

A Procedure for Combining Zero and Sign Restrictions in a VAR-Identification Scheme *

Alex Haberis[†] and Andrej Sokol[‡]

5 June 2014

Abstract

In this paper we describe a procedure for implementing zero restrictions within the context of a sign restrictions identification scheme for VARs. The procedure introduces an additional step into the algorithm outlined in Fry and Pagan (2011) and Rubio-Ramirez et al. (2006) for implementing sign restrictions. This extra step involves rotating a candidate identification matrix using Givens rotation matrices to introduce zero restrictions. We then check whether the elements of the candidate matrix satisfy the sign restrictions as usual. We illustrate how our procedure works by generating artificial data from the theoretical model of An and Schorfheide (2007), which implies certain restrictions on the impact of its structural shocks on the model's endogenous variables. We exploit our knowledge of that pattern to identify structural shocks from the reduced-form errors of a VAR estimated on the simulated data.

JEL: C32, C51, E12

*We would like to thank Luca Benati, Richard Harrison, Francesca Monti, Roland Meeks, Haroon Mumtaz, Matthias Paustian, Konstantinos Theodoridis, and Tony Yates for useful discussions. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Bank of England.

[†]Bank of England and Centre for Macroeconomics. Email: alex.haberis@bankofengland.co.uk

[‡]Bank of England. Email: andrej.sokol@bankofengland.co.uk

1 Introduction

A central aim of structural vector-autoregression analysis is to uncover economically-interpretable – ‘structural’ – shocks that drive dynamics of macroeconomic variables in the data. Economic theory is widely used to motivate the restrictions that are necessary to identify structural shocks from the reduced-form innovations to a vector autoregression (VAR). In recent years, a growing literature has used sign restrictions on the impulse response functions (IRFs) of a VAR to identify economic shocks in the data. The source of these restrictions can be a formal model. For example, the IRFs of a dynamic stochastic general equilibrium (DSGE) model will imply a pattern of co-movement among the model’s endogenous variables that can be used to inform the sign restrictions on the IRFs of a VAR. Alternatively, the restrictions might be based on a view of the effects of a shock on certain variables that is not derived from a fully-specified model¹.

In either case, as well as having implications about the signs of the responses, economic theory can imply that certain variables do not respond at all to some shocks. This suggests that, in addition to sign restrictions, zero restrictions can be used to identify the effects of economic shocks in the data. Indeed, it may be the case that zero restrictions are *essential* to identify the effects of some shocks. An example of a shock that induces a zero response among economic variables can be found in the New Keynesian literature, where monetary policy can be specified to entirely offset the impact of demand shocks on inflation. Another policy example is the assessment of the size of the ‘balanced-budget’ fiscal multiplier, i.e. the impact on output from a policy that involves exogenously changing government spending and taxes in a way that leaves the primary balance unchanged. The impulse response of the primary balance to such a shock would be zero on impact.²

In this paper, we outline a procedure for implementing zero restrictions in the context of a sign restrictions identification scheme. This extends the methodology outlined in Baumeister and Benati (2013) to allow for multiple zero restrictions. In related work, Arias et al. (2014) also propose an algorithm to implement zero restrictions on the IRFs of a VAR. The difference between their procedure and ours is that we use a combination of Givens and Householder transformations to generate candidate structures that satisfy the zero and sign restrictions, whereas Arias et al. (2014) use only Householder transformations, although with additional orthogonality restrictions, to deliver the zero restrictions. In this paper we do not attempt a comparison of the relative performance of the two algorithms, which we leave for future research.

We illustrate our procedure by applying it to the simple New Keynesian model described in An and Schorfheide (2007). The rational expectations equilibrium of this model has a first-order VAR representation so it is, in principle, possible to obtain empirical estimates of the model’s IRFs using a VAR estimated on data for the model’s observable variables. However, the model’s specification of monetary policy implies inflation does not respond at all to government spending shocks. As a result, there are zero responses in the model’s IRFs, and it would not be possible to identify the government spending shock using only sign restrictions – zero

¹Fry and Pagan (2011) review papers in the sign restrictions literature and distinguish between studies in which the sign restrictions come from a formal model or otherwise.

²Though not necessarily thereafter, depending on the implementation of the policy.

restrictions are essential to identify this shock. So we simulate data from the model, estimate a VAR on the simulated data for the observable variables, and identify the structural shocks using a combination of sign and zero restrictions implied by the DSGE model. We show that our procedure does well in recovering the correct impulse responses from the original model.

The paper is organised as follows. In Section 2, we illustrate why VAR impulse responses may include zero responses; in Section 3, we outline our algorithm for introducing zero restrictions using an application to the An and Schorfheide (2007) model; in Section 4 we conclude.

2 When zero restrictions on VAR impulse response functions are required

To illustrate how economic theory can provide the motivation for zero restrictions in a VAR-identification scheme, we focus on a special case in which the shocks in an economic model match up with the shocks associated with a VAR. This will only be true when the model has a finite-order VAR representation. In this case, a structural VAR will give accurate estimates of the IRFs to the economic shocks that are of interest.

Our approach for introducing zeros, however, is not restricted to this case. It is possible that the sign and zero restrictions used to identify the VAR are not derived from a well-defined model; Fry and Pagan (2011) discuss papers that use this approach. That said, since in Section 3 we illustrate our procedure using data artificially generated from a model, it is convenient to consider an economic model with a finite-order VAR representation. This allows us to side-step two issues in VAR analysis which, although of practical and theoretical significance, are not the focus of this paper. First, if a model does not have a well-defined VAR representation, a VAR estimated on data for the model’s observable variables may be relatively uninformative about economic shocks one wishes to identify, as discussed by Fernandez-Villaverde et al. (2007). Second, considering a model whose VAR representation is of finite order allows us to avoid issues related to the potential ‘truncation bias’ introduced by estimating a finite-order VAR when the ‘true’ model has a VAR representation of infinite order, as discussed by Ravenna (2007).

2.1 VAR model

The issue of identifying shocks in a VAR can be outlined, without loss of generality, by considering a first-order VAR model:

$$\mathbf{y}_t = \beta \mathbf{y}_{t-1} + \mathbf{e}_t \tag{1}$$

where $(\mathbf{y}_t$ is an $n_y \times 1$) vector of observable variables and \mathbf{e}_t is an $(n_y \times 1)$ vector of Normally-distributed errors, for which $E(\mathbf{e}_t) = \mathbf{0}$, $E(\mathbf{e}_t \mathbf{e}_t') = \Sigma$, and $E(\mathbf{e}_t \mathbf{e}_{t-s}') = \mathbf{0}$ for $s \neq 0$.

To extract a vector of economically-interpretable shocks from the VAR errors, we need an identification scheme. This can be represented by the matrix $\tilde{\mathbf{F}}$, which transforms the VAR errors into an $(n_y \times 1)$ vector of

orthogonal ‘structural’ shocks, ϵ_t , such that $\epsilon_t = \tilde{\mathbf{F}}^{-1}\mathbf{e}_t$. The structural shocks are Normally distributed, with $E(\epsilon_t) = \mathbf{0}$, $E(\epsilon_t\epsilon_t') = \mathbf{I}$, and $E(\epsilon_t\epsilon_{t-s}') = \mathbf{0}$ for $s \neq 0$. Furthermore, the matrix $\tilde{\mathbf{F}}$ is constructed so that the structural shocks are uncorrelated with each other, $E(\epsilon_{i,t}\epsilon_{j,t}) = 0$, and are consistent with a theory about the effects of the economic shocks being estimated. Having constructed a matrix $\tilde{\mathbf{F}}$, the empirical model can take the form of a structural VAR:

$$\mathbf{y}_t = \beta\mathbf{y}_{t-1} + \tilde{\mathbf{F}}\epsilon_t \quad (2)$$

2.2 Theoretical model

The equilibrium of a theoretical model can take the following state-space representation for \mathbf{y}_t :³

$$\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{B}\mathbf{w}_t \quad (3)$$

$$\mathbf{y}_t = \mathbf{C}\mathbf{x}_{t-1} + \mathbf{D}\mathbf{w}_t \quad (4)$$

where \mathbf{y}_t continues to denote an $(n_y \times 1)$ vector of observable variables and \mathbf{x}_t is an $(n_x \times 1)$ vector of (possibly unobserved) state variables. The vector \mathbf{w}_t ($n_w \times 1$) includes both measurement errors and innovations to the economically-interpretable structural shocks (e.g. shocks to technology, preferences, policy, and so on).⁴ They are Normally distributed, with $E(\mathbf{w}_t) = \mathbf{0}$, $E(\mathbf{w}_t\mathbf{w}_t') = \mathbf{I}$, and $E(\mathbf{w}_t\mathbf{w}_{t-s}') = \mathbf{0}$ for $s \neq 0$. The coefficient matrices \mathbf{A} ($n_x \times n_x$), \mathbf{B} ($n_x \times n_w$), \mathbf{C} ($n_y \times n_x$), and \mathbf{D} ($n_y \times n_w$) are functions of the deep parameters of the model: technology, preferences, the persistence of the structural shocks, etc..

Provided certain conditions hold,⁵ economic models of the ABCD form – equations (3) and (4) – have a finite-order VAR representation. This maps innovations to the structural shocks to the observable variables and their lags. In a particular case, which is convenient for our exposition, the VAR representation of the economic model will be first-order:⁶

$$\mathbf{y}_t = \mathbf{C}\mathbf{A}\mathbf{C}^{-1}\mathbf{y}_{t-1} + \mathbf{D}\mathbf{w}_t. \quad (5)$$

2.3 Theory-based identification of the VAR shocks

From equations (2) and (5) it is clear that there is a correspondence between the economic model and the SVAR. Therefore, to identify the SVAR shocks – i.e. to construct the matrix $\tilde{\mathbf{F}}$ – it is possible to draw on the information in \mathbf{D} about the pattern of responses – including zero responses – of the model’s endogenous variables to the economic shocks, \mathbf{w}_t . For instance, restricting $\tilde{\mathbf{F}}$ to match the pattern of signs and zeros in \mathbf{D} would allow the structural shocks in the VAR, ϵ_t , to be interpreted as the economic shocks in the model, \mathbf{w}_t .

³This representation is based on Fernandez-Villaverde et al (2007).

⁴Here we are considering the square case in which the number of shocks is equal to the number of observable variables.

⁵As outlined by Fernandez-Villaverde et al. (2007) and Ravenna (2007).

⁶This holds when $\mathbf{B} = \mathbf{A}\mathbf{C}^{-1}\mathbf{D}$. See Morris (2013) for further details on the conditions under which a DSGE model has a VAR(1) representation.

3 An algorithm for imposing zero and sign restrictions: illustrative example

In this section we outline, using a simple application based on the model of An and Schorfheide (2007), our algorithm for implementing zero and sign restrictions in a VAR-identification scheme.⁷ The aim of the algorithm is to find a set of matrices that satisfy the sign and zero restrictions implied by \mathbf{D} in the model's VAR representation, which we will denote by $\tilde{\mathbf{F}}_m$, for $m = 1, 2, \dots, n_M$. Each matrix $\tilde{\mathbf{F}}_m$ represents an identified *structure* rather than an identified *model*. The number of identified structures, n_M , is chosen by the econometrician.⁸ Generalising Baumeister and Benati (2013), our approach combines the procedure for imposing sign restrictions of Rubio-Ramirez et al. (2006) with the imposition of $n_z \leq \frac{n_y(n_y-1)}{2}$ zero restrictions using an equal number of Givens rotations matrices.⁹ The algorithm involves repeating a few steps many times until the required number of matches, n_M , are found.

Although our procedure is applicable to VAR-shock identification based on zero and sign restrictions in general, the An and Schorfheide (2007) model is useful for expositional purposes insofar as its rational expectations equilibrium has a VAR(1) representation which implies zero as well as sign restrictions. We can, therefore, generate artificial data from the model, estimate a VAR(1) on the simulated data, and then identify the structural shocks using a combination of zero and sign restrictions. Before turning to the algorithm, we briefly describe the model.

3.1 An and Schorfheide (2007) model

3.1.1 Theoretical model

The log-linearised equilibrium conditions and exogenous processes of the An and Schorfheide (2007) model take the following form:

$$\hat{y}_t = \mathbb{E}_t \hat{y}_{t+1} - \frac{1}{\tau} (\hat{r}_t - \mathbb{E}_t \hat{\pi}_{t+1} - \mathbb{E}_t \hat{z}_{t+1}) + \hat{g}_t - \mathbb{E}_t \hat{g}_{t+1} \quad (6)$$

$$\hat{\pi}_t = \beta \mathbb{E}_t \hat{\pi}_{t+1} + \kappa (\hat{y}_t - \hat{g}_t) \quad (7)$$

$$\hat{c}_t = \hat{y}_t - \hat{g}_t \quad (8)$$

$$\hat{r}_t = \rho_r \hat{r}_{t-1} + (1 - \rho_r) (\psi_1 \hat{\pi}_t + \psi_2 (\hat{y}_t - \hat{g}_t)) + \epsilon_{r,t} \quad (9)$$

$$\hat{g}_t = \rho_g \hat{g}_{t-1} + \epsilon_{g,t} \quad (10)$$

$$\hat{z}_t = \rho_z \hat{z}_{t-1} + \epsilon_{z,t} \quad (11)$$

⁷The general algorithm is described in Appendix A.

⁸Fry and Pagan (2011) discuss the problem of structure versus model identification.

⁹This maximum number of zero restrictions corresponds to the number of zeros in a triangular matrix. Indeed, the algorithm described in this section can be used to triangularise a square matrix of size n_y .

where \hat{y}_t is output, $\hat{\pi}_t$ is inflation, \hat{c}_t is consumption, \hat{r}_t is the nominal interest rate, \hat{g}_t is government spending, \hat{z}_t is technology, $\epsilon_{r,t}$ is a monetary policy shock, $\epsilon_{g,t}$ is a government spending shock, and $\epsilon_{z,t}$ is a technology shock. The shocks are normally distributed with means zero and standard deviations σ_r , σ_g , and σ_z respectively. The structure of the model is standard in the New Keynesian literature: equation (6) is the dynamic *IS* curve; equation (7) is the New Keynesian Phillips curve; equation (8) is the aggregate goods market equilibrium; equation (9) is the monetary policy rule; equations (10) and (11) are the exogenous processes for government spending and technology.¹⁰

We adopt the same parameterisation as in An and Schorfheide (2007). The value of the intertemporal elasticity of substitution, τ , is set to 2; the discount factor, β , is 0.9975; the elasticity of demand for intermediate goods, ν^{-1} , is 10; the degree of price stickiness, ϕ , is 53.6797; the steady state inflation rate, $\bar{\pi}$, is 1.008; the weight on inflation in policy rule, ψ_1 , is 1.5; the weight on consumption in policy rule, ψ_2 , is 0.125; the persistence of government spending shock process, ρ_g , is 0.95; the persistence of technology shock process, ρ_z , is 0.9; and the Phillips curve slope, $\kappa \equiv \frac{\tau(1-\nu)}{\nu\bar{\pi}^2\phi}$, is 0.33.

The VAR(1) representation of the model's rational expectations equilibrium is given by a system in which the vector of observable variables is $\mathbf{y}_t \equiv [r_t \ y_t \ \pi_t]'$, and the vector of shocks is $\mathbf{w}_t \equiv [\epsilon_{z,t} \ \epsilon_{g,t} \ \epsilon_{r,t}]'$:¹¹

$$\begin{bmatrix} r_t \\ y_t \\ \pi_t \end{bmatrix} = \begin{bmatrix} 0.7902 & 0 & 0.2535 \\ 0.1944 & 0.95 & -0.4642 \\ 0.1195 & 0 & 0.6242 \end{bmatrix} \begin{bmatrix} r_{t-1} \\ y_{t-1} \\ \pi_{t-1} \end{bmatrix} + \begin{bmatrix} 0.6055 & 0 & 0.6858 \\ 1.4863 & 1 & -1.1011 \\ 1.4909 & 0 & -0.7462 \end{bmatrix} \begin{bmatrix} \epsilon_{z,t} \\ \epsilon_{g,t} \\ \epsilon_{r,t} \end{bmatrix} \quad (12)$$

From the \mathbf{D} matrix in equation (12), it is clear that the model's assumptions imply a set of contemporaneous responses of the observable variables to the shocks that includes zeros. In particular, neither the nominal interest rate nor inflation respond to the government spending shock. This is because, in the model, consumption – and hence inflation via the Phillips curve, and nominal rates via the policy rule – are left unchanged by government spending shocks. This reflects the model's unit government-spending multiplier on output.

3.1.2 VAR estimation and identification scheme

To illustrate our procedure, we estimate a structural VAR(1) for the model's observable variables on data simulated from the theoretical model, equation (12).¹² To identify the economic shocks from the reduced-form VAR errors we use sign and zero restrictions consistent with the \mathbf{D} matrix in equation (12), which are summarised in Table 1. Clearly, only imposing the sign restriction that the government spending shock increases output would be insufficient to identify this shock. And imposing additional sign restrictions would be inconsistent with the theory. Therefore, the zero restrictions are essential to identify the government spending shock in this model.

¹⁰The formulation of the Phillips curve and policy rule are slightly different to common specifications for these equations; nevertheless, it is equivalent since, $y_t - g_t$, which is equal to consumption, is proportional to real marginal cost and the output gap.

¹¹See Morris (2012). See also Komunjer and Ng (2011) for further details on the model's ABCD representation.

¹²The parameter estimates for the VAR estimation are the same as the true data generating process in equation (12) to 4 decimal places.

	Technology shock	Government spending shock	Monetary policy shock
Nominal interest rate	+	0	+
Output	+	+	-
Inflation	+	0	-

Table 1: Sign and zero restrictions for identification scheme

3.2 Implementing the algorithm for zero and sign restrictions

The algorithm for imposing zero and sign restrictions involves iterating through four steps until the desired number of candidate structures, n_m , is found. Before starting the iterative stage of the algorithm, it is necessary to take the Cholesky decomposition, \mathbf{F} , of the estimated covariance matrix of the VAR errors, $\hat{\Sigma}$:

$$\mathbf{F}^{-1} \hat{\Sigma} \mathbf{F}'^{-1} = \mathbf{I}_{n_y}.$$

This defines a set of shocks $\epsilon_t \equiv \mathbf{F}^{-1} \hat{\mathbf{e}}_t$ that are uncorrelated and have unit variance, but do not necessarily satisfy the sign and zero restrictions, where $\hat{\mathbf{e}}_t$ are the estimated VAR errors.

In the iterative part of the algorithm, at each stage of the iteration, we generate orthonormal matrices \mathbf{Q}_G (i.e. $\mathbf{Q}_G \mathbf{Q}'_G = \mathbf{I}_{n_y}$) such that the uncorrelated, unit-variance shocks, $\tilde{\epsilon}_t \equiv (\mathbf{F} \mathbf{Q}_G)^{-1} \mathbf{e}_t \equiv \tilde{\mathbf{F}}^{-1} \mathbf{e}_t$ satisfy the sign and zero restrictions in Table 1, and have the same covariance matrix as the reduced-form errors, $E(\mathbf{F} \mathbf{Q}_G \tilde{\epsilon}_t \tilde{\epsilon}'_t \mathbf{Q}'_G \mathbf{F}') = \mathbf{F} \mathbf{F}' = \Sigma$. In particular, we repeat the following steps until n_M matches have been found:

1. We take the QR decomposition of a random ($n_y \times n_y$) Normal matrix, \mathbf{W} :

$$\mathbf{W} = \mathbf{Q} \mathbf{R}$$

where \mathbf{Q} is an orthogonal matrix (i.e. $\mathbf{Q} \mathbf{Q}' = \mathbf{Q}' \mathbf{Q} = \mathbf{I}_{n_y}$), and \mathbf{R} is a triangular matrix. We now have a matrix $\mathbf{F} \mathbf{Q}$ and some uncorrelated, unit-variance shocks, ϵ_t :

$$\begin{aligned} \hat{\mathbf{e}}_t &= \mathbf{F} \mathbf{Q} \epsilon_t \\ &= \begin{bmatrix} f_{11} & f_{12} & f_{13} \\ f_{21} & f_{22} & f_{23} \\ f_{31} & f_{32} & f_{33} \end{bmatrix} \begin{bmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \\ \epsilon_{3,t} \end{bmatrix} \end{aligned}$$

2. We rotate $\mathbf{F} \mathbf{Q}$ to introduce the zero restrictions, $f_{12} = f_{32} = 0$, by post-multiplying with two Givens rotations matrices:

$$\mathbf{F} \mathbf{Q} \mathbf{G}_1(\theta_1^*) \mathbf{G}_2(\theta_2^*) \equiv \mathbf{F} \mathbf{Q}_G \equiv \tilde{\mathbf{F}}$$

where the Givens rotations are functions of the angles θ_1^* and θ_2^* . These two angles can be found by solving the system of two simultaneous equations obtained by setting the (1, 2) and (3, 2) elements of $\tilde{\mathbf{F}}$ to zero. Writing this out for clarity, starting with the first restriction, we can choose the following Givens matrix to set the element in (1, 2) to zero:

$$\mathbf{G}_1(\theta_1) = \begin{bmatrix} \cos \theta_1 & -\sin \theta_1 & 0 \\ \sin \theta_1 & \cos \theta_1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

For the second zero restriction in (3, 2) we use the following Givens rotation:

$$\mathbf{G}_2(\theta_2) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_2 & -\sin \theta_2 \\ 0 & \sin \theta_2 & \cos \theta_2 \end{bmatrix}$$

We then compute the product:

$$\mathbf{FQ}\mathbf{G}_1(\theta_1)\mathbf{G}_2(\theta_2) = \begin{bmatrix} f_{11} & f_{12} & f_{13} \\ f_{21} & f_{22} & f_{23} \\ f_{31} & f_{32} & f_{33} \end{bmatrix} \begin{bmatrix} \cos \theta_1 & -\sin \theta_1 & 0 \\ \sin \theta_1 & \cos \theta_1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta_2 & -\sin \theta_2 \\ 0 & \sin \theta_2 & \cos \theta_2 \end{bmatrix}$$

Setting the expressions in the elements in (1, 2) and (3, 2) of this product gives the following pair of simultaneous equations, which can be solved for θ_1^* and θ_2^* :

$$\begin{aligned} f_{13} \sin \theta_2 + \cos \theta_2 (f_{12} \cos \theta_1 - f_{11} \sin \theta_1) &= 0 \\ f_{33} \sin \theta_2 + \cos \theta_2 (f_{32} \cos \theta_1 - f_{31} \sin \theta_1) &= 0 \end{aligned}$$

Having solved for the angles θ_1^* and θ_2^* , and used them to construct the Givens rotations, $\mathbf{G}_1(\theta_1^*)$ and $\mathbf{G}_2(\theta_2^*)$, we have an orthonormal matrix, $\tilde{\mathbf{F}}$, that satisfies the zero restrictions.

3. In this step, we check whether the non-zero elements in $\tilde{\mathbf{F}}$ match the sign restrictions in Table 1, and retain the matrix if it does.

4. We go back to step 1 and repeat until n_M matches have been found.

In our application we set n_M to 500. As a result, we have 500 identified structures given by:

$$\tilde{\mathbf{F}}_m \tilde{\epsilon}_t = \mathbf{e}_t.$$

for $m = 1, 2, \dots, 100$, where $\tilde{\epsilon}_t$ can be interpreted as the shocks in the economic model.

3.2.1 Identified shocks

Figure 1 shows the estimated IRFs to the shocks for the 500 identified structures. In the figure, the swathes of blue (plain) lines are the IRFs for each identified structure. The red (circled) lines are the ‘‘median target’’ IRF calculated following Fry and Pagan (2011), i.e. the IRF – generated by one of our m structures – with the smallest average distance from the median responses across the technology, government spending, and monetary policy shocks. Finally, the green (crossed) lines are the IRFs from the model itself. The figure shows the Fry-Pagan medians of the estimated IRFs match well the IRFs from the theoretical model, including the zero responses for the government spending shock on inflation and the nominal interest rate.

3.3 Discussion

As noted above, Arias et al. (2014) also propose a procedure for implementing sign and zero restrictions, as well as surveying earlier contributions. The difference between the algorithm proposed by Arias et al. (2014)




Figure 1: Impulse response functions for shocks identified with sign and zero restrictions

Notes: The swathes of blue (plain) lines are the IRFs for all the draws that match the sign and zero restrictions. The red (circled) lines are the medians calculated as in Fry and Pagan (2011). The green (crossed) lines are the model responses.

and the one described here is that in the former, the zeros are introduced to the Q matrix via additional linear restrictions in the Householder transformation that generates Q . In our procedure, we introduce the zeros in step 2 by rotating the Q matrix using Givens rotation matrices. We intend to compare the relative performance of the two algorithms in future work.

Furthermore, in a Bayesian setting, our iterative steps could be adapted to the ones described by Arias et al. (2014) for drawing from the posterior distribution of the VAR's structural parameters. Indeed, it would involve drawing a matrix $\hat{\Sigma}$ from the reduced-form posterior (e.g. using Gibbs sampling), finding its Cholesky decomposition \mathbf{F} , and then running through steps 1-4 until the desired number of candidate structures for each posterior draw are found. From a theoretical point of view, our algorithm therefore does not suffer from the shortcomings of other approaches to introduce zero restrictions surveyed by Arias et al. (2014).

4 Conclusion

In this paper we have described an algorithm for implementing zero restrictions within the context of a sign-restrictions identification scheme for VARs that exploits the properties of Givens rotation matrices. We show

how the algorithm works with an application to the DSGE model of An and Schorfheide (2007), which has precise implications for the pattern of sign and zero restrictions that would be necessary to identify the structural shocks from the reduced form errors in a VAR estimated on data for its observable variables. We generate artificial data from the model and identified the shocks using our procedure, showing that these can be recovered in a satisfactory manner. In future work, we intend to compare the performance of the algorithm to alternative approaches for introducing zero restrictions, such as using a Cholesky decomposition on an appropriately partitioned matrix and the algorithm of Arias et al. (2014), and also to explore empirical applications.

References

- An, S. and Schorfheide, F. (2007). Bayesian analysis of DSGE models. *Econometric Reviews*, 26(2-4):113–172.
- Arias, J. E., Rubio-Ramrez, J. F., and Waggoner, D. F. (2014). Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications. *Federal Reserve Board International Finance Discussion Papers*.
- Baumeister, C. and Benati, L. (2013). Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound. *International Journal of Central Banking*, 9(2):165–212.
- Fernandez-Villaverde, J., Rubio-Ramirez, J. F., Sargent, T. J., and Watson, M. W. (2007). ABCs (and Ds) of understanding VARs. *American Economic Review*, 97(3):1021–1026.
- Fry, R. and Pagan, A. (2011). Sign restrictions in structural vector autoregressions: A critical review. *Journal of Economic Literature*, 49(4):938–60.
- Komunjer, I. and Ng, S. (2011). Dynamic identification of dynamic stochastic general equilibrium models. *Econometrica*, 79(6):1995–2032.
- Morris, S. (2013). VAR(1) representation of DSGE models. In Giacomini, R., editor, *Advances in Econometrics*, volume 31.
- Ravenna, F. (2007). Vector autoregressions and reduced form representations of dsge models. *Journal of Monetary Economics*, 54(7):2048–2064.
- Rubio-Ramirez, J. F., Waggoner, D., and Zha, T. (2006). Markov-switching structural vector autoregressions: Theory and application. *Computing in Economics and Finance 2006* 69, Society for Computational Economics.

A Algorithm for imposing multiple zero restrictions

A.1 General algorithm

This section describes the generalisation of the algorithm outlined in Section 3.2 in the main text. The objective is to find a set of matrices that satisfy the sign and zero restrictions in the econometrician's VAR-identification scheme, which we will denote by $\tilde{\mathbf{F}}_m$, for $m = 1, 2, \dots, n_M$. The algorithm involves iterating over a number of steps many times until the required number of matches, n_M , are found. Before starting the iterative stage of the algorithm, it is necessary to take the Cholesky decomposition, \mathbf{F} , of the estimated covariance matrix of the VAR errors, $\hat{\Sigma}$:

$$\mathbf{F}^{-1} \hat{\Sigma} \mathbf{F}'^{-1} = \mathbf{I}_{n_y}.$$

This defines a set of shocks $\epsilon_t \equiv \mathbf{F}^{-1} \mathbf{e}_t$ that are uncorrelated and have unit variance, but do not necessarily satisfy the sign and zero restrictions. In the iterative stage of the algorithm, we generate orthogonal matrices \mathbf{Q}_G (i.e. $\mathbf{Q}_G \mathbf{Q}'_G = \mathbf{I}_{n_y}$) such that the uncorrelated, unit-variance shocks, $\tilde{\epsilon}_t \equiv (\mathbf{F} \mathbf{Q}_G)^{-1} \mathbf{e}_t \equiv \tilde{\mathbf{F}}^{-1} \mathbf{e}_t$ satisfy the sign and zero restrictions, and $E(\mathbf{F} \mathbf{Q}_G \tilde{\epsilon}_t \tilde{\epsilon}'_t \mathbf{Q}'_G \mathbf{F}') = \mathbf{F} \mathbf{F}' = \Sigma$. In particular, we repeat the following steps until n_M matches have been found:

1. Take the QR decomposition of an $(n_y \times n_y)$ random Normal matrix \mathbf{W} :

$$\mathbf{W} = \mathbf{Q} \mathbf{R}$$

where \mathbf{Q} is an orthogonal matrix (i.e. $\mathbf{Q} \mathbf{Q}' = \mathbf{Q}' \mathbf{Q} = \mathbf{I}_{n_y}$), and \mathbf{R} is a triangular matrix.

2. Rotate $\mathbf{F} \mathbf{Q}$ to introduce the zero restrictions by post-multiplying by the product of n_z Givens-rotation matrices, where n_z is the number of zero restrictions to be imposed:

$$\tilde{\mathbf{F}} = \mathbf{F} \mathbf{Q} \prod_{k=1}^{n_z} \mathbf{G}_k(\theta_k^*)$$

where $\mathbf{G}_k(\theta_k)$ denotes the Givens rotation matrix parametrised by angle θ_k ¹³. The product of \mathbf{Q} and the Givens-rotation matrices gives the orthogonal matrix $\mathbf{Q}_G \equiv \mathbf{Q} \prod_{k=1}^{n_z} \mathbf{G}_k(\theta_k^*)$. Let $\Omega = \{(m, n) \in \{1, \dots, N\} \times \{1, \dots, N\}\}$ denote the set of coordinate pairs that index the elements of $\mathbf{F} \mathbf{Q}_G$ that are required to satisfy zero restrictions.¹⁴ Then the angles $\theta^* = [\theta_1^*, \theta_2^*, \dots, \theta_{n_z}^*]'$ solve the following system of n_z equations:

$$vec([\mathbf{F} \mathbf{Q}_G]_{\Omega}) = \mathbf{0}$$

That is, the angles are chosen such that the impact of certain shocks on some variables will be zero.

3. Check the signs in $\tilde{\mathbf{F}}$ match up with the sign restrictions and retain the matrix if so.
4. Go back to step 1 and repeat until the n_M matches have been found.

¹³Recall that Givens rotations are also orthogonal, $\mathbf{G}_k \mathbf{G}'_k = \mathbf{G}'_k \mathbf{G}_k = \mathbf{I}_{n_y}$.

¹⁴The Givens matrices are chosen such that the k -th zero restriction always lies on either a row or column where the non-zero elements of G_k are located.

This will result in n_M matrices such that:

$$\begin{aligned} \mathbf{e}_t &= \tilde{\mathbf{F}}_m \tilde{\epsilon}_t \\ &= \mathbf{FQ}_m \prod_{k=1}^{n_z} \mathbf{G}_k^m(\theta_k^*) \tilde{\epsilon}_t. \end{aligned}$$

A.2 Multi-Period Zero Restrictions

A similar algorithm can also be employed to impose zero restrictions in the periods following impact. We explore two extensions: imposing zero restrictions for additional periods beyond impact, and imposing zero restrictions further out in time, without constraining a variable to be zero on impact.

The total number of zero restrictions that can in principle be imposed across all periods still remains $n_y(n_y - 1)/2$. However, the fact that restrictions in successive periods have the same location poses an additional constraint, because there are only $n_y - 1$ Givens matrices that rotate any given column of a square matrix of size n_y .

The algebra is the same for both extensions. Using a notation similar to that in the previous section, our problem now becomes that of finding a set of Givens matrices that deliver a matrix $\tilde{\mathbf{F}}$ that solves

$$\left[\mathbf{HF} \right]_{\Omega} = \mathbf{0} \tag{A.1}$$

where

$$\Omega = \{(m, n) \in \{1, \dots, NP\} \times \{1, \dots, N\}\} \tag{A.2}$$

is the set of coordinate pairs that index the zero restrictions in each of the $p = 1, 2, \dots, P$ periods over which one wishes to impose zero restrictions. \mathbf{H} is a $n_y P \times n_y$ matrix that stacks the propagation mechanism of the VAR in each period and is constructed as follows:

$$\mathbf{H} = \begin{bmatrix} \mathbf{I}_{n_y} \\ \mathbf{JMU} \\ \vdots \\ \mathbf{JM}^{P-1}\mathbf{U} \end{bmatrix}$$

where

$$\mathbf{J} \equiv \begin{bmatrix} \mathbf{I}_{n_y} & \mathbf{0}_{n_y \times (L-1)} \end{bmatrix} \quad \mathbf{M} \equiv \begin{bmatrix} & \beta' \\ \mathbf{I}_{n_y(L-1)} & \mathbf{0}_{(L-1) \times n_y} \end{bmatrix} \quad \mathbf{U} = \begin{bmatrix} \mathbf{I}_{n_y} \\ \mathbf{0}_{n_y \times (L-1)} \end{bmatrix}$$

This allows for the imposing multi-period zero restrictions, both on impact and starting from an arbitrary period.

A.3 A More Efficient but Less General Variant of the Algorithm

The need to solve a system of nonlinear equation can be computationally costly, especially when the algorithm is coupled with a sign restrictions routine that requires to impose the zero restrictions at each iteration. When

there are only relatively few restrictions, and subject to some feasibility constraints, there is an alternative to step 2 the algorithm described in the main text which doesn't require the joint solution of a system of equations, but instead works recursively.

The following example illustrates how this alternative algorithm works. As before, the goal is to find m identified structures $\tilde{\mathbf{F}}_m$ which satisfy

$$\tilde{\mathbf{F}}_m \epsilon_t = \mathbf{e}_t$$

and suppose we wanted to set two elements of $\tilde{\mathbf{F}}_m$ equal to zero on impact:

$$\Omega = \{(3, 1), (1, 3)\}$$

Recall that a Givens matrix solves the following problem:

$$G \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} c \\ 0 \end{bmatrix}$$

and let us again start from the Cholesky decomposition of Σ , \mathbf{F} , post-multiplied by an orthonormal matrix \mathbf{Q} .

We can impose $[\mathbf{FQ}]_{31} = 0$ using the second column of \mathbf{FQ} as a "pivot", by solving:

$$\begin{bmatrix} g_{11}^1 & g_{12}^1 \\ g_{21}^1 & g_{22}^1 \end{bmatrix} \begin{bmatrix} [\mathbf{FQ}]_{32} \\ [\mathbf{FQ}]_{31} \end{bmatrix} = \begin{bmatrix} c \\ 0 \end{bmatrix}$$

and then post-multiplying \mathbf{FQ} as follows:

$$\mathbf{FQ}\mathbf{G}_1 = \mathbf{FQ} \begin{bmatrix} g_{11}^1 & g_{12}^1 & 0 \\ g_{21}^1 & g_{22}^1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

Similarly, $[\mathbf{FQ}\mathbf{G}_1]_{13} = 0$ can be imposed by solving:

$$\begin{bmatrix} g_{11}^2 & g_{12}^2 \\ g_{21}^2 & g_{22}^2 \end{bmatrix} \begin{bmatrix} [\mathbf{FQ}\mathbf{G}_1]_{32} \\ [\mathbf{FQ}\mathbf{G}_1]_{31} \end{bmatrix} = \begin{bmatrix} c \\ 0 \end{bmatrix}$$

and then postmultiplying by $G_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & g_{11}^2 & g_{12}^2 \\ 0 & g_{21}^2 & g_{22}^2 \end{bmatrix}$.

Thus,

$$\tilde{\mathbf{F}}_m = \mathbf{FQ}_m \mathbf{G}_1^m \mathbf{G}_2^m.$$

Several things should be noted. First of all, the algorithm is recursive: G_2 is computed on the basis of the postmultiplied matrix $\mathbf{FQ}\mathbf{G}_1$, where one zero restriction had already been imposed. This is possible because a Givens rotation only affects two rows and columns of a matrix, so the first restriction is preserved when we

then postmultiply by G_2 . However, the flip side of this is that once a zero restriction has been imposed in a particular location, neither the row nor the column of that location can be used as a "pivot" for the following restrictions. For example, it would not have been possible to use:

$$G_2 = \begin{bmatrix} g_{11}^2 & 0 & g_{12}^2 \\ 0 & 1 & 0 \\ g_{21}^2 & 0 & g_{22}^2 \end{bmatrix}$$

to impose the second restriction (though the above is in principle a valid Givens matrix), because that would have changed the value of the bottom left element of the matrix, that we had previously set to zero.

This means that the algorithm allows to impose *at most* $2N - 3$ restrictions, but a smaller number might actually be feasible depending on the location of the desired zero restriction. If the restrictions are feasible, however, the gain in computational speed is notable, especially with larger matrices.