced-form errors $\mathbf{u}_t$ are generated by the following spatial model,

$$\mathbf{u}_t = \mathbf{D}\mathbf{e}_t,$$

and the vector of structural innovations is $\mathbf{e}_t \sim IIDN(\mathbf{0}, \mathbf{I}_n)$. We abstract from any dynamics, since the number of lagged terms is not pertinent to the identification problem. The matrix of structural shocks $\mathbf{D}$ is generated as follows. The "global" shock, which is to be estimated, is the first structural shock and the first column of $\mathbf{D}$ given by $n \times 1$ vector of ones $\mathbf{d}_1 = (1, 1, ..., 1)'$, and all signs of vector $\mathbf{d}_1$ are assumed to be known in the identification procedure. The lower right $(n-1) \times (n-1)$-dimensional submatrix of matrix $\mathbf{D}$, denoted as $\mathbf{D}_{22}$, is generated as

$$\mathbf{D}_{22} = (\mathbf{I}_{n-1} - \rho \mathbf{S})^{-1},$$

where the spatial matrix $\mathbf{S}$ is given by

$$\mathbf{S} = \begin{pmatrix} 0 & 1 & 0 & & 0 \\ \frac{1}{2} & 0 & \frac{1}{2} & & \\ 0 & \frac{1}{2} & 0 & \ddots & 0 \\ & & \ddots & \ddots & \frac{1}{2} \\ 0 & & 0 & 1 & 0 \end{pmatrix},$$

that is, the process $\mathbf{D}_{22}(e_{2t}, ..., e_{nt})'$ is a bilateral spatial autoregressive process. Dependence of the first reduced error u_{1t} on the remaining structural shocks is generated as $d_{1j} \sim U(-\rho/N, \rho/N)$, where d_{1j} denotes the j-th element of the first row of $\mathbf{D}$. Note that for $\rho = 0$, matrix $\mathbf{D}$ becomes (A-6) discussed in the Appendix, whereas as $\rho \to 1$, the degree of cross-section dependence increases. Two experiments are considered: $\rho = 0$ (benchmark case) and $\rho = 0.4$ (stronger cross-section dependence). Matrix $\mathbf{D}$ is generated at the beginning of the experiment and $R = 2000$ replications are carried out. The structural IR vector at the time of impact is estimated in the following three different ways:

1. As a traditional median IR, denoted as $\widehat{\mathbf{r}}_1^{med} = med\{\mathbf{B}\boldsymbol{\alpha}, \boldsymbol{\alpha} \in \mathcal{S}\}$, where $\mathbf{B}\mathbf{B}'$ is the Cholesky decomposition of the known covariance matrix $\boldsymbol{\Sigma}$ and $\mathcal{S}$ is the set of successful draws from the unit circle that satisfy the sign restrictions.[2]
2. As an IR that is closest to the median IR (as proposed by Fry and Pagan 2007), denoted as $\widehat{\mathbf{r}}_1^{FP} = \mathbf{B}\boldsymbol{\alpha}_m$, where $\boldsymbol{\alpha}_m = \arg\min_{\boldsymbol{\alpha} \in \mathcal{S}} \left\| \mathbf{B}\boldsymbol{\alpha} - \widehat{\mathbf{r}}_1^{med} \right\|$.
3. As a scaled estimate $\widehat{\mathbf{r}}_1^{scaled} = \mathbf{B}\widehat{\boldsymbol{\alpha}}$, where $\widehat{\boldsymbol{\alpha}} = \frac{med(\mathcal{S})}{\|med(\mathcal{S})\|}$.

The covariance matrix $\boldsymbol{\Sigma}$ is taken to be known (i.e., we abstract from the estimation uncertainty). Matrix $\mathbf{D}$ is treated as unknown, except for the signs of $\mathbf{d}_1$. The vector $\boldsymbol{\alpha}$ is generated randomly as $\boldsymbol{\alpha} = \boldsymbol{\eta}/\|\boldsymbol{\eta}\|$,where $\boldsymbol{\eta} \sim N(0, \mathbf{I}_n)$. The number of draws that satisfy the sign restrictions are chosen to be $s \in \{100, 500, 1000, 5000\}$.

Monte Carlo Findings

To provide a first impression of the performance of the summary measures under inspection, Figures 1 and 2 plot the histograms of the errors of the estimated structural contemporaneous impulse response

[2] The decomposition of $\boldsymbol{\Sigma}$ is arbitrary and other decompositions should yield similar results since no other information besides the signs of responses is imposed; see, for instance, Uhlig (2005).

vector $\mathbf{r}_1 = \mathbf{d}_1$ for 2000 repetitions of the benchmark experiment ($\rho = 0$) and two choices for the dimension of the model, $(n, s) \in \{(5, 100); (20, 1000)\}$, where recall n is the number of variables and s is the number of successful draws in the set $\mathcal{S}$.[3] Two different metrics for computing the error as the distance between estimated structural impulse responses (SIRs) and true SIRs are considered: the maximum absolute row sum norm, denoted as $\|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\|_r$, and the euclidean norm, denoted as $\|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\|$.

As can be inferred from Figure 1, which plots the absolute frequency of the errors computed using Euclidian norm, both the median impulse response as well as the measure suggested by Fry and Pagan (2007) perform less well than the scaled response as proposed here.

Figure 1: Error Histograms of Estimated Impulse Responses for Small Benchmark Model

($\rho = 0$, $n = 5$ variables, $s = 100$ draws)

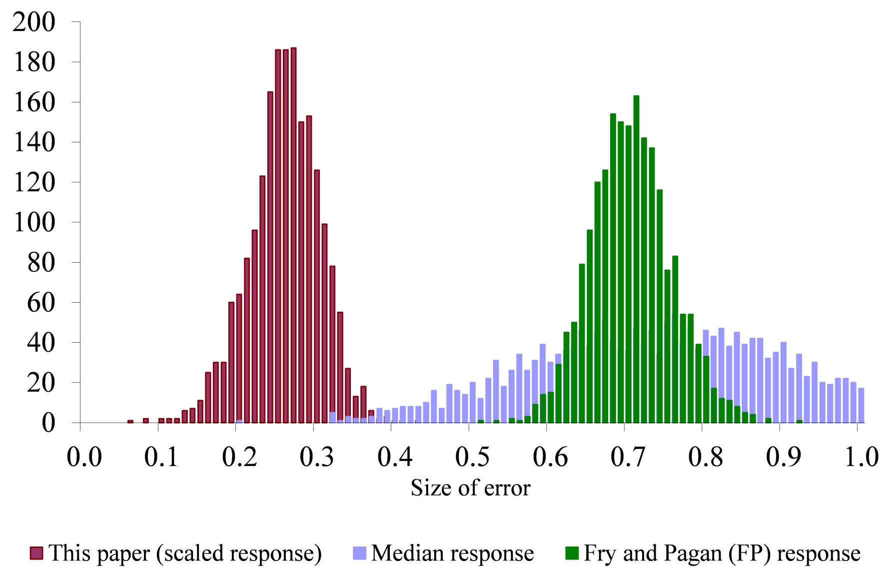

NOTE: The estimator whose distribution is centered closer to zero is preferable.

Furthermore, comparing the errors obtained in the relatively smaller system, as shown in Figure 1, with the errors obtained in a relatively larger system embedding more information on the true signs of the IRs, as shown in Figure 2, shows that the error in estimating the SIRs actually increases in n and hence the number of restrictions imposed when using the conventional summary measures. The scaled response estimator $\widehat{\mathbf{r}}_1^{scaled}$, however, improves with rising n, and hence fulfills the intuitive consistency property outlined in the Technical Appendix. The latter result is robust to the degree of cross-section dependence.[4] The reason behind the increase in the error of estimation when using median response or FP response is described in the Appendix. Intuition is that medians in our setting behave similarly to averaging and averaging across the set of admissible impulse responses does not yield a proper impulse response. As the dimension of the model increases, this problem becomes more serious (as evidenced in Figure 2).

[3] For the sake of clarity of exposition and to be compatible with the notations in the Technical Appendix, we refer to the contemporaneous impulse response vector as $\mathbf{r}_1$, whereas we refer to the structural shocks as $\mathbf{d}_1$.

[4] This is evidenced by Figure 3 reported in the working paper version of this paper; see Chudik and Fidora (2011).

Figure 2: Error Histograms of Estimated Impulse Responses for Large Benchmark Model

($\rho = 0, n = 20$ variables, $s = 1000$ draws)

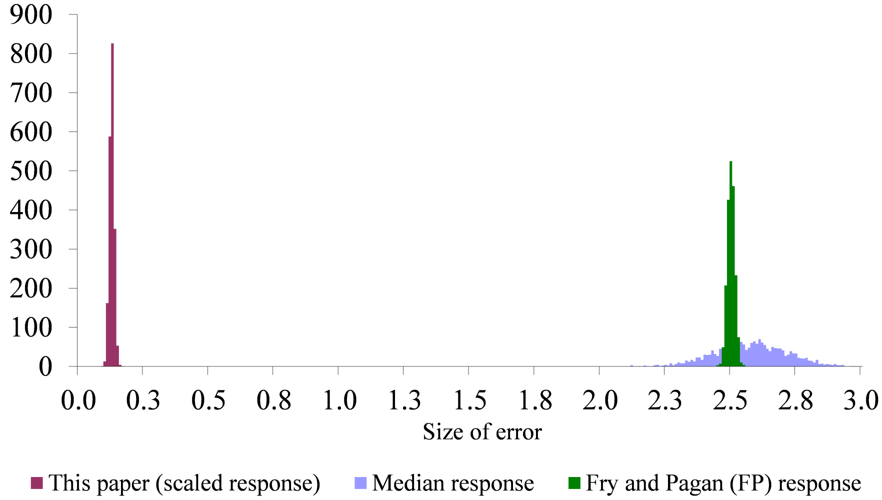

NOTE: The estimator whose distribution is centered closer to zero is preferable.

The preliminary findings from inspection of the error histograms are confirmed when considering a wide range of alternative values for the model dimension. Tables 1 and 2 report the average error of the individual estimates of the structural impulse response vector for $(n, s) \in \{(5, 10, 20, 50, 100); (100, 500, 1000, 5000)\}$ for both the benchmark experiment ($\rho = 0$) as well as the case of higher cross-section dependence ($\rho = 0.4$). It can be inferred from Table 1 that the median response ($\widehat{\mathbf{r}}_1^{med}$) as well as the response which is closest to the median response ($\widehat{\mathbf{r}}_1^{FP}$) do not provide a good estimate of the structural impulse response when n is large. First, performance does not improve in the number of models that satisfy sign restrictions. Second, the estimated responses become even worse as the number of variables and hence the number of sign restrictions increases. These results confirm that even if a large amount of sign restrictions is imposed, the conventional summary measures $\widehat{\mathbf{r}}_1^{med}$ and $\widehat{\mathbf{r}}_1^{FP}$ can be highly misleading when interpreted as a "consensus" view. In particular, both errors $\left\|\widehat{\mathbf{r}}_1^{FP} - \mathbf{r}_1\right\|$ and $\left\|\widehat{\mathbf{r}}_1^{med} - \mathbf{r}_1\right\|$ seem to diverge to infinity with n, and therefore the estimation error is arbitrarily large for n large enough. On the other hand, the scaled response $\widehat{\mathbf{r}}_1^{scaled}$, clearly gets closer to the true response as n and the number of draws increase. Finally, the results also show that introducing a "global dimension" helps gain accuracy: for a system of some tens of variables, $\widehat{\mathbf{r}}_1^{scaled}$ quickly approaches the "true" impulse response keeping the number of successful draws large enough.

Table 1: Monte Carlo Results for Benchmark Case ($\rho = 0$): Error in Estimating Structural Vector r_1

Metric:	Row norm, $\|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\|_r$				Euclidean norm, $\|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\|$			
	Median response $\widehat{\mathbf{r}}_1^{med}$							
n/s	**100**	**500**	**1000**	**5000**	**100**	**500**	**1000**	**5000**
5	0.42	0.41	0.41	0.41	0.70	0.69	0.69	0.69
10	0.52	0.51	0.51	0.51	1.43	1.43	1.43	1.43
20	0.61	0.59	0.59	0.59	2.51	2.50	2.50	2.50
50	0.72	0.69	0.69	0.69	4.78	4.78	4.78	4.78
100	0.78	0.76	0.76	0.75	7.47	7.47	7.47	7.47
	Response closest to median response $\widehat{\mathbf{r}}_1^{FP}$							
5	0.57	0.49	0.45	0.38	0.74	0.68	0.66	0.62
10	0.77	0.76	0.75	0.73	1.52	1.50	1.49	1.47
20	0.87	0.85	0.84	0.83	2.59	2.58	2.58	2.57
50	0.94	0.94	0.94	0.93	4.84	4.84	4.84	4.84
100	0.97	0.96	0.96	0.96	7.52	7.52	7.52	7.52
	Response corresponding to the scaled $med(\mathcal{S})$, $\widehat{\mathbf{r}}_1^{scaled}$							
5	0.17	0.14	0.13	0.13	0.26	0.24	0.24	0.23
10	0.13	0.10	0.10	0.10	0.24	0.18	0.17	0.17
20	0.10	0.06	0.06	0.05	0.28	0.15	0.13	0.11
50	0.12	0.05	0.04	0.02	0.56	0.18	0.12	0.08
100	0.15	0.05	0.03	0.02	1.11	0.30	0.18	0.07

NOTES: n is the number of variables (equal to the number of sign restrictions imposed) and s is the chosen number of successful draws. Estimator with smaller average error is preferable. See Section 1 for detailed description of Monte Carlo design.

Table 2: Monte Carlo Results for High Cross-Section Dependence ($\rho = 0.4$): Error in Estimating Structural Vector r_1

Metric:	Row norm, $\|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\|_r$				Euclidean norm, $\|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\|$			
	Median response $\widehat{\mathbf{r}}_1^{med}$							
n/s	**100**	**500**	**1000**	**5000**	**100**	**500**	**1000**	**5000**
5	0.49	0.49	0.49	0.49	0.62	0.62	0.61	0.61
10	0.56	0.56	0.56	0.56	1.34	1.34	1.34	1.34
20	0.61	0.60	0.60	0.60	2.41	2.40	2.40	2.40
50	0.70	0.68	0.68	0.68	4.64	4.64	4.64	4.64
100	0.76	0.74	0.74	0.74	7.29	7.29	7.29	7.29
	Response closest to median response $\widehat{\mathbf{r}}_1^{FP}$							
5	0.48	0.40	0.38	0.35	0.64	0.55	0.52	0.49
10	0.76	0.74	0.73	0.69	1.43	1.39	1.38	1.35
20	0.86	0.85	0.84	0.84	2.50	2.48	2.48	2.47
50	0.93	0.93	0.92	0.92	4.72	4.71	4.70	4.70
100	0.96	0.96	0.96	0.95	7.35	7.35	7.35	7.34
	Response corresponding to the scaled $med(\mathcal{S})$, $\widehat{\mathbf{r}}_1^{scaled}$							
5	0.33	0.31	0.31	0.30	0.56	0.57	0.57	0.57
10	0.22	0.19	0.18	0.17	0.44	0.44	0.44	0.44
20	0.15	0.12	0.11	0.11	0.34	0.31	0.31	0.31
50	0.11	0.06	0.06	0.04	0.47	0.21	0.19	0.19
100	0.14	0.05	0.04	0.03	0.94	0.25	0.17	0.13

NOTES: n is the number of variables (equal to the number of sign restrictions imposed) and s is the chosen number of successful draws. Estimator with smaller average error is preferable. See Section 1 for detailed description of Monte Carlo design.

2. IDENTIFICATION OF OIL SUPPLY SHOCKS IN A GVAR MODEL

The Basic GVAR Setup

In this section, we implement the sign restriction approach in the context of a large-scale model of the global economy that allows us to impose a large number of sign restrictions. The application illustrates how our scaled median approach can be implemented in a straightforward way in the context of a GVAR and compares results obtained from both the traditional summary measure of impulse responses as well as the proposed scaled median summary measure. Overall, the observed patterns of impulse responses tend to be qualitatively similar, although the scaled median approach systematically yields quantitatively larger point estimates of impulse responses, as we will explain below.

We follow Pesaran, Schuermann, and Weiner (2004) and estimate the following standard country-specific VARX* (p, q) models:

$$\mathbf{\Phi}_{ii}\left(L, p_i\right)\mathbf{x}_{it} = \mathbf{a}_{i0} + \mathbf{a}_{i1}t + \mathbf{\Lambda}_i\left(L, q_i\right)\mathbf{x}_{it}^{*} + \mathbf{u}_{it}, \tag{1}$$

where $\mathbf{x}_{it}$ denotes a $k_i \times 1$ vector of macroeconomic variables belonging to country $i \in \{1, .., N\}$ with N being the number of countries, $\mathbf{x}_{it}^{*} = \sum_{j=1}^{N} w_{ij}\mathbf{x}_{jt}$ denotes country-specific foreign variables (foreign cross-section averages), $\{w_{ij}\}$ are the foreign trade weights, $\mathbf{\Phi}_{ii}\left(L, p_i\right) = \sum_{\ell=0}^{p_i} \mathbf{\Phi}_{ii,\ell}L^{\ell}$ and $\mathbf{\Lambda}_i\left(L, q_i\right) = \sum_{\ell=0}^{q_i} \mathbf{\Lambda}_{i,\ell}L^{\ell}$ are matrix polynomials of unknown coefficients, L is the lag operator, and error terms, collected in a $k \times 1$ vector $\mathbf{u}_t = \left(\mathbf{u}_{1t}', \mathbf{u}_{2t}', ..., \mathbf{u}_{Nt}'\right)'$ with $k = \sum_{i=1}^{N} k_i$ are allowed to be cross-sectionally weakly dependent.[5]

Once estimated on a country-by-country basis, individual VARX* models for $i = 1, .., N$, can be stacked together and solved as one system:

$$\mathbf{G}\left(L, p\right)\mathbf{x}_t = \mathbf{a}_0 + \mathbf{a}_1 t + \mathbf{u}_t, \tag{2}$$

where $\mathbf{a}_{\ell} = \left(\mathbf{a}_{\ell 1}', ..., \mathbf{a}_{\ell N}'\right)'$ for $\ell = 0, 1$, and the construction of the matrix polynomial $\mathbf{G}\left(L, p\right)$ is in detail explained in Pesaran, Schuermann, and Weiner (2004). GVAR model (2) can be used for impulse response or persistence profile analysis in the usual manner.

Data and Sign Restrictions

We set $p_i = p = 2$ and $q_i = q = 1$ and estimate a GVAR in first differences with quarterly data over the period 1979:Q3–2003:Q3 for $N = 26$ countries.[6] We include four country-specific variables —real output,

[5] See Chudik, Pesaran, and Tosetti (2011) for a definition of strong and weak cross-section dependence.

[6] The sample includes Argentina, Australia, Brazil, Canada, Chile, China, euro area, India, Indonesia, Japan, Korea, Malaysia, Mexico, New Zealand, Norway, Peru, Philippines, Saudi Arabia, Singapore, South Africa, Sweden, Switzerland,

inflation, a short-term interest rate, and the real exchange rate—as well as real oil prices as an observable global factor. All variables except the real exchange rates are taken from Dées et al. (2007).[7] Real effective exchange rates and the foreign trade weights are taken from Bussière, Chudik, and Sestieri (2012).

To identify oil supply shocks we rely on a simple identification scheme that allows us to discriminate oil supply shocks from a large set of alternative shocks. In particular, we require *negative* oil supply shocks to be contemporaneously associated with (i) a decrease in real output across all oil importers and (ii) an increase in real oil prices. In all, we impose 21 contemporaneous sign restrictions. We do not impose any restrictions on real output in countries that have been significant oil exporters over the sample period.[8] To the extent that no other economically meaningful shocks are able to produce a negative correlation between real output and real oil prices across *all* oil-importing economies, this identification scheme identifies oil supply shocks.

Empirical Results

In the following, we analyze the effect of negative oil supply shocks on (i) real output in different regions of the global economy, including oil-exporting economies themselves, as well as on (ii) global exchange rate configurations.

Figure 3 gives a general overview of the reaction of output to a one-standard-deviation shock to oil supply. On average, mature economies—including the United States and the euro area—tend to record a decline in real output by between 0.5 percent and 0.75 percent cumulated over the four quarters following a one-standard-deviation oil supply shock (that causes oil prices to increase by 2.4 percent). Emerging economies—both in Latin America and Asia—record somewhat higher declines in growth of, on average, between 1 percent and 1.5 percent. The stronger effect on emerging economies could be a reflection of higher energy intensity of production in these countries, on the one hand, and dependence on external demand from mature economies, on the other hand. In this respect, China stands out as a notable exception with a relatively modest reaction of output to the oil supply shock, possibly reflecting the fact that, despite high-energy intensity, a large part of energy demand is met domestically. Interestingly, though not largely unexpectedly, real output also declines across all oil-exporting economies in our sample—including Saudi Arabia and Norway, although to a lesser extent—despite the reaction of real output for these economies being unrestricted in the identification procedure.

Figure 5 reported in the working paper version of this paper, Chudik and Fidora (2011), provides further detail and compares the impulse response based on the scaled median approach to impulse responses computed as in Uhlig (2005). A main finding from this exercise is that—while the patterns of impulse responses are qualitatively comparable—our scaled median impulse response tends to systematically yield larger point estimates of the shocks under inspection. The quantitative difference is sizeable as impulse responses based on the scaled median tend to be about three times larger than those based on the traditional measure. This is because the traditional measure is based on a median value of $\boldsymbol{\alpha}$ that—at least in the case of a single shock—is bound to lie *within* the unit circle (and therefore does not belong to the set of impulse responses), while our scaling approach avoids this caveat (and $\widehat{\mathbf{r}}_1^{scaled}$ belongs to the set of IRs).

Finally, results for the real exchange rate are summarized in Figure 4, and the complete set of responses is reported in Figure 7 of the working paper version, Chudik and Fidora (2011). As expected, oil importers' real exchange rates mostly tend to depreciate in response to the negative terms of trade shock by on average between 0.5 percent and 1.5 percent, whereas the large oil exporters' appreciate by up to 1 percent. Maybe surprisingly, the United States appreciates by 1.3 percent. One rationalization for this unexpected result could be recycling of oil exporters' increased revenues in U.S. financial markets. As a final note, Figure 7 reported in the working paper shows that, as in the case of real output before, our measure yields similar but systematically larger point estimates of impulse responses, confirming our findings above.

3. CONCLUDING REMARKS

Identification of structural VARs by means of sign restrictions has become increasingly popular in applied econometrics over the recent past. Maybe surprisingly, the performance of identification schemes using sign restriction has only received limited attention. Sign restrictions do not pin down a unique structural model, and it is therefore not surprising that any sign restriction identification procedure is bound to be imperfect.

Thailand, Turkey, United Kingdom, United States.

[7]Updated version of this database is available online on professor M.H. Pesaran's website http://www.econ.cam.ac.uk/faculty/pesaran/.

[8]We take these countries to be Saudi Arabia, Norway, Indonesia, Mexico and the United Kingdom.

Figure 3: Output Decline in Response to One-Standard-Deviation Negative Oil Supply Shock

(Four Quarters Cumulated in Percent)

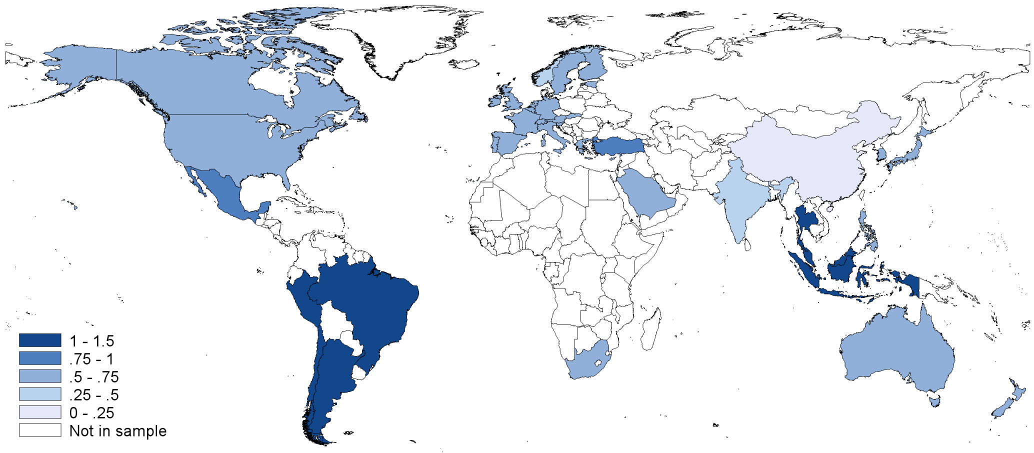

Figure 4: Appreciation of Real Exchange Rate

(Four Quarters Cumulated in Percent)

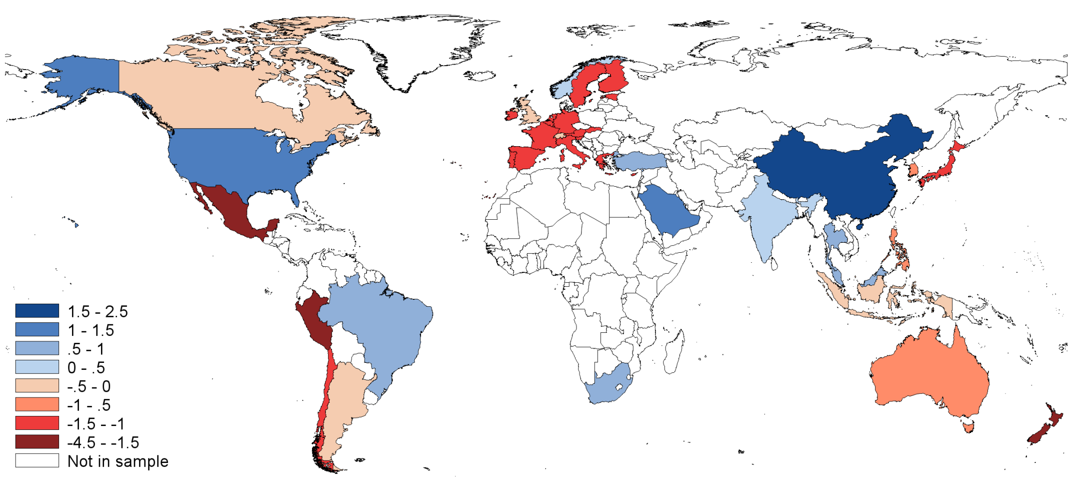

Nevertheless, one would expect that, with an increasing number of sign restrictions imposed, one should obtain a better understanding of the structural shock in question. The global or cross-section dimension offers an intuitive and straightforward way of imposing a large number of sign restrictions to identify shocks that are global in nature—i.e., shocks that affect many cross-section units at the same time.

We implement the sign restriction approach in the context of a global VAR model of the world economy, which, given its global dimension, allows for imposing a large number of sign restrictions, and we identify the effect of oil shocks on the global economy. Our results suggest that negative oil supply shocks (i) have a stronger impact on emerging economies' real output as compared with mature economies, (ii) have a negative impact on real growth in oil-exporting economies as well, and (iii) tend to cause an appreciation (depreciation) of oil exporters' (oil importers') real exchange rates but also lead to an appreciation of the U.S. dollar. One possible explanation would be recycling of oil exporters' increased revenues in U.S. financial markets.

REFERENCES

Bernanke, Ben S. (1986), "Alternative Explanations of the Money-Income Correlation," *Carnegie-Rochester Conference Series on Public Policy* 25 (1): 49–99.

Blanchard, Olivier J., and Danny Quah (1989), "The Dynamic Effects of Aggregate Supply and Demand Disturbances," *American Economic Review* 79 (4): 655–73.

Blanchard, Olivier J., and Mark W. Watson (1986), "Are Business Cycles All Alike?" in *The American Business Cycle: Continuity and Change* (National Bureau of Economic Research), 123–80.

Bussière, M., A. Chudik, and G. Sestieri (2012), "Modelling Global Trade Flows: Results from a GVAR Model," Globalization and Monetary Policy Institute Working Paper no. 119 (Federal Reserve Bank of Dallas, June).

Canova, Fabio, and Gianni de Nicoló (2002), "Monetary Disturbances Matter for Business Fluctuations in the G-7," *Journal of Monetary Economics* 49: 1131–59.

Canova, Fabio, and Joaquim Pina (1999), "Monetary Policy Misspecification in VAR Models," Centre for Economic Policy Research Discussion Paper no. 2333.

Chudik, A., M. H. Pesaran, and E. Tosetti (2011), "Weak and Strong Cross Section Dependence and Estimation of Large Panels," *Econometrics Journal* 14 (1): C45–C90.

Chudik, Alexander, and Michael Fidora (2011), "Using the Global Dimension to Identify Shocks with Sign Restrictions," ECB Working Paper no. 1318.

Dées, S., F. di Mauro, M. H. Pesaran, and L. V. Smith (2007), "Exploring the International Linkages of the Euro Area: A Global VAR Analysis," *Journal of Applied Econometrics* 22 (1): 1–38.

Faust, Jon (1998), "The Robustness of Identified VAR Conclusions About Money," *Carnegie-Rochester Conference Series in Public Policy* 49 (1): 207–44.

Fry, Renée, and Adrian Pagan (2007), "Some Issues in Using Sign Restrictions for Identifying Structural VARs," National Centre for Econometric Research Working Paper no. 14.

Inoue, Atsushi, and Lutz Kilian (2012), "Inference on Impulse Response Functions in Structural VAR Models" (unpublished paper, August 15, 2012).

Mountford, Andrew, and Harald Uhlig (2009), "What Are the Effects of Fiscal Policy Shocks?" *Journal of Applied Econometrics* 24 (6): 960–92.

Pesaran, M. H., Til Schuermann, and Scott M. Weiner (2004), "Modelling Regional Interdependencies Using A Global Error-Correcting Macroeconometric Model," *Journal of Business and Economics Statistics* 22: 129–62.

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.

APPENDIX A: Technical Appendix

Identification of Shocks in a Reduced-Form VAR Model

Suppose that n endogenous variables collected in a vector $\mathbf{x}_t = (x_{1t}, ..., x_{nt})'$ are generated from the following structural VAR(1) model,

$$\mathbf{A}_0\mathbf{x}_t = \mathbf{A}_1\mathbf{x}_{t-1} + \mathbf{e}_t, \tag{A-1}$$

where one lag is assumed for simplicity of exposition, $\mathbf{A}_0$ and $\mathbf{A}_1$ are $n \times n$ matrices of unknown structural coefficients, and $\mathbf{e}_t$ is a vector of orthogonal structural innovations with individual elements e_{it}, for $i = 1, .., N$, having unit variance. Assuming the matrix $\mathbf{A}_0$ is invertible, the reduced form of the structural VAR model (A-1) is:

$$\mathbf{x}_t = \mathbf{\Phi}\mathbf{x}_{t-1} + \mathbf{u}_t, \tag{A-2}$$

where the vector of reduced form errors is given by the "spatial" model $\mathbf{u}_t = \mathbf{D}\mathbf{e}_t$, $\mathbf{D} = \mathbf{A}_0^{-1}$, and $\mathbf{\Phi} = \mathbf{A}_0^{-1}\mathbf{A}_1$.

There is not much disagreement on how to estimate reduced-form VARs, namely the coefficient matrix $\mathbf{\Phi}$ and the covariance matrix of errors, denoted as $\mathbf{\Sigma} = E\left(\mathbf{u}_t\mathbf{u}_t'\right) = \mathbf{D}\mathbf{D}'$. The identification problem is how to decompose reduced-form errors into economically meaningful shocks, i.e., to construct $\mathbf{D}$. The decomposition $\mathbf{\Sigma} = \mathbf{D}\mathbf{D}'$ is not a unique decomposition of the covariance matrix $\mathbf{\Sigma}$. In particular, for any $n \times n$ orthogonal matrix $\mathbf{Q}$, we have $\mathbf{\Sigma} = \left(\mathbf{D}\mathbf{Q}'\right)\left(\mathbf{Q}\mathbf{D}'\right)$ and hence an infinite number of candidate matrices $\mathbf{B}$, where $\mathbf{Q}'\mathbf{Q} = \mathbf{I}$ (orthogonality conditions) and $\mathbf{B} = \mathbf{D}\mathbf{Q}'$. Orthogonality of $\mathbf{Q}$ implies $n(n+1)/2$ restrictions on the elements of $\mathbf{Q}$, leaving thus $n(n-1)/2$ free parameters in constructing the matrix $\mathbf{Q}$. Thus, $n(n-1)/2$ restrictions need to be imposed for exact identification.

Structural shocks in the structural VAR model (A-1) are given by the individual column vectors of matrix $\mathbf{D}$. Let $r_{ij\ell}$ denote response of variable x_{jt} to the structural shock i in the period ℓ after the shock, that is:

$$r_{ij\ell} = \mathbf{s}_j'\mathbf{\Phi}^{\ell}\mathbf{d}_i, \text{ for } \ell \geq 0 \text{ and } i, j \in \{1, .., n\}, \tag{A-3}$$

where $\mathbf{s}_j$ is $n \times 1$ selection vector that selects the j-th element and $\mathbf{d}_i$ is the j-th column of matrix $\mathbf{D}$. For future reference, we denote the IR vector to shock $\mathbf{g}$ as

$$\mathbf{r}_{\ell}(\mathbf{g}) = \mathbf{\Phi}^{\ell}\mathbf{g}, \text{ for } \ell \geq 0,$$

and the structural IR (SIR) vector to the structural shock $\mathbf{d}_i$, as

$$\mathbf{r}_{i\ell} = \mathbf{\Phi}^{\ell}\mathbf{d}_i, \text{ for } \ell \geq 0 \text{ and } i \in \{1, .., n\}.$$

Sign Restriction Approach in the Simplest Setting

To focus on the issue of identification in the simplest and most transparent way, let us abstract from the estimation uncertainty of the reduced-form parameters (coefficients and covariance matrix of errors) of VAR models. In addition, we assume the absence of any dynamics, since the number of lagged terms is not pertinent to the identification problem. The focus of the exposition is thus on the simplest case where

$$\mathbf{x}_t = \mathbf{u}_t, \tag{A-4}$$

and the covariance matrix of the reduced form errors $\mathbf{\Sigma} = \mathbf{D}\mathbf{D}'$ is assumed to be known, but the matrix of structural shocks $\mathbf{D}$ is unknown. Let us denote the contemporaneous IR of structural shock j on variable i as r_{ij}, for $i, j, \in \{1, .., n\}$. In the absence of any dynamics, it follows that $r_{ij} = d_{ij}$ for any $i, j, \in \{1, .., n\}$.[9]

Suppose we are interested in identifying only the first structural shock $\mathbf{d}_1$ (the first column of $\mathbf{D}$) and some or all signs of the vector of the contemporaneous IR $\mathbf{r}_1 = (r_{11}, r_{21}, ..., r_{N1})'$ are known. For any decomposition of the covariance matrix $\mathbf{\Sigma}$ into the product of matrices $\mathbf{B}\mathbf{B}'$, the potential candidates for the structural shock $\mathbf{d}_1$ are fully characterized by the set of vectors $\mathbf{g}_\alpha = \mathbf{B}\boldsymbol{\alpha}$, where $\|\boldsymbol{\alpha}\| = 1$. Typically there are many different shocks $\mathbf{g}_\alpha$ that satisfy the sign restrictions. Denote $\mathcal{S}$ to be a set of $n \times 1$ vectors $\boldsymbol{\alpha}$ of unit length such that the responses $\mathbf{r}(\mathbf{g}_\alpha)$, for $\boldsymbol{\alpha} \in \mathcal{S}$, satisfy the imposed sign restrictions. A measure that somehow summarizes the available information is then required as one often wishes to know what is a "consensus" view. This measure could be interpreted as a point estimate of the SIR at each point in time following the shock. A common way to summarize the available models in the literature is to compute the median of the IRs that satisfy the sign restrictions, that is

$$\widehat{\mathbf{r}}_1^{med} = med\left\{\mathbf{r}(\mathbf{g}_\alpha), \boldsymbol{\alpha} \in \mathcal{S}\right\}. \tag{A-5}$$

[9] For the sake of clarity of exposition, we continue to refer to the contemporaneous impulse responses as r_{ij}, whereas we refer to the structural shocks as d_{ij}.

Global Dimension to the Rescue of Weak Information

One criticism of the sign restriction approach is that it delivers only a weak (imprecise) identification since it uses only weak information. The fact that weak information delivers weak results is a powerful wisdom, and we do not object to this view. Instead, we argue that introducing a "global dimension" can considerably improve the character of information embodied in the sign restrictions in a way that the resulting identification procedure is arbitrarily good as the number of variables (and the number of sign restrictions) tends to infinity. The following simple example suffices to make this point clear.

Suppose that the structural shocks are given by the following matrix

$$\underset{n\times n}{\mathbf{D}} = \begin{pmatrix} 1 & 0 & 0 & 0 & \cdots & 0 \\ 1 & 1 & 0 & 0 & & 0 \\ 1 & 0 & 1 & 0 & & 0 \\ 1 & 0 & 0 & 1 & & 0 \\ \vdots & & & & \ddots & \\ 1 & 0 & 0 & 0 & \cdots & 1 \end{pmatrix}, \tag{A-6}$$

and the objective is to estimate a "global" shock $\mathbf{d}_1$, using the knowledge about the signs of $\mathbf{d}_1$. As before, consider any decomposition of known matrix $\boldsymbol{\Sigma} = \mathbf{BB}'$. For $\boldsymbol{\alpha}$ uniformly distributed on a unit sphere, the distribution of random vectors $\mathbf{B}\boldsymbol{\alpha}$ and $\mathbf{D}\boldsymbol{\alpha}$ is the same. For fixed n, information embodied in the sign restrictions is imperfect and rather weak if n is relatively small. In particular, the probability that a linear combination of structural shocks $\{\mathbf{d}_2, ..., \mathbf{d}_n\}$ will yield the same signs as the last $n-1$ elements of the structural shock $\mathbf{d}_1$ is exponentially decaying in n:

$$p\left\{\sum_{j=2}^{n} \alpha_j \mathbf{d}_j > 0\right\} = \frac{1}{2^{n-1}}. \tag{A-7}$$

Clearly, there are many models that satisfy the sign restrictions, and equation (A-7) suggests that chances that structural shocks $\{\mathbf{d}_2, ..., \mathbf{d}_n\}$ distort the estimate of $\mathbf{d}_1$ are good for small values of n. As $n \to \infty$, this probability goes to zero at an exponential rate. This provides the basic intuition for the "global dimension" to be able to help with the estimation of global structural shocks or—to be more precise—structural shocks which have unbounded maximum absolute column-sum norm in n.

The Problem of Summarizing Available Models

As Fry and Pagan (2007) argue, the median IR (denoted as $\widehat{\mathbf{r}}_1^{med}$) does not necessarily belong to the space of impulse responses. To illustrate the problem of using the median IR, suppose we have two draws of impulse responses $\{\mathbf{r}(\mathbf{g}_{\alpha 1}), \mathbf{r}(\mathbf{g}_{\alpha 2})\}$ which satisfy the sign restrictions. There is no guarantee that there exists a model where $\widehat{\mathbf{r}}_1^{med} = med\{\mathbf{r}(\mathbf{g}_{\alpha 1}), \mathbf{r}(\mathbf{g}_{\alpha 2})\}$ is a structural impulse response. In particular, the solution for $\boldsymbol{\alpha}$ in system $\widehat{\mathbf{r}}_1^{med} = \mathbf{B}\boldsymbol{\alpha}$ need not lie on the unit circle, and therefore $\widehat{\mathbf{r}}_1^{med}$ need not even belong to the space of IRs.

Since the median is not a monotonic transformation, i.e., $med\{\mathbf{g}_{\alpha 1}, ..., \mathbf{g}_{\alpha s}\} \neq \mathbf{B}{\cdot}med\{\boldsymbol{\alpha}_1, ..., \boldsymbol{\alpha}_s\}$, let us for the moment and for the sake of ease of exposition consider arithmetic averages instead. Averaging is a monotonic transformation and the average of the candidate IR (at the time of impact), denoted as $\widehat{\mathbf{r}}_1^{ave} = [\mathbf{r}(\mathbf{g}_{\alpha 1}) + ... + \mathbf{r}(\mathbf{g}_{\alpha s})]/s$, can be written as $\widehat{\mathbf{r}}_1^{ave} = \mathbf{B}\overline{\boldsymbol{\alpha}} = \mathbf{B}(\boldsymbol{\alpha}_1 + ... + \boldsymbol{\alpha}_s)/s$, where $\overline{\boldsymbol{\alpha}} = s^{-1}\sum_{i=1}^{s}\boldsymbol{\alpha}_i$. Note that triangle inequality implies that $\|\overline{\boldsymbol{\alpha}}\| < 1$ (unless $\boldsymbol{\alpha}_1 = \boldsymbol{\alpha}_2 = ... = \boldsymbol{\alpha}_s$). Therefore, as long as the set of draws $\mathcal{S}$ has different elements, $\widehat{\mathbf{r}}_1^{ave}$ (the average of the candidate IR) *never* belongs to the space of IRs. The more candidate draws for the structural shock one obtains and the higher dimensionality of the model, the smaller $\|\overline{\boldsymbol{\alpha}}\|$ could get. Similar arguments also apply to the median as opposed to averaging. The Monte Carlo experiments using the data-generating process described in Section 1 reveal that the median behaves similarly to the averaging.

This paper looks at the same problem of multiplicity of models from a slightly different perspective than Fry and Pagan (2007). As opposed to summarizing the available information by means of quantiles and medians of the impulse responses, we suggest finding an IR vector, which would be as close as possible to the true (unknown) IR. Ideally we would like to find $\widehat{\boldsymbol{\alpha}}$ on the unit circle that would satisfy

$$\widehat{\boldsymbol{\alpha}} = \arg\min_{\boldsymbol{\alpha}'\boldsymbol{\alpha}=1} \|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\|, \tag{A-8}$$

where $\widehat{\mathbf{r}}_1 = \mathbf{B}\boldsymbol{\alpha}$. However, it is not possible to solve (A-8) without knowing the matrix $\mathbf{D}$, and for this reason, we opt for the following weaker consistency requirement,

$$\lim_{n\to\infty} \|\widehat{\mathbf{r}}_1 - \mathbf{r}_1\| = 0, \tag{A-9}$$

that is, the error in estimating IRs vanishes as $n \to \infty$. General treatment of the problem (A-9) goes beyond the analysis presented here; instead, we propose to compute the impulse response corresponding to the scaled median of draws $\boldsymbol{\alpha} \in \mathcal{S}$, that is,

$$\widehat{\mathbf{r}}_1^{scaled} = \mathbf{B}\widehat{\boldsymbol{\alpha}}, \tag{A-10}$$

where

$$\widehat{\boldsymbol{\alpha}} = \frac{\boldsymbol{\alpha}_{med}}{\|\boldsymbol{\alpha}_{med}\|}, \text{ and } \boldsymbol{\alpha}_{med} = med\left\{\mathcal{S}\right\}.$$

In this paper, we refer to $\widehat{\mathbf{r}}_1^{scaled}$ as a “scaled” estimate of the SIR.