The Unit-effect Normalisation in Set-identified Structural Vector Autoregressions

Matthew Read

Research Discussion Paper
2022–[##]

January 2022

Economic Research Department
Reserve Bank of Australia

I thank Adam Cagliarini for useful discussions. The views expressed in this paper are my own and do not necessarily reflect the views of the Reserve Bank of Australia.

Author: ReadM@rba.gov.au

Media Office: rbainfo@rba.gov.au

DRAFT

Abstract

Structural vector autoregressions (SVARs) that are set-identified (e.g. using sign restrictions) are typically estimated under the normalisation that the structural shocks have unit standard deviation, in which case the estimated impulse responses are to a standard-deviation shock. However, impulse responses to a unit shock – a shock that raises a particular variable by one unit – are naturally of greater interest to policymakers. For example, policymakers are interested in answering questions like 'what is the effect of a 100 basis point increase in the federal funds rate?' I show that identified sets for impulse responses obtained under the 'unit-effect normalisation' may be unbounded and discuss issues that this raises for conducting inference. I provide an easily verifiable sufficient condition to check for unboundedness and explain how to check for unboundedness numerically when the sufficient condition is not satisfied. I discuss how to draw useful posterior inferences about impulse responses even when the identified sets for these impulse responses are unbounded at some values of the reduced-form parameters. I illustrate the empirical relevance of these issues by estimating the macroeconomic effects of a 100 basis point shock to the federal funds rate under different identifying restrictions.

JEL Classification Numbers: C32, E52

Keywords: Bayesian inference, impulse responses, monetary policy, set-identified models, sign and zero restrictions

Table of Contents

DRAFT

1. Introduction

When estimating the effects of macroeconomic shocks using structural vector autoregressions (SVARs), it has become increasingly common to use sign restrictions and/or a set of zero restrictions that are insufficient to point-identify the parameters of interest (e.g. Uhlig 2005; Arias, Rubio-Ramírez and Waggoner 2018). Set-identified SVARs are typically estimated under the normalisation that the structural shocks have unit standard deviation. The impulse responses that are obtained under this normalisation consequently represent impulse responses to a standard-deviation shock. However, as argued by Stock and Watson (2016, 2018), impulse responses to a unit shock – a shock that raises a particular variable by one unit – are naturally of more interest to policymakers. For example, when estimating the effects of monetary policy, central bankers are interested in answering questions like 'what is the effect of a 100 basis point increase in the policy rate?' In this paper, I explore issues that arise when conducting inference about the impulse responses to a unit shock.

Under the standard approach to Bayesian inference in set-identified SVARs (e.g. Uhlig 2005; Rubío-Ramirez, Waggoner and Zha 2010; Arias *et al* 2018), it is straightforward to transform from the standard-deviation normalisation to the unit-effect normalisation. As is the case under point-identifying restrictions, this transformation simply requires dividing the impulse responses obtained under the standard-deviation normalisation by the impulse response that is to be normalised to one (e.g. the impact response of the policy rate). Repeating this at each draw of the reduced-form parameters generates a posterior distribution for the impulse responses to a unit shock. However, there are well-documented problems with the standard approach to Bayesian inference in set-identified models. In particular, because the model is set-identified, the likelihood function is flat with respect to a particular parameter. As a consequence, a component of the prior is 'unrevisable' in the sense that it is never updated by the data (Poirier 1998). The prior is also (unintentionally) informative about parameters of interest, such as impulse responses (Baumeister and Hamilton 2015). Together, these features raise the concern that posterior inferences may be driven by the choice of prior rather than by the data and identifying restrictions. This would not necessarily be problematic if the prior reflected credible subjective beliefs about the parameters of interest, but this is not the case; the prior is primarily chosen because it makes Bayesian inference computationally straightforward.

To address these issues, Giacomini and Kitagawa (2021) propose conducting 'robust' (multiple-prior) Bayesian inference in set-identified models. This approach eliminates the source of posterior sensitivity arising due to the unrevisable component of the prior. The key feature of this approach is that it replaces the prior with a class of priors that are consistent with the identifying restrictions (given a fixed prior for the reduced-form parameters). The class of priors generates a class of posteriors, which can be summarised in various ways. For example, rather than generating a single posterior mean, the class of posteriors generates a *set* of posterior means, which is an interval that contains every posterior mean that could be obtained under the class of priors. The class of posteriors can also be summarised using a 'robust credible interval', which is an interval that is assigned at least a given posterior probability under all posterior in the class. Additionally, given a particular hypothesis of interest (e.g. that the output response is negative at some horizon), the class of posteriors generates a set of posterior probabilities for this hypothesis. This set can be summarised by the posterior lower and upper probabilities, which are, respectively, the smallest and largest posterior probabilities of the hypothesis over the class of posteriors.

In the context of set-identified SVARs, implementing the robust Bayesian approach to inference requires computing the lower and upper bounds of the identified set – the set of parameter values corresponding to the same value of the likelihood function – for each impulse response. Doing this at each draw of the reduced-form parameters from their posterior generates posterior distributions for the bounds of the identified set. These posterior distributions can be used to compute the quantities that summarise the class of posteriors. For example, the set of posterior means is an interval with lower (upper) bound equal to the posterior mean of the lower (upper) bound of the identified set.

I highlight complications that arise when using this robust Bayesian approach to conduct inference about impulse responses to a unit shock. In particular, if zero is contained within the identified set for the impulse response that is being normalised to one (e.g. the impact response of the policy rate), the identified sets for the impulse responses to a unit shock will be unbounded. In turn, if these identified sets are unbounded within any region of the reduced-form parameter space that receives positive posterior probability, the sets of posterior means will be unbounded. At face value, this suggests that set-identifying restrictions may be extremely uninformative about the impulse responses to a unit shock. Nevertheless, it still may be possible to draw useful posterior inferences about impulse responses to a unit shock when the identified set is unbounded with positive posterior probability. For example, even when the set of posterior *means* is unbounded, the set of posterior *medians* or some other quantile may be bounded; as a result, it may still be possible to construct robust credible intervals if the credibility level is not too extreme. Moreover, the posterior lower and upper probabilities remain well defined. Consequently, it is always possible to draw inferences such as 'the posterior probability that output declines by more than x per cent at horizon h is at least y per cent and at most z per cent'.[1]

Given the ubiquity of Bayesian methods in the literature on set-identified SVARs, I focus on issues that arise when conducting (robust) Bayesian inference under the unit-effect normalisation. However, similar issues will also arise when conducting frequentist inference about impulse responses to a unit shock. Existing approaches to frequentist inference in set-identified SVARs focus on impulse responses to a standard-deviation shock as the parameters of interest (e.g. Gafarov, Meier and Montiel Olea 2018; Granziera, Moon and Schorfheide 2018). If the maximum-likelihood estimator (MLE) of the reduced-form parameters is such that zero is included within the identified set for the impulse response that is to be normalised, frequentist estimates of the identified sets for all other impulse responses will be unbounded. Even when the identified set at the MLE of the reduced-form parameters is bounded, the bootstrapped distribution of the MLE may include values of the reduced-form parameters that yield unbounded identified sets. In this case, bootstrapped frequentist confidence intervals may or may not be bounded; whether this is the case will depend on the chosen confidence level and the probability that the identified set is unbounded under the bootstrapped sampling distribution of the MLE.

To make these issues clear, I use a static bivariate SVAR in which I can analytically characterise identified sets under a set of sign restrictions on impulse responses. I describe how the intuition in the bivariate example extends to higher-dimensional SVARs identified with both sign and zero restrictions and explain how to verify whether identified sets for impulse responses to a unit shock

1 Giacomini, Kitagawa and Read (2022) describe an algorithm for conducting robust Bayesian inference in proxy SVARs under the unit-effect normalisation. They note that the identified set may be unbounded at some values of the reduced-form parameters, but do not draw out the implications of this issue for conducting inference.

are unbounded in this setting. I provide an easily verifiable sufficient condition for these identified sets to be unbounded; when there is an insufficient number of sign restrictions, the identified set is always unbounded. When this sufficient condition is not satisfied, I explain how to check whether the identified set is unbounded at a given value of the reduced-form parameters by adapting algorithms that have been previously used to check whether identified sets are nonempty (e.g. Giacomini and Kitagawa 2021; Read, forthcoming).

To illustrate the importance of these issues in practice, I estimate the macroeconomic effects of a 100 basis point shock to the federal funds rate under different combinations of identifying restrictions: the sign restrictions on impulse responses to a monetary policy shock proposed in Uhlig (2005); the sign and zero restrictions on the systematic component of monetary policy proposed in Arias, Caldara and Rubio-Ramírez (2019); and the 'narrative restrictions' proposed in Antolín-Díaz and Rubio-Ramírez (2018).

Under the restrictions considered in Arias *et al* (2019), the sufficient condition for unboundedness applies, so the identified sets for impulse responses to a 100 basis point shock are always unbounded. These restrictions are therefore extremely uninformative about the effects of a 100 basis point shock, and outputs obtained using standard Bayesian inference are misleading about the informativeness of the data and identifying restrictions. Combining these restrictions with the sign restrictions on impulse responses considered in Uhlig (2005) yields identified sets that are bounded with posterior probability approaching, but less than, 100 per cent. In this case, since the identified sets are unbounded with positive posterior probability, the sets of posterior means for the output response to a 100 basis point shock are unbounded. Nevertheless, the sets of posterior *medians* remain bounded, because the identified sets are unbounded with low posterior probability. However, the set of posterior medians for the output response includes zero at almost all horizons of interest, and the restrictions are unable to rule out large *increases* in output following a positive 100 basis point shock.

Additionally imposing narrative restrictions on the sign of the monetary policy shock in October 1979 and its contribution to the change in the federal funds rate in this month (as in Antolín-Díaz and Rubio-Ramírez (2018)) results in the identified sets being bounded with 100 per cent posterior probability. This implies that the sets of posterior means and all posterior quantiles are bounded. The additional restrictions substantially tighten the set of posterior medians and robust credible intervals. The results under this set of restrictions are consistent with the largest effects of monetary policy on output occurring after about two years and lying towards the lower end of the range of existing estimates summarised in Ramey (2016).

Finally, I discuss the possibility of using alternative identifying restrictions to ensure that the identified sets for the impulse responses to a unit shock are bounded. A straightforward solution would be to directly restrict the impulse response of the normalising variable to be nonzero; however, I argue that it may be difficult to justify such restrictions and inferences may be sensitive to changes in the imposed bounds. Bounds on the forecast error variance decomposition, such as those proposed in Volpicella (forthcoming), may indirectly constrain the impulse response of the normalising variable to be nonzero, but – again – such restrictions may be difficult to justify in practice and yield results that are highly sensitive to the imposed bounds. Another alternative is to impose a prior directly over the structural parameterisation (as in Baumeister and Hamilton (2015)) that rules out unbounded impulse responses to unit shocks. However, a component of this prior will never be updated by the data, in which case there may again be concerns about the sensitivity of posterior inference to the choice of prior.

The remainder of the paper is structured as follows. Section 2 outlines the SVAR framework and explains the robust Bayesian approach that is used to conduct inference about impulse responses to a unit shock. Section 3 uses a static bivariate SVAR to outline the issues associated with conducting inference about impulse responses to a unit shock. Section 4 describes how to verify whether identified sets for impulse responses to a unit shock are unbounded in a more general setting. Section 5 estimates the macroeconomic effects of a 100 basis point shock to the federal funds rate under different sets of identifying restrictions. Section 6 discusses using alternative restrictions to ensure boundedness of identified sets for impulse responses to a unit shock. Section 7 concludes. Proofs and additional details are contained in the appendices.

Notation. For a matrix $\boldsymbol{X}$, $\mathrm{vec}(\boldsymbol{X})$ is the vectorisation of $\boldsymbol{X}$. When $\boldsymbol{X}$ is symmetric, $\mathrm{vech}(\boldsymbol{X})$ is the half-vectorisation of $\boldsymbol{X}$, which stacks the elements of $\boldsymbol{X}$ that lie on or below the diagonal into a vector. $\boldsymbol{e}_{i,n}$ is the ith column of the $n \times n$ identity matrix, $\boldsymbol{I}_n$. $\boldsymbol{0}_{n\times m}$ is an $n \times m$ matrix of zeros.

2. Framework

This section describes the SVAR model, outlines the concept of identifying restrictions and identified sets, and explains how I conduct robust Bayesian inference about impulse responses to a unit shock.

2.1 SVAR and Orthogonal Reduced Form

Let $\boldsymbol{y}_t$ be an $n \times 1$ vector of variables following the SVAR(p) process:

$$\boldsymbol{A}_0\boldsymbol{y}_t = \boldsymbol{A}_+\boldsymbol{x}_t + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim N(\boldsymbol{0}_{n\times 1}, \boldsymbol{I}_n),$$

where $\boldsymbol{A}_0$ is an invertible $n \times n$ matrix with positive diagonal elements (which is a normalisation on the signs of the structural shocks) and $\boldsymbol{x}_t = (\boldsymbol{y}'_{t-1}, \dots, \boldsymbol{y}'_{t-p})'$. The orthogonal reduced form of the model is:

$$\boldsymbol{y}_t = \boldsymbol{B}\boldsymbol{x}_t + \boldsymbol{\Sigma}_{tr}\boldsymbol{Q}\boldsymbol{\varepsilon}_t,$$

where $\boldsymbol{B} = (\boldsymbol{B}_1, \dots, \boldsymbol{B}_p) = \boldsymbol{A}_0^{-1}\boldsymbol{A}_+$ is the matrix of reduced-form coefficients, $\boldsymbol{\Sigma}_{tr}$ is the lower-triangular Cholesky factor of the variance-covariance matrix of the reduced-form VAR innovations, $\boldsymbol{\Sigma} = E(\boldsymbol{u}_t\boldsymbol{u}_t') = \boldsymbol{A}_0^{-1}(\boldsymbol{A}_0^{-1})'$ with $\boldsymbol{u}_t = \boldsymbol{y}_t - \boldsymbol{B}\boldsymbol{x}_t$, and $\boldsymbol{Q}$ is an $n \times n$ orthonormal matrix (i.e. $\boldsymbol{Q}\boldsymbol{Q}' = \boldsymbol{I}_n$). The reduced-form parameters are denoted by $\boldsymbol{\phi} = (\mathrm{vec}(\boldsymbol{B})', \mathrm{vech}(\boldsymbol{\Sigma}_{tr})')' \in \boldsymbol{\Phi}$ and the space of $n \times n$ orthonormal matrices by $\mathcal{O}(n)$. I assume that the space of reduced-form parameters $\boldsymbol{\Phi}$ is such that the infinite-order vector moving average representation of the VAR exists.[2]

The impulse responses to standard-deviation shocks are obtained from the coefficients of the vector moving average representation:

2 This is the case when the eigenvalues of the companion matrix lie inside the unit circle (e.g. Hamilton 1994; Kilian and Lütkepohl 2017).

$$\boldsymbol{y}_t = \sum_{h=0}^{\infty} \boldsymbol{C}_h \boldsymbol{\Sigma}_{tr} \boldsymbol{Q} \boldsymbol{\varepsilon}_{t-h},$$

where $\boldsymbol{C}_h$ is defined recursively by $\boldsymbol{C}_h = \sum_{l=1}^{\min\{h,p\}} \boldsymbol{B}_l \boldsymbol{C}_{h-l}$ for $h \geq 1$ with $\boldsymbol{C}_0 = \boldsymbol{I}_n$. The (i,j)th element of the matrix $\boldsymbol{C}_h \boldsymbol{\Sigma}_{tr} \boldsymbol{Q}$ is the horizon-h impulse response of the ith variable to the jth structural shock, denoted by $\eta_{i,j,h}(\boldsymbol{\phi}, \boldsymbol{Q}) = \boldsymbol{c}'_{ih}(\boldsymbol{\phi})\boldsymbol{q}_j$, where $\boldsymbol{c}'_{ih}(\boldsymbol{\phi}) = \boldsymbol{e}'_{i,n}\boldsymbol{C}_h$ is the ith row of $\boldsymbol{C}_h$ and $\boldsymbol{q}_j = \boldsymbol{Q}\boldsymbol{e}_{j,n}$ is the jth column of $\boldsymbol{Q}$. The horizon-h impulse response of the ith variable to a shock in the first variable that raises the first variable by one unit on impact is then

$$\tilde{\eta}_{i,1,h}(\boldsymbol{\phi}, \boldsymbol{Q}) = \frac{\eta_{i,1,h}(\boldsymbol{\phi}, \boldsymbol{Q})}{\eta_{1,1,0}(\boldsymbol{\phi}, \boldsymbol{Q})} = \frac{\boldsymbol{c}'_{ih}(\boldsymbol{\phi})\boldsymbol{q}_1}{\mathbf{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1}.$$

In what follows, I refer to this parameter as an 'impulse response to a unit shock'. The assumption that the normalising impulse response is the impact response of the first variable to the first shock is made without loss of generality. In some contexts, it may be natural to normalise the impulse responses such that a specific variable increases by one unit at a longer horizon; for example, when estimating the effects of news shocks, the natural normalising variable may not respond much at shorter horizons.

2.2 Identifying Restrictions and Identified Sets

Imposing restrictions on (functions of) the structural parameters is equivalent to imposing restrictions on $\boldsymbol{Q}$ given $\boldsymbol{\phi}$; for example, consider a sign restriction on an impulse response such that $\eta_{i,j,h}(\boldsymbol{\phi}, \boldsymbol{Q}) = \boldsymbol{c}'_{ih}(\boldsymbol{\phi})\boldsymbol{q}_j \geq 0$. This is a linear inequality restriction on $\boldsymbol{q}_j$, where the coefficients in the restriction are a function of $\boldsymbol{\phi}$. More generally, let $S(\boldsymbol{\phi}, \boldsymbol{Q}) \geq \mathbf{0}_{s\times 1}$ represent a collection of s sign restrictions (including the sign normalisation $\mathrm{diag}(\boldsymbol{A}_0) \geq \mathbf{0}_{n\times 1}$). Similarly, represent a collection of f zero restrictions by $F(\boldsymbol{\phi}, \boldsymbol{Q}) = \mathbf{0}_{f\times 1}$. For example, these could include zero restrictions on impulse responses, elements of $\boldsymbol{A}_0$ or long-run cumulative impulse responses.[3]

Let f_i represent the number of zero restrictions constraining the ith column of $\boldsymbol{Q}$, so $\sum_{i=1}^{n} f_i = f$. I assume that the variables are ordered such that f_i is weakly decreasing and that $f_i \leq n - i$ for $i = 1, \ldots, n$ with strict inequality for at least one i; this is a sufficient condition for the model to be set-identified (Rubio-Ramírez *et al* 2010; Bacchiocchi and Kitagawa 2021).

Given a collection of sign and zero restrictions, the identified set for $\boldsymbol{Q}$ – which collects observationally equivalent parameter values (i.e. parameter values with the same value of the likelihood function) – is

$$\mathcal{Q}(\boldsymbol{\phi}|S,F) = \left\{\boldsymbol{Q} \in \mathcal{O}(n) : S(\boldsymbol{\phi}, \boldsymbol{Q}) \geq \mathbf{0}_{s\times 1}, F(\boldsymbol{\phi}, \boldsymbol{Q}) = \mathbf{0}_{f\times 1}\right\}.$$

The identified set for a particular impulse response is then the set of values of that impulse response as $\boldsymbol{Q}$ varies over its identified set; that is, $\eta_{i,j,h}(\boldsymbol{\phi}|S,F) = \left\{\eta_{i,j,h}(\boldsymbol{\phi}, \boldsymbol{Q}) : \boldsymbol{Q} \in \mathcal{Q}(\boldsymbol{\phi}|S,F)\right\}$ or $\tilde{\eta}_{i,j,h}(\boldsymbol{\phi}|S,F) = \left\{\tilde{\eta}_{i,j,h}(\boldsymbol{\phi}, \boldsymbol{Q}) : \boldsymbol{Q} \in \mathcal{Q}(\boldsymbol{\phi}|S,F)\right\}$. Note that these identified sets may be empty at particular values of $\boldsymbol{\phi}$.

3 See Stock and Watson (2016) or Kilian and Lütkepohl (2017) for overviews of approaches to identification in SVARs.

2.3 Robust Bayesian Inference in Set-identified SVARs

The standard approach to conducting Bayesian inference in set-identified SVARs involves specifying a prior for the reduced-form parameters $\boldsymbol{\phi}$ and a uniform prior for the orthonormal matrix $\boldsymbol{Q}$ (Uhlig 2005; Rubio-Ramírez *et al* 2010; Arias *et al* 2018).[4] To draw from the resulting posterior in practice, one samples values of $\boldsymbol{\phi}$ from its posterior and $\boldsymbol{Q}$ from a uniform distribution over $Q(\boldsymbol{\phi}|F)$ and discards draws that violate the sign restrictions.[5] Assume there is a scalar parameter of interest that is a function of the structural parameters, η (e.g. a particular impulse response). Draws of η are obtained by transforming the draws of $\boldsymbol{\phi}$ and $\boldsymbol{Q}$, and the posterior is summarised using quantities such as the posterior mean and quantiles.

Let $\pi_{\boldsymbol{\phi}}$ be a prior for $\boldsymbol{\phi} \in \boldsymbol{\Phi}$, where $\boldsymbol{\Phi}$ is the space of reduced-form parameters such that $Q(\boldsymbol{\phi}|S,F)$ is nonempty. A joint prior for the full set of parameters $\boldsymbol{\theta} = (\boldsymbol{\phi}',\text{vec}(\boldsymbol{Q})')'$ can be decomposed as $\pi_{\boldsymbol{\theta}} = \pi_{\boldsymbol{Q}|\boldsymbol{\phi}}\pi_{\boldsymbol{\phi}}$, where $\pi_{\boldsymbol{Q}|\boldsymbol{\phi}}$ is the conditional prior for $\boldsymbol{Q}$ given $\boldsymbol{\phi}$, which is zero outside of $Q(\boldsymbol{\phi}|S,F)$. After observing the data $\boldsymbol{Y}$, the posterior is $\pi_{\boldsymbol{\theta}|\boldsymbol{Y}} = \pi_{\boldsymbol{\phi}|\boldsymbol{Y}}\pi_{\boldsymbol{Q}|\boldsymbol{\phi}}$, where $\pi_{\boldsymbol{\phi}|\boldsymbol{Y}}$ is the posterior for $\boldsymbol{\phi}$. The prior for $\boldsymbol{\phi}$ is updated by the data (through the likelihood), whereas the conditional prior for $\boldsymbol{Q}$ given $\boldsymbol{\phi}$ is not, because $\boldsymbol{Q}$ does not appear in the likelihood. Baumeister and Hamilton (2015) show that a uniform conditional prior for $\boldsymbol{Q}$ may be unintentionally informative about parameters of interest, such as impulse responses. Combining this observation with the fact that a component of the prior is never updated by the data raises the concern that posterior inferences may be sensitive to the choice of $\pi_{\boldsymbol{Q}|\boldsymbol{\phi}}$. Given that this prior is chosen primarily for computational convenience and does not necessarily reflect credible prior information about the parameters of interest, it is desirable that researchers eliminate or quantify the sensitivity of posterior inference to the choice of unrevisable conditional prior.

To this end, I adopt a 'robust' (multiple-prior) Bayesian approach proposed by Giacomini and Kitagawa (2021) to conduct posterior inference in set-identified models.[6] In the context of an SVAR, this approach eliminates the source of posterior sensitivity arising due to the unrevisable conditional prior for $\boldsymbol{Q}$ given $\boldsymbol{\phi}$. Importantly, this helps to disentangle the information in the posterior that is contributed by the data and the identifying restrictions from the information that is contributed by the prior. The key feature of the approach is that it replaces the unrevisable conditional prior for $\boldsymbol{Q}$ given $\boldsymbol{\phi}$ with the class of all conditional priors that satisfy the identifying restrictions:

$$\Pi_{\boldsymbol{Q}|\boldsymbol{\phi}} = \left\{\pi_{\boldsymbol{Q}|\boldsymbol{\phi}} : \pi_{\boldsymbol{Q}|\boldsymbol{\phi}}\big(Q(\boldsymbol{\phi}|S,F)\big) = 1\right\}.$$

---

4 This approach to conducting inference in SVARs under sign restrictions is ubiquitous in the empirical literature. Alternative approaches are to conduct Bayesian inference under a prior that is specified directly over the structural parameters of the SVAR (Baumeister and Hamilton, 2015) or to conduct frequentist inference (e.g. Gafarov *et al* 2018; Granziera *et al* 2018).

5 When conducting standard Bayesian inference below, I impose a *conditionally* uniform prior for $\boldsymbol{Q}$ given $\boldsymbol{\phi}$ by obtaining a single draw of $\boldsymbol{Q}$ from $Q(\boldsymbol{\phi}|S,F)$ at each draw of $\boldsymbol{\phi}$ (whenever $Q(\boldsymbol{\phi}|S,F)$ is nonempty). In contrast, rejecting joint draws of $(\boldsymbol{\phi},\boldsymbol{Q})$ that violate the identifying restrictions implicitly places more weight (relative to the notional prior for $\boldsymbol{\phi}$) on values of $\boldsymbol{\phi}$ that are less likely to satisfy the sign restrictions under the uniform distribution over $Q(\boldsymbol{\phi}|F)$. See Uhlig (2017) for a discussion of this point.

6 See Giacomini, Kitagawa and Read (2021b) for a review of the literature on robust Bayesian analysis in econometrics. This review includes a detailed description of different approaches to conducting robust Bayesian inference in set-identified SVARs.

Combining the class of priors $\Pi_{\boldsymbol{Q}|\boldsymbol{\phi}}$ with the posterior for $\boldsymbol{\phi}$, $\pi_{\boldsymbol{\phi}|\boldsymbol{Y}}$, generates a class of posteriors for $\boldsymbol{\theta}$:

$$\Pi_{\boldsymbol{\theta}|\boldsymbol{Y}} = \left\{\pi_{\boldsymbol{\theta}|\boldsymbol{Y}} = \pi_{\boldsymbol{Q}|\boldsymbol{\phi}}\pi_{\boldsymbol{\phi}|\boldsymbol{Y}} : \pi_{\boldsymbol{Q}|\boldsymbol{\phi}} \in \Pi_{\boldsymbol{Q}|\boldsymbol{\phi}}\right\}.$$

The class of posteriors for $\boldsymbol{\theta}$ induces a class of posteriors for η, $\Pi_{\eta|\boldsymbol{Y}}$. Giacomini and Kitagawa (2021) suggest summarising this class of posteriors by reporting the 'set of posterior means':

$$\left[\int_{\boldsymbol{\Phi}} \ell(\boldsymbol{\phi}) d\pi_{\boldsymbol{\phi}|\boldsymbol{Y}}, \int_{\boldsymbol{\Phi}} u(\boldsymbol{\phi}) d\pi_{\boldsymbol{\phi}|\boldsymbol{Y}}\right],$$

where $\ell(\boldsymbol{\phi}) = \inf\{\eta(\boldsymbol{\phi}, \boldsymbol{Q}) : \boldsymbol{Q} \in \boldsymbol{Q}(\boldsymbol{\phi}|S, F)\}$ is the lower bound of the identified set for η given $\boldsymbol{\phi}$ and $u(\boldsymbol{\phi}) = \sup\{\eta(\boldsymbol{\phi}, \boldsymbol{Q}) : \boldsymbol{Q} \in \boldsymbol{Q}(\boldsymbol{\phi}|S, F)\}$ is the upper bound. The set of posterior means is an interval that contains all possible posterior means that could be obtained under the class of priors $\Pi_{\boldsymbol{Q}|\boldsymbol{\phi}}$. They also suggest reporting a robust credible region with credibility level α, which is an interval estimate for η such that the posterior probability put on the interval is at least α for all posteriors in the class $\Pi_{\eta|\boldsymbol{Y}}$. Given a particular hypothesis of interest (e.g. that the output response to a monetary policy shock is negative at some horizon), the set of posteriors also generates a set of posterior probabilities for this hypothesis. This set can be summarised by the posterior lower and upper probabilities, which are, respectively, the smallest or largest posterior probabilities of the hypothesis over all posteriors in $\Pi_{\eta|\boldsymbol{Y}}$. Each of these quantities can be computed from the posterior distributions of $\ell(\boldsymbol{\phi})$ and $u(\boldsymbol{\phi})$.

3. The Unit-effect Normalisation in a Bivariate SVAR

To illustrate the issues that arise when conducting inference about impulse responses to a unit shock, I consider the simplest possible SVAR – a bivariate SVAR with no dynamics – identified using simple sign restrictions on impulse responses. This allows me to analytically derive identified sets for the impulse responses. See Appendix A for derivations of the results in this section.

The model is $\boldsymbol{A}_0\boldsymbol{y}_t = \boldsymbol{\varepsilon}_t$, where $\boldsymbol{y}_t = (y_{1t}, y_{2t})'$, $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t})'$ and $E(\boldsymbol{\varepsilon}_t\boldsymbol{\varepsilon}_t') = \boldsymbol{I}_2$. The orthogonal reduced form of the model is $\boldsymbol{y}_t = \boldsymbol{\Sigma}_{tr}\boldsymbol{Q}\boldsymbol{\varepsilon}_t$, where $\boldsymbol{\Sigma}_{tr}$ is the lower-triangular Cholesky factor of $\boldsymbol{\Sigma} = E(\boldsymbol{y}_t\boldsymbol{y}_t')$ and $\boldsymbol{Q}$ is a 2×2 orthonormal matrix. I denote the reduced-form parameter as $\boldsymbol{\phi} = \mathrm{vech}(\boldsymbol{\Sigma}_{tr}) = (\sigma_{11}, \sigma_{21}, \sigma_{22})'$ with $\sigma_{11}, \sigma_{22} > 0$. In the bivariate case, the space of 2×2 orthonormal matrices can be represented as

$$\mathcal{O}(2) = \left\{\begin{bmatrix}\cos\theta & -\sin\theta \\ \sin\theta & \cos\theta\end{bmatrix} : \theta \in [-\pi, \pi]\right\} \cup \left\{\begin{bmatrix}\cos\theta & \sin\theta \\ \sin\theta & -\cos\theta\end{bmatrix} : \theta \in [-\pi, \pi]\right\},$$

where the first set is the set of 'rotation' matrices and the second is the set of 'reflection' matrices. Henceforth, I leave it implicit that $\theta \in [-\pi, \pi]$.[7]

In the absence of any identifying restrictions, the identified set for $\boldsymbol{A}_0^{-1}$ (the matrix of impact impulse responses) is

$$\boldsymbol{A}_0^{-1} \in \left\{\begin{bmatrix} \sigma_{11}\cos\theta & -\sigma_{11}\sin\theta \\ \sigma_{21}\cos\theta + \sigma_{22}\sin\theta & \sigma_{22}\cos\theta - \sigma_{21}\sin\theta \end{bmatrix}\right\} \cup \left\{\begin{bmatrix} \sigma_{11}\cos\theta & \sigma_{11}\sin\theta \\ \sigma_{21}\cos\theta + \sigma_{22}\sin\theta & \sigma_{21}\sin\theta - \sigma_{22}\cos\theta \end{bmatrix}\right\},$$

and the identified set for $\boldsymbol{A}_0$ is

$$\boldsymbol{A}_0 \in \left\{\frac{1}{\sigma_{11}\sigma_{22}}\begin{bmatrix} \sigma_{22}\cos\theta - \sigma_{21}\sin\theta & \sigma_{11}\sin\theta \\ -\sigma_{21}\cos\theta - \sigma_{22}\sin\theta & \sigma_{11}\cos\theta \end{bmatrix}\right\} \cup \left\{\frac{1}{\sigma_{11}\sigma_{22}}\begin{bmatrix} \sigma_{22}\cos\theta - \sigma_{21}\sin\theta & \sigma_{11}\sin\theta \\ \sigma_{21}\cos\theta + \sigma_{22}\sin\theta & -\sigma_{11}\cos\theta \end{bmatrix}\right\}.$$

Throughout, I impose the sign normalisation $\text{diag}(\boldsymbol{A}_0) \geq \mathbf{0}_{2\times 1}$.

Consider the case where the impact response of the first variable to the first shock is restricted to be nonnegative ($\eta_{1,1,0} \equiv \mathbf{e}'_{1,2}\mathbf{A}_0^{-1}\mathbf{e}_{1,2} \geq 0$) and the impact response of the second variable to the first shock is restricted to be nonpositive ($\eta_{2,1,0} \equiv \mathbf{e}'_{2,2}\mathbf{A}_0^{-1}\mathbf{e}_{1,2} \leq 0$). These sign restrictions and the sign normalisation restrict the set of values that θ can take given the reduced-form parameters; that is, they generate an identified set for θ, which can in turn be used to obtain an identified set for $\eta_{1,1,0}$:

$$\eta_{1,1,0} \in \begin{cases} \left[\sigma_{11}\cos\left(\arctan\left(\min\left\{\frac{\sigma_{22}}{\sigma_{21}}, \frac{\sigma_{21}}{\sigma_{22}}\right\}\right)\right), \sigma_{11}\right] & \text{if } \sigma_{21} < 0 \\ \left[0, \sigma_{11}\cos\left(\arctan\left(-\frac{\sigma_{21}}{\sigma_{22}}\right)\right)\right] & \text{if } \sigma_{21} \geq 0. \end{cases}$$

When $\sigma_{21} < 0$, the identified set for $\eta_{1,1,0}$ excludes zero. In contrast, when $\sigma_{21} \geq 0$, the identified set for $\eta_{1,1,0}$ includes zero.

The impulse response of the second variable to a unit shock in the first variable is

$$\tilde{\eta}_{2,1,0} \equiv \frac{\eta_{2,1,0}}{\eta_{1,1,0}} = \frac{\sigma_{21}\cos\theta + \sigma_{22}\sin\theta}{\sigma_{11}\cos\theta} = \frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}}{\sigma_{11}}\tan\theta.$$

The identified set for this impulse response is

$$\tilde{\eta}_{2,1,0} \in \begin{cases} \left[\frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}^2}{\sigma_{11}\sigma_{21}}, 0\right] & \text{if } \sigma_{21} < 0 \\ (-\infty, 0] & \text{if } \sigma_{21} \geq 0. \end{cases}$$

7 Baumeister and Hamilton (2015) use a similar example to show that the standard uniform prior over $\boldsymbol{Q}$ is informative about impulse responses. In particular, the implicit prior over the impulse response to a unit shock is a Cauchy distribution that is truncated by the sign restrictions, where the points of truncation depend on the reduced-form parameters. In contrast to their work, my example focuses on the *identified set* for the impulse response to a unit shock as the object of interest. Giacomini, Kitagawa and Read (2021a) use a similar example to illustrate issues that arise when imposing 'narrative restrictions', which are restrictions that involve the values of the structural shocks in specific periods.

When $\sigma_{21} < 0$, the lower bound of this identified set is negative and finite, while the upper bound is zero. The identified set for $\tilde{\eta}_{2,1,0}$ is therefore bounded. In contrast, when $\sigma_{21} \geq 0$, the identified set for $\tilde{\eta}_{2,1,0}$ is unbounded (below); $\tilde{\eta}_{2,1,0}$ diverges to $-\infty$ as θ approaches $-\pi/2$ (which is the lower bound of the identified set for θ) from above. The upper bound of the identified set for this impulse response is equal to zero, so the sign restrictions are completely uninformative about $\tilde{\eta}_{2,1,0}$ outside of its sign (which is imposed).

This example highlights that the identified set for the impulse response to a unit shock is unbounded if the identified set for the impact response of the normalising variable to a standard-deviation shock includes zero. The exercise also demonstrates that sign restrictions may or may not be informative about impulse responses to a unit shock. Whether this is the case depends on the values of the reduced-form parameters. The following sections discuss implications of an unbounded identified set for conducting inference about the impulse responses to a unit shock.[8]

3.1 Issues Conducting Robust Bayesian Inference

In the context of conducting robust Bayesian inference in this example, the set of posterior means for the impulse response to a unit shock will be bounded only if $\pi_{\boldsymbol{\phi}|Y}(\sigma_{21} \geq 0) = 0$. However, even when $\pi_{\boldsymbol{\phi}|Y}(\sigma_{21} \geq 0) > 0$, the set of posterior *medians* (or other specific quantiles) may be bounded. The posterior lower and upper probabilities of some hypothesis (e.g. that $\tilde{\eta}_{2,1,0}$ exceeds some value) also remain well defined. To elaborate on these claims, I make the simplifying assumption that $\pi_{\boldsymbol{\phi}|Y}$ is supported only on two values of the reduced-form parameters: $\boldsymbol{\phi}^a = (\sigma_{11}, \sigma_{21}^a, \sigma_{22})'$ and $\boldsymbol{\phi}^b = (\sigma_{11}, \sigma_{21}^b, \sigma_{22})'$, where $\sigma_{21}^a < 0 \leq \sigma_{21}^b$. I denote the lower bound of the identified set for $\tilde{\eta}_{2,1,0}$ when $\sigma_{21} = \sigma_{21}^a$ by $L(\boldsymbol{\phi}^a)$ and the posterior probability that $\sigma_{21} = \sigma_{21}^a$ by α.

Under this assumption, the identified set for $\tilde{\eta}_{2,1,0}$ is $[L(\boldsymbol{\phi}^a), 0]$ with posterior probability α and it is $(-\infty, 0]$ with posterior probability $1-\alpha$. The set of posterior means, which has bounds equal to the posterior means of the bounds of the identified set, will therefore be $(-\infty, 0]$ unless $\alpha = 1$. Consequently, if $\pi_{\boldsymbol{\phi}|Y}$ places positive probability on the event $\sigma_{21} = \sigma_{21}^b \geq 0$, the set of posterior means is completely uninformative about the impulse response to a unit shock (other than its sign, which is imposed by the sign restrictions).

The median of the upper bound of the identified set is zero regardless of the value of α. When $\alpha \geq 0.5$, the posterior median of the lower bound of the identified set is $L(\boldsymbol{\phi}^a)$. The set of posterior *medians* – which is an interval with lower (upper) bound equal to the posterior median of the lower (upper) bound of the identified set – will therefore be bounded despite the set of posterior *means* being unbounded. In contrast, when $\alpha < 0.5$, the posterior median of the lower bound of the identified set is $-\infty$, so the set of posterior medians is unbounded. By similar logic, the set of posterior τ-quantiles will be bounded so long as $\alpha \geq \tau$. The class of posteriors may therefore still contain useful information about particular posterior quantiles of the impulse responses to a unit shock even when the identified set is unbounded with positive posterior probability.

---

8 Similar issues will arise when the parameter of interest is the structural coefficient on a particular variable after normalising the coefficient on another variable to equal unity, which is the ratio of elements of $\boldsymbol{A}_0$. If the identified set for the normalising coefficient includes zero, the identified set for the ratio of coefficients will be unbounded. Such ratios can be interpreted as behavioural elasticities (e.g. the slope of a demand or supply curve). In the bivariate model, this type of parameter can equivalently be expressed as an impulse response to a unit shock (i.e. the ratio of elements of $\boldsymbol{A}_0^{-1}$). This equivalence does not hold in larger-dimensional models; see the discussion in Baumeister and Hamilton (2022).

Consider the hypothesis that $\tilde{\eta}_{2,1,0} \leq x$ for some $x < 0$. The posterior *lower* probability of this hypothesis is equal to the posterior probability that the identified set is contained within the interval $(-\infty, x]$. This probability is zero for all $x < 0$. The posterior *upper* probability of the hypothesis is equal to the posterior probability that the identified set intersects the interval $(-\infty, x]$. The posterior upper probability is one for $L(\boldsymbol{\phi}^a) \leq x < 0$ and is $1-\alpha$ for $x < L(\boldsymbol{\phi}^a)$. The set of posterior probabilities for the hypothesis $\tilde{\eta}_{2,1,0} \leq x$ is therefore $[0,1]$ for $L(\boldsymbol{\phi}^a) \leq x < 0$ and is $[0,1-\alpha]$ for $x < L(\boldsymbol{\phi}^a)$. As α approaches zero, so that the identified set is almost always unbounded, the set of posterior probabilities converges to the unit interval for all values of x. In this case, the sign restrictions are not informative about the hypothesis regardless of the value of x. In contrast, as α approaches one, the set of posterior probabilities converges to zero for sufficiently negative values of x (i.e. for $x < L(\boldsymbol{\phi}^a)$). In this case, the sign restrictions can rule out 'large' responses to a unit shock with high posterior probability even though the identified set for the response is sometimes unbounded.

This discussion illustrates that it is still possible to extract information about the impulse responses to a unit shock using the robust Bayesian approach to inference when the identified set is unbounded with positive posterior probability. The takeaways from this stylised model extend to the general case of an n-dimensional SVAR with dynamics and/or where the posterior for $\boldsymbol{\phi}$ has continuous support.

3.1.1 Frequentist validity of robust Bayesian approach

For general set-identified models, Giacomini and Kitagawa (2021) provide high-level conditions under which their robust Bayesian approach to inference has a valid frequentist interpretation, in the sense that the set of posterior means is consistent for the true identified set (i.e. the identified set when $\boldsymbol{\phi}$ is equal to its true value, $\boldsymbol{\phi}_0$) and the robust credible interval has correct frequentist coverage for the true identified set. In the context of SVARs and when the parameter of interest is an impulse response to a standard-deviation shock, Giacomini and Kitagawa (2021) provide sufficient conditions under which these high-level conditions will hold. In particular, the set of posterior means can be interpreted as a consistent estimator of the true identified set if the identified set is convex and continuous at $\boldsymbol{\phi} = \boldsymbol{\phi}_0$. Additionally, if the endpoints of the identified set ($\ell(\boldsymbol{\phi})$ and $u(\boldsymbol{\phi})$) are differentiable in $\boldsymbol{\phi}$ at $\boldsymbol{\phi} = \boldsymbol{\phi}_0$ with nonzero derivatives, the robust credible interval has valid frequentist coverage of the true identified set.[9]

When the parameter of interest is an impulse response to a unit shock, the high-level conditions for frequentist validity of the robust Bayesian approach are not necessarily satisfied. For example, these conditions include the assumption that the true identified set is bounded. Consequently, the robust Bayesian approach to inference is not guaranteed to have an asymptotically valid frequentist interpretation when the parameter of interest is an impulse response to a unit shock.

To illustrate, consider the bivariate model and assume that $\boldsymbol{\phi}_0$ is such that $\sigma_{21} > 0$, so the true identified set is unbounded. In general, a robust credible region with credibility $1-\tau$ can be constructed be taking the $\tau/2$ quantile of $\ell(\boldsymbol{\phi})$ and the $1-\tau/2$ quantile of $u(\boldsymbol{\phi})$. For values of $\boldsymbol{\phi}$ in

9 If the identified set is not convex, the robust Bayesian output can be interpreted as providing valid frequentist inference about the convex hull of the identified set. Non-convexity of impulse-response identified sets can arise when the identifying restrictions constrain multiple columns of $\boldsymbol{Q}$ (see Example B.5 in Giacomini and Kitagawa (2020)).

a small neighbourhood of $\boldsymbol{\phi}_0$, $l(\boldsymbol{\phi}) = -\infty$ and $u(\boldsymbol{\phi}) = 0$, so naively applying the robust Bayesian approach in this case will yield a robust credible interval of $(-\infty, 0]$. Clearly, this interval always (weakly) includes the true identified set, so the asymptotic frequentist coverage probability will be trivially equal to one, which is greater than the nominal credibility level τ (i.e. the robust credible interval is conservative).[10]

3.2 Issues Conducting Frequentist Inference

Similar issues will also arise when conducting frequentist inference about impulse responses to a unit shock. Let $\widehat{\boldsymbol{\phi}}$ be the MLE of $\boldsymbol{\phi}$. In the current bivariate example, if $\widehat{\boldsymbol{\phi}}$ is such that $\hat{\sigma}_{21} \geq 0$, a frequentist estimate of the identified set for $\tilde{\eta}_{2,1,0}$ – which simply plugs the MLE of $\widehat{\boldsymbol{\phi}}$ into the expression for the identified set given above – will be unbounded. If $\widehat{\boldsymbol{\phi}}$ is such that $\hat{\sigma}_{21} < 0$, the frequentist estimate of the identified set for $\tilde{\eta}_{2,1,0}$ will be bounded. However, even then, issues may arise when attempting to conduct frequentist inference about the impulse response. For example, if one were to apply a frequentist bootstrap (which is common in the context of point-identified SVARs), confidence intervals at some confidence levels will be unbounded when there is a positive probability that $\hat{\sigma}_{21} \geq 0$ under the bootstrapped sampling distribution of $\widehat{\boldsymbol{\phi}}$.

4. Checking for Unboundedness in SVARs

As noted above, the lessons from the bivariate model of Section 3 extend to the general setting of an n-dimensional SVAR with dynamics. They also extend to the case where there are both sign and zero restrictions on the structural parameters. In this general setting, analytical expressions for impulse-response identified sets are typically not available and it is necessary to compute or approximate the identified set numerically. This section explains how to check whether the identified sets for the impulse responses to a unit shock are unbounded in this setting.

Accurately verifying whether the identified set is unbounded is helpful for understanding whether particular inferential outputs (e.g. sets of posterior means or quantiles) are themselves unbounded. From a practical standpoint, it is important to check whether the identified set is unbounded before attempting to compute the bounds of the identified set. One approach to computing these bounds is to use a numerical optimisation routine where the objective function to be minimised or maximised is $\tilde{\eta}_{i,j,h}(\boldsymbol{\phi}, \boldsymbol{Q})$ and the constraints are the set of identifying restrictions. If the identified set is unbounded given $\boldsymbol{\phi}$, standard gradient-based numerical optimisation routines (e.g. an interior-point algorithm) will terminate at some large, but arbitrary, value of the objective function. Another approach to computing the bounds is to obtain many random draws of $\boldsymbol{Q}$ from a distribution over $\mathcal{Q}(\boldsymbol{\phi}|S,F)$ and compute the minimum and maximum over these draws. When the identified set is bounded, the approximation error from this approach will vanish as the number of draws increases, but this will not be the case when the identified set is unbounded.

In the n-variable SVAR (described in Section 2), assume that the sign restrictions $S(\boldsymbol{\phi}, \boldsymbol{Q}) \geq \mathbf{0}_{s\times 1}$ include the restriction that the impact response of the first variable to the first shock is nonnegative, $\eta_{1,1,0} = \boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 \geq 0$. For example, in the context of identifying the effects of

10 In the case where $\boldsymbol{\phi}_0$ is such that $\sigma_{21} < 0$, the true identified set is bounded, but the robust credible interval has an asymptotic frequentist coverage probability equal to $1 - \tau/2 > 1 - \tau$ (i.e. the robust credible interval is conservative). This arises because the upper bound of the identified set is degenerate and is therefore not differentiable in $\boldsymbol{\phi}$ with non-zero derivative, which means that asymptotic normality of the reduced-form parameters does not translate into asymptotic normality of $u(\boldsymbol{\phi})$ via the delta method.

monetary policy shocks, this restriction would require that a positive monetary policy shock (the first shock) does not decrease the federal funds rate (the first variable) on impact. Such a restriction seems natural. The identified set for $\tilde{\eta}_{i,1,h}$, $(i,h) \neq (1,0)$, will be unbounded if and only if the identified set for $\eta_{1,1,0}$ includes zero.[11] This will be the case if there exists $\boldsymbol{Q}$ satisfying the zero restrictions, the 'binding' sign restriction on $\eta_{1,1,0}$ ($\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 = 0$), and any remaining sign restrictions. The following proposition formalises this claim.

Proposition 3.1. *(Necessary and sufficient condition for unbounded identified sets.) Assume* $\eta_{1,1,0} = \boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 \geq 0$ *is included within the set of sign restrictions* $S(\boldsymbol{\phi},\boldsymbol{Q}) \geq \boldsymbol{0}_{s\times 1}$ *and that* $Q(\boldsymbol{\phi}|S,F)$ *is nonempty. The identified set for the impulse response to a unit shock to the first variable,* $\tilde{\eta}_{i,1,h}(\boldsymbol{\phi}|S,F)$*, is unbounded for* $(i,h) \neq (1,0)$ *if and only if* $0 \in \eta_{1,1,0}(\boldsymbol{\phi}|S,F)$.

Consider imposing a set of zero and sign restrictions constraining $\boldsymbol{q}_1$ only, $F(\boldsymbol{\phi},\boldsymbol{Q}) = F(\boldsymbol{\phi})\boldsymbol{q}_1 = \boldsymbol{0}_{f\times 1}$ and $S(\boldsymbol{\phi},\boldsymbol{Q}) = S(\boldsymbol{\phi})\boldsymbol{q}_1 \geq \boldsymbol{0}_{s\times 1}$. Again, assume that the sign restrictions represented in $S(\boldsymbol{\phi})$ include the restriction that the impact response of the first variable to the first shock is nonnegative, $\eta_{1,1,0} = \boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 \geq 0$. In this setting, the following proposition states a sufficient condition for the identified set for $\eta_{1,1,0}$ to include zero.

Proposition 3.2. *(Sufficient condition for identified set for* $\eta_{1,1,0}$ *to include zero.) Assume any sign and zero restrictions constrain* $\boldsymbol{q}_1$ *only,* $\eta_{1,1,0} = \boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 \geq 0$ *is contained within the set of sign restrictions* $S(\boldsymbol{\phi})\boldsymbol{q}_1 \geq \boldsymbol{0}_{s\times 1}$ *and the number of zero restrictions in* $F(\boldsymbol{\phi})\boldsymbol{q}_1 = \boldsymbol{0}_{f\times 1}$ *satisfies* $0 \leq f < n-1$ *with* $\mathrm{rank}(F(\boldsymbol{\phi})) = r$. *If* $s + f \leq n$, *then* $0 \in \eta_{1,1,0}(\boldsymbol{\phi}|S,F)$ *for any value of* $\boldsymbol{\phi} \in \boldsymbol{\Phi}$.

The sufficient condition in Proposition 3.2 can be used to verify whether the identified set for $\tilde{\eta}_{i,1,h}$ ($(i,h) \neq (1,0)$) is always unbounded at any value of $\boldsymbol{\phi}$. The condition is easily verifiable; it simply requires counting the number of sign and zero restrictions imposed. Although the proposition only applies when the identifying restrictions constrain a single column of $\boldsymbol{Q}$, this is the case in many empirical applications; examples include Uhlig (2005) and Arias *et al* (2019) (see also the references in Gafarov *et al* (2018)). The assumption that $0 \leq f < n-1$ rules out the case where $\boldsymbol{q}_1$ (and thus any impulse response to the first shock) is point-identified.[12]

It immediately follows from Proposition 3.2 that a necessary condition for the identified set for $\eta_{1,1,0}$ to exclude zero is that $s + f > n$. In this case, whether it is possible to construct a vector that satisfies $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 = 0$ and the remaining identifying restrictions depends on the reduced-form parameters. Geometrically, the condition $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 = 0$ and the zero restrictions are jointly satisfied when $\boldsymbol{q}_1$ lies in an $(n-f-1)$-dimensional hyperplane that is orthogonal to $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}$ and the rows of $F(\boldsymbol{\phi})$, while the remaining sign restrictions in $S(\boldsymbol{\phi})$ require $\boldsymbol{q}_1$ to lie within the intersection of $s-1$ half-spaces. The identified set for $\eta_{1,1,0}$ will include zero if and only if the intersection of this hyperplane and these half-spaces is nonempty. When $s + f > n$, the hyperplane and half-spaces are not guaranteed to intersect; whether they intersect depends on the values of the reduced-form parameters, which determine the orientations of the hyperplane and half-spaces.

11 I exclude the possibility of imposing sign restrictions with strict inequality (i.e. $S(\boldsymbol{\phi}) > \boldsymbol{0}_{s\times 1}$). In this case, the identified set for $\eta_{1,1,0}$ will be an open interval. Consequently, the identified set for $\tilde{\eta}_{i,1,h}$ could be unbounded without the identified set for $\eta_{1,1,0}$ including zero.

12 When $f = n-1$ and $\mathrm{rank}(F(\boldsymbol{\phi})) = n-1$, the identified set for $\eta_{1,1,0}$ (which is a singleton) excludes zero so long as $\mathrm{rank}((F(\boldsymbol{\phi})', \boldsymbol{\Sigma}'_{tr}\boldsymbol{e}_{1,n,})') = n$. This condition would be violated in the (unrealistic) instance where the zero restrictions in $F(\boldsymbol{\phi})$ include the restriction that $\eta_{1,1,0} = 0$.

This geometric intuition suggests an approach for numerically verifying whether $\eta_{1,1,0}(\boldsymbol{\phi}|S,F)$ includes zero and thus whether $\tilde{\eta}_{i,j,h}(\boldsymbol{\phi}|S,F)$ is unbounded. Let $\tilde{F}(\boldsymbol{\phi},\boldsymbol{Q}) = \mathbf{0}_{(f+1)\times 1}$ represent an augmented set of zero restrictions that includes the sign restriction on $\eta_{1,1,0}$ with equality (i.e. $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 = 0$) and let $\tilde{S}(\boldsymbol{\phi},\boldsymbol{Q}) \geq \mathbf{0}_{(s-1)\times 1}$ collect the remaining sign restrictions. The identified set for $\eta_{1,1,0}$ includes zero if and only if the identified set for $\boldsymbol{Q}$ given the augmented set of restrictions, $Q(\boldsymbol{\phi}|\tilde{S},\tilde{F})$, is nonempty. In the case where the identifying restrictions constrain $\boldsymbol{q}_1$ only, the algorithm proposed in Read (forthcoming) can be used to determine whether $Q(\boldsymbol{\phi}|\tilde{S},\tilde{F})$ is nonempty.[13] When the identifying restrictions constrain multiple columns of $\boldsymbol{Q}$, one can determine whether $Q(\boldsymbol{\phi}|\tilde{S},\tilde{F})$ is nonempty by drawing from a uniform distribution over $Q(\boldsymbol{\phi}|\tilde{F})$ (e.g. using the algorithms in Arias *et al* (2018) or Giacomini and Kitagawa (2021)) until a draw is obtained satisfying the remaining sign restrictions. If no such draw can be obtained after a large number of draws, $Q(\boldsymbol{\phi}|\tilde{S},\tilde{F})$ can be approximated as empty, in which case $\tilde{\eta}_{i,j,h}(\boldsymbol{\phi}|S,F)$ is bounded.

5. Estimating the Effects of a 100 Basis Point Federal Funds Rate Shock

This section illustrates the importance of the issues discussed above by estimating the macroeconomic effects of a 100 basis point shock to the federal funds rate under different sets of identifying restrictions.

I use the reduced-form VAR considered in Uhlig (2005), Antolín-Díaz and Rubio-Ramírez (2018) and Arias *et al* (2019). The model's endogenous variables are real GDP (GDP_t), the GDP deflator ($GDPDEF_t$), a commodity price index (COM_t), total reserves (TR_t), nonborrowed reserves (NBR_t) (all in natural logarithms) and the federal funds rate (FFR_t). The data are monthly and run from January 1965 to November 2007.[14] The VAR includes 12 lags of the variables and a constant. I assume a Jeffreys' (improper) prior over the reduced-form parameters, so $\pi_{\boldsymbol{\phi}} \propto |\boldsymbol{\Sigma}|^{-(n-1)/2}$. This means that the posterior for $\boldsymbol{\phi}$ is a normal-inverse-Wishart distribution, from which it is straightforward to obtain independent draws (e.g. Del Negro and Schorfheide 2011). Under each set of restrictions below, I obtain 10,000 draws from the posterior for $\boldsymbol{\phi}$ such that the identified set is nonempty.[15]

The papers listed above conduct Bayesian inference under a uniform prior for the orthonormal matrix and primarily present impulse responses to a standard-deviation shock.[16] Given concerns around posterior sensitivity to the choice of prior, I instead employ the robust Bayesian approach to inference proposed in Giacomini and Kitagawa (2021). I investigate whether the restrictions are informative about impulse responses to a 100 basis point shock – which are naturally of greater interest to policymakers than the impulse responses to a standard-deviation shock – once the source of posterior sensitivity to the choice of prior is eliminated.

13 If $f = n-2$, the unit-length vector $\tilde{\boldsymbol{q}}_1$ satisfying $\tilde{F}(\boldsymbol{\phi})\tilde{\boldsymbol{q}}_1 = \mathbf{0}_{(f+1)\times 1}$ is pinned down up to sign; such a vector can be found by computing an orthonormal basis for the null space of $\tilde{F}(\boldsymbol{\phi})$. If either $\tilde{S}(\boldsymbol{\phi})\tilde{\boldsymbol{q}}_1 \geq \mathbf{0}_{(s-1)\times 1}$ or $-\tilde{S}(\boldsymbol{\phi})\tilde{\boldsymbol{q}}_1 \geq \mathbf{0}_{(s-1)\times 1}$, then $Q(\boldsymbol{\phi}|\tilde{S},\tilde{F})$ is nonempty. For $0 \leq f < n-2$, the algorithm described in Read (forthcoming) is applicable.

14 I use the updated version of the dataset from Antolín-Díaz and Rubio-Ramírez (2018). The monthly series for GDP_t and $GDPDEF_t$ are obtained by interpolation; see Arias *et al* (2019) for details.

15 When the restrictions constrain a single column of $\boldsymbol{Q}$ only (i.e. under Restrictions (1) and (2)), I check whether the identified set is nonempty at each draw of $\boldsymbol{\phi}$ using Algorithm 4.1 in Read (forthcoming).

16 Uhlig (2005) and Arias *et al* (2019) present impulse responses to a standard-deviation shock. Antolín-Díaz and Rubio-Ramírez (2018) describe their impulse responses as being to a 25 basis point shock, but the normalisation is imposed incorrectly. This is evident from the fact that the credible intervals for the impact impulse response of the federal funds rate fan out around 25 basis points (e.g. see their Figure 6).

First, I consider the identifying restrictions proposed in Arias *et al* (2019), who impose sign and zero restrictions on the structural equation for the federal funds rate, which they interpret as the monetary policy reaction function. The restrictions impose that the coefficients on TR_t and NBR_t in the structural equation for FFR_t are zero, which means that the Federal Reserve does not react to changes in reserves when setting the federal funds rate. They also impose sign restrictions on the coefficients of GDP_t and $GDPDEF_t$ such that the Federal Reserve does not increase the federal funds rate in response to lower output or prices, which is consistent with the types of policy rules typically specified in New Keynesian dynamic stochastic general equilibrium models. Finally, the impact response of FFR_t to the monetary policy shock is restricted to be nonnegative, so that a monetary policy shock always raises FFR_t on impact, which seems natural. I denote this set of identifying restrictions as Restriction (1).

Figure 1: Impulse Responses to 100 Basis Point Shock – Restriction (1)

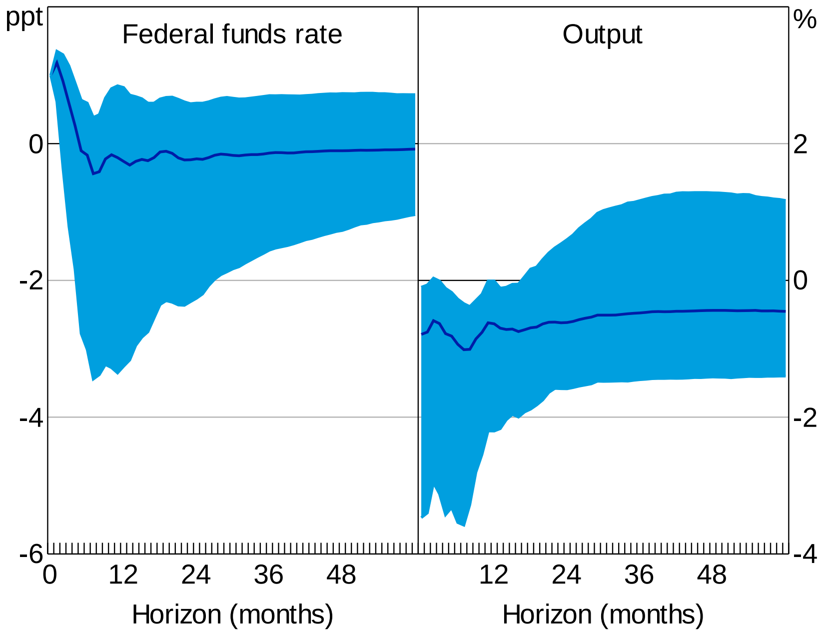

Notes: Blue line is posterior median and blue shading represents 68 per cent equi-tailed credible intervals obtained under the identifying restrictions in Arias *et al* (2019) and a conditionally uniform prior for $\boldsymbol{Q}|\boldsymbol{\phi}$.

Figure 1 presents the impulse responses of FFR_t and GDP_t to a 100 basis point shock obtained under these identifying restrictions and a conditionally uniform prior for $\boldsymbol{Q}$ given $\boldsymbol{\phi}$ (i.e. a standard approach to Bayesian inference). Based on the posterior median, output falls by a maximum of about one per cent. The 68 per cent credible intervals include declines in output of close to 4 per cent, so there is considerable posterior probability assigned to very large declines in output. This set of restrictions involves four sign restrictions (including the sign normalisation on the $(1,1)$ element of $\boldsymbol{A}_0$) and two zero restrictions, so $s + f = 4 + 2 \leq n = 6$. This means that the sufficient condition in Proposition 3.2 is satisfied and the identified sets for impulse responses to a unit monetary policy shock are always unbounded. In turn, this means that the set of posterior means and the sets of posterior τ-quantiles are unbounded for all $\tau \in [0,1]$. In this case, the identifying restrictions cannot rule out the possibility that the federal funds rate does not respond to a monetary policy shock on impact. Consequently, the restrictions are extremely uninformative about the impulse responses to a 100 basis point shock to the federal funds rate. An implication is that the results obtained under the conditionally uniform prior in Figure 1 are driven by the choice of conditional prior rather than information in the data and identifying restrictions.

Next, I combine the restrictions from Arias *et al* (2019) with the sign restrictions on impulse responses proposed in Uhlig (2005). These sign restrictions impose that the impulse response of FFR_{t+h} to the monetary policy shock is nonnegative and the impulse responses of $GDPDEF_{t+h}$, COM_{t+h} and NBR_{t+h} are nonpositive for $h = 0,1,\ldots,5$. I refer to this set of restrictions as Restriction (2). Under a conditionally uniform prior, the additional sign restrictions appreciably tighten the posterior distribution of the impulse responses to a 100 basis point shock (Figure 2). The posterior median suggests that output falls by a maximum of 0.4 per cent about two years after the shock and the 68 per cent credible intervals no longer contain extremely large values; for example, at the two-year horizon the credible intervals span declines in output of 0.1–0.6 per cent. To what extent is the unrevisable component of the prior driving these results?

Under these restrictions, there are two zero restrictions and 27 sign restrictions, so the sufficient condition in Proposition 3.2 is not satisfied and the identified sets for the impulse responses to a unit shock are not necessarily unbounded. I therefore check whether the identified set for $\eta_{1,1,0}$ includes zero at each draw from the reduced-form posterior using a modification of Algorithm 4.1 from Read (forthcoming), which is applicable here because only a single column of $\boldsymbol{Q}$ is restricted. The identified set for the impact response of the cash rate includes zero in only 0.06 per cent of draws from the posterior, which implies that the identified sets for the impulse responses to a unit shock are bounded close to 100 per cent of the time. However, since the identified sets are unbounded with positive probability, the sets of posterior means for the output response to a 100 basis point shock are unbounded. Nevertheless, the set of posterior medians remains bounded because the identified sets are unbounded with low posterior probability. The robust credible intervals are also bounded so long as the credibility level is not too extreme. I approximate the bounds of the identified set at each draw of $\boldsymbol{\phi}$ by obtaining 10,000 draws of $\boldsymbol{Q}$ from a uniform distribution over $\mathcal{Q}(\boldsymbol{\phi}|S,F)$, computing the impulse responses to a unit shock at each draw of $\boldsymbol{Q}$ and taking the minimum and maximum over the draws.[17] To summarise the set of posteriors, I present the set of posterior medians and a 68 per cent robust credible interval.[18]

17 I draw from the uniform distribution over $\mathcal{Q}(\boldsymbol{\phi}|S,F)$ using the Gibbs sampler described in Read (forthcoming), which extends the Gibbs sampler proposed in Amir-Ahmadi and Drautzburg (2021) to additionally allow for zero restrictions.

18 For each horizon and variable of interest, I construct the 68 per cent robust credible interval by computing the 16th percentile of the posterior distribution of the lower bound of the identified set and the 84th percentile of the posterior distribution of the upper bound. Note that this differs to the *shortest* robust credible interval that is proposed in Giacomini and Kitagawa (2021); computing the shortest credible interval requires searching over a grid of possible values, which can be computationally difficult when the identified set is unbounded.

Figure 2: Impulse Responses to 100 Basis Point Shock – Restriction (2)

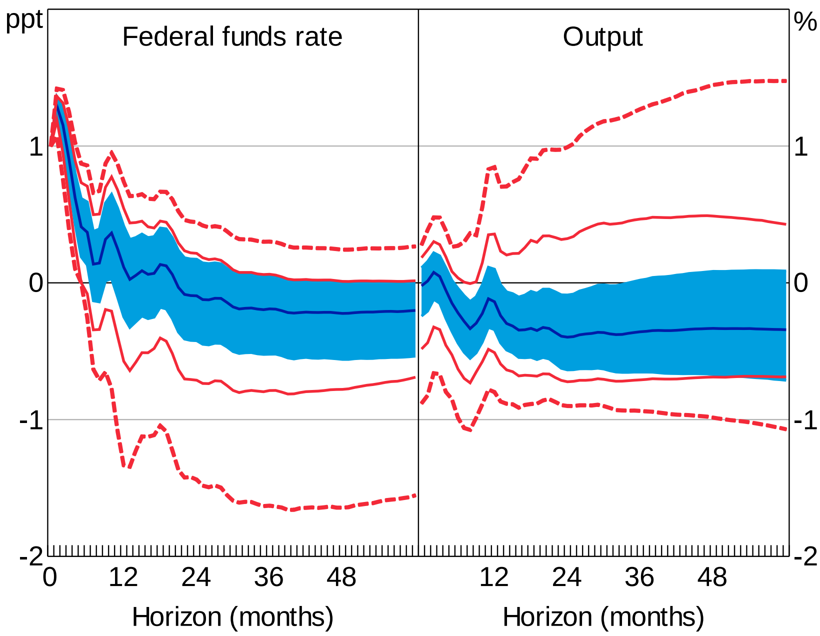

Notes: Blue line is posterior median and blue shading represents 68 per cent equi-tailed credible intervals obtained using a conditionally uniform prior for $\boldsymbol{Q}|\boldsymbol{\phi}$ given a combination of the identifying restrictions in Uhlig (2005) and Arias *et al* (2019); solid red lines represent sets of posterior medians and dashed red lines represent 68 per cent robust credible intervals.

The set of posterior medians for the output response to a 100 basis point shock includes zero at essentially all horizons. The 68 per cent robust credible intervals for the output response include both large negative and large positive responses. Hence, under Restriction (2), there is substantial uncertainty about the output response to a 100 basis point monetary policy shock after eliminating the effect of the unrevisable component of the prior.

Finally, I add the 'narrative restrictions' proposed in Antolín-Díaz and Rubio-Ramírez (2018) to Restriction (2). I refer to this set of restrictions as Restriction (3). Narrative restrictions are restrictions on functions of the structural *shocks* in specific periods (as opposed to restrictions on functions of the structural *parameters*) that represent information about the nature of the shocks hitting the economy during particular historical episodes.[19] The specific narrative restrictions imposed are that the monetary policy shock was positive and was the 'overwhelming' contributor to the observed unexpected change (i.e. the forecast error) in the federal funds rate in October 1979. This is the month in which the Federal Reserve unexpectedly and dramatically raised the federal funds rate following Paul Volcker becoming chairman, and is widely considered an example of a monetary policy shock (e.g. Romer and Romer 1989).

The contribution of the jth structural shock to the one-step-ahead forecast error in variable i in period t is $H_{i,j,t} = \eta_{i,j,0}(\boldsymbol{\phi}, \boldsymbol{Q})\varepsilon_{j,t} = \boldsymbol{e}'_{i,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_j\boldsymbol{q}'_j\boldsymbol{\Sigma}_{tr}^{-1}\boldsymbol{u}_t$. The restriction that the monetary policy shock was the 'overwhelming' contributor to the observed unexpected change in the federal funds rate

19 Giacomini *et al* (2021a) discuss identification and inference under narrative restrictions. They propose extending the robust Bayesian approach to inference from Giacomini and Kitagawa (2021) to this setting to avoid undesirable features of the standard Bayesian approach to inference that arise when using narrative restrictions.

means that the absolute contribution of the monetary policy shock to the forecast error in the federal funds rate is greater than the sum of the absolute contributions of all other shocks, or $|H_{i,j,t}| \geq \sum_{k \neq j} |H_{i,k,t}|$. This is a restriction on the historical decomposition that simultaneously constrains all columns of $\boldsymbol{Q}$. Consequently, it is necessary to numerically approximate whether the identified sets for the impulse responses to a unit shock are unbounded. Following the approach described in Section 4, I use 100,000 draws of $\boldsymbol{Q}$ from a uniform distribution over $\mathcal{Q}(\boldsymbol{\phi}|\tilde{F})$ to check for unboundedness. I approximate the identified set as being empty if I cannot obtain a draw of $\boldsymbol{Q}$ satisfying the identifying restrictions after 100,000 draws. I approximate the bounds of the identified set by obtaining 10,000 draws from a uniform distribution over $\mathcal{Q}(\boldsymbol{\phi}|\boldsymbol{S}, F)$ and computing the minimum and maximum impulse responses over these draws.[20] The numerical methods used to obtain the results under these restrictions are computationally burdensome, so I base the results on 1,000 (rather than 10,000) draws of $\boldsymbol{\phi}$ such that the identified set is nonempty.

Figure 3: Impulse Responses to 100 Basis Point Shock – Restriction (3)

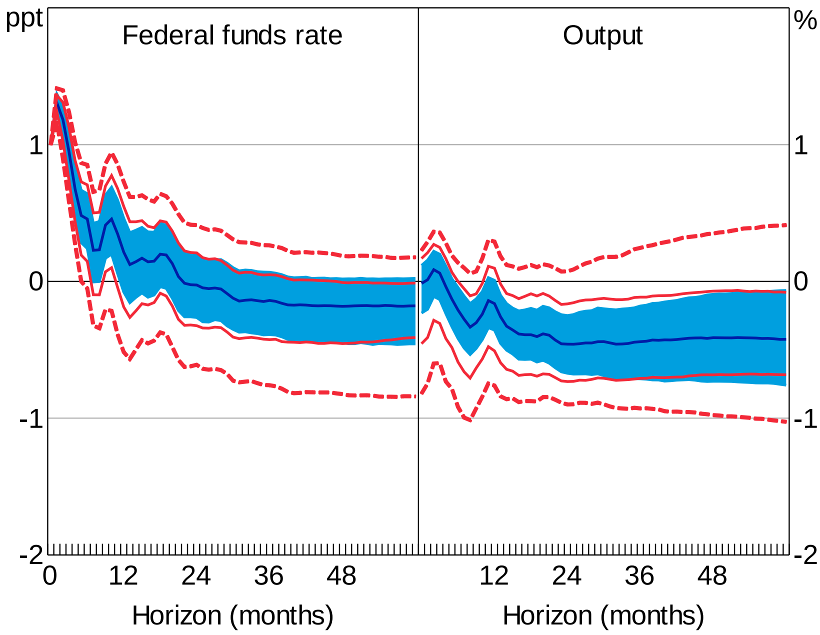

Notes: Blue line is posterior median and blue shading represents 68 per cent equi-tailed credible intervals obtained using a conditionally uniform prior for $\boldsymbol{Q}|\boldsymbol{\phi}$ given a combination of the identifying restrictions in Uhlig (2005), Antolín-Díaz and Rubio-Ramírez (2018) and Arias *et al* (2019); solid red lines represent sets of posterior medians and dashed red lines represent 68 per cent robust credible intervals.

Under Restriction (3), the identified set is bounded in 100 per cent of draws from the reduced-form posterior. Consequently, the sets of posterior means and all posterior quantiles for the impulse responses to a 100 basis point shock are bounded (although, as discussed above, the numerical tools used to check for unboundedness in this case are only approximate). The set of posterior medians excludes zero at most horizons and the 68 per cent robust credible intervals are substantially narrower than under Restriction (2) (Figure 3). Nevertheless, the robust credible

20 Under narrative restrictions, the sign restrictions are functions of the data through the reduced-form VAR innovations that enter the restrictions. Consequently, the standard definition of an identified set does not apply; Giacomini *et al* (2021a) instead introduce the concept of a 'conditional' identified set, which is the identified set that would be obtained after conditioning on the data that directly enter the narrative restrictions. I leave this dependence implicit in the notation and refer to the identified set interchangeably with the conditional identified set.

intervals continue to include zero at all horizons and there remains substantial uncertainty about the output response to a 100 basis point shock.

Table 1 tabulates the posterior lower and upper probabilities that the decline in output is more extreme than a given threshold at selected horizons. Under Restriction (2), the posterior lower and upper probabilities that the output response is negative include both small values and values close to one at all horizons, which indicates that the data and identifying restrictions are fairly uninformative about the sign of the output response. In contrast, under Restriction (3), the posterior lower probability that the output response is negative at the two-year horizon is around 75 per cent and the posterior upper probability of this hypothesis is 100 per cent. The hypothesis that output declines following a positive monetary policy shock therefore receives reasonably high posterior probability uniformly over the class of posteriors that are consistent with the identifying restrictions. Both sets of identifying restrictions effectively rule out relatively large declines in output following a 100 basis point shock; for example, under Restriction (3), the posterior lower probability that output declines by more than 1 per cent two years after the shock is zero and the posterior upper probability is around 5 per cent.

Table 1: Posterior Probability that Decline in Output Exceeds Threshold Following 100 Basis Point Shock[(a)]

Horizon\threshold (%)	Lower probability				Upper probability			
	0	**–0.25**	**–0.5**	**–1**	**0**	**–0.25**	**–0.5**	**–1**
				Restriction (2)				
Impact	0.00	0.00	0.00	0.00	0.95	0.75	0.48	0.10
One year	0.13	0.01	0.00	0.00	0.99	0.84	0.48	0.06
Two years	0.27	0.11	0.03	0.00	1.00	1.00	0.93	0.07
Three years	0.23	0.11	0.04	0.00	1.00	0.99	0.88	0.11
Four years	0.23	0.12	0.05	0.00	1.00	0.98	0.80	0.15
				Restriction (3)				
Impact	0.00	0.00	0.00	0.00	0.95	0.73	0.45	0.04
One year	0.27	0.03	0.00	0.00	0.99	0.84	0.47	0.04
Two years	0.76	0.36	0.08	0.00	1.00	1.00	0.93	0.06
Three years	0.64	0.33	0.11	0.01	1.00	0.99	0.88	0.10
Four years	0.57	0.30	0.12	0.01	1.00	0.98	0.79	0.13

Note: (a) Posterior lower (upper) probability is the smallest (largest) posterior probability obtainable within the class of posteriors consistent with the identifying restrictions.

The existing literature contains a wide range of estimates for the output effects of a 100 basis point shock to the federal funds rate; for example, Ramey (2016) reports a range of existing estimates for the trough in the response of output under different samples, specifications and approaches to identification. These estimates range from as low as 0.6 per cent to as high as 5 per cent. The estimates tend to suggest that the trough in the response of output occurs around two years after the shock, which is consistent with the estimates obtained under Restriction (3). The results under Restriction (3) are broadly consistent with the effects of monetary policy being at the smaller end of the range of existing estimates.

6. Ruling Out Unboundedness Using Alternative Restrictions

In the context of estimating the effects of monetary policy, unboundedness of the impulse responses to a unit shock arises when the identified set for the impact response of the federal funds rate to the monetary policy shock includes zero. A zero value for this impulse response may strike some researchers as implausible. Imposing sign, zero or narrative restrictions of the types considered above can indirectly rule this possibility out. This section discusses alternative restrictions that could potentially be used to rule out the possibility that the monetary policy shock has no impact effect on the federal funds rate. Although this discussion is framed in the context of estimating the effects of monetary policy, it also applies more generally to other settings.

6.1 Direct Bounds on Impulse Responses

One possibility is to directly restrict the impact response of the federal funds rate to the monetary policy shock so that it is greater than some (strictly positive) number (i.e. $\boldsymbol{e}_{1,n}'\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 \geq \lambda$, where $\lambda > 0$ is a specified scalar). However, it seems difficult to justify such restrictions on the basis of economic theory – what is the smallest plausible impact effect of a 'standard-deviation' monetary policy shock on the federal funds rate? Restrictions of this type could potentially be justified on the basis of prior estimates (e.g. from other SVARs or from estimated dynamic stochastic general equilibrium (DSGE) models), but prior estimates may themselves be based on assumptions that lack credibility. Conversely, one could impose bounds on the responses of variables to a unit shock such that unbounded impulse responses are ruled out by assumption. However, it seems similarly difficult to come up with hard bounds on the responses of variables to a 100 basis point shock without these bounds being somewhat arbitrary. Moreover, in either case, when identified sets are unbounded in the absence of such restrictions, inferences may be highly sensitive to changes in the imposed bounds.

To illustrate, return to the bivariate example of Section 3 and consider imposing (in addition to the sign restrictions) the restriction that $\eta_{1,1,0} \geq \lambda$ for some $\lambda > 0$. When $\sigma_{21} \geq 0$ and $0 < \lambda \leq \sigma_{11}$, the identified set for $\tilde{\eta}_{2,1,0}$ is[21]

$$\tilde{\eta}_{2,1,0} \in \left[\frac{\sigma_{21}}{\sigma_{11}} - \frac{\sigma_{22}}{\lambda}\sqrt{\left(1 - \left(\frac{\lambda}{\sigma_{11}}\right)^2\right)}, 0\right].$$

The additional restriction therefore results in the identified set being bounded; in the absence of this restriction (or as λ converges to zero from above), the identified set is $(-\infty, 0]$. However, the lower bound of the identified set is sensitive to the choice of λ, particularly for small values of λ; the derivative of the lower bound tends to ∞ as λ approaches zero from above. Setting λ to some small positive number to rule out an unbounded identified set for $\tilde{\eta}_{2,1,0}$ will therefore yield an identified set that is highly sensitive to the exact choice of λ.

---

21 If $\lambda > \sigma_{11}$, the identified set is empty. See Appendix A.2 for details about this example.

6.2 Bounds on the Forecast Error Variance Decomposition

Rather than directly restricting the impact effect of the monetary policy shock on the federal funds rate, one could instead consider restricting the magnitude of the one-step-ahead forecast error variance decomposition (FEVD) of the federal funds rate with respect to the monetary policy shock. This is the contribution of the monetary policy shock to the one-step-ahead forecast error variance (FEV) of the federal funds rate. For example, Volpicella (forthcoming) proposes imposing bounds on the FEVD, where the bounds are elicited from a range of estimated DSGE models.

Such restrictions may indirectly rule out the possibility that the monetary policy shock has no impact effect on the federal funds rate; intuitively, if one imposes a strictly positive lower bound on the contribution of the monetary policy shock to the one-step-ahead FEV of the federal funds rate, the impact effect of the shock itself must be strictly positive. More formally, the horizon-h FEVD of the ith variable with respect to the jth shock is

$$FEVD_{i,j,h}(\boldsymbol{\phi},\boldsymbol{Q}) = \frac{\sum_{l=0}^{h-1} \boldsymbol{c}_{il}'(\boldsymbol{\phi})\boldsymbol{q}_j\boldsymbol{q}_j\boldsymbol{c}_{il}(\boldsymbol{\phi})}{\sum_{l=0}^{h-1} \boldsymbol{c}_{il}'(\boldsymbol{\phi})\boldsymbol{c}_{il}(\boldsymbol{\phi})}.$$

The impact effect of the jth shock on the ith variable ($\boldsymbol{c}_{i0}'(\boldsymbol{\phi})\boldsymbol{q}_j = \boldsymbol{e}_{i,n}'\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_j$) is zero if and only if $FEVD_{i,j,0}(\boldsymbol{\phi},\boldsymbol{Q}) = 0$, so bounding $FEVD_{i,j,0}(\boldsymbol{\phi},\boldsymbol{Q})$ away from zero indirectly bounds the impact response away from zero. However, if the assumptions underlying the DSGE models that are used to elicit these bounds lack credibility, the derived bounds on the FEVDs will also lack credibility. As in the case where the normalising impulse response is directly bounded away from zero, the identified set obtained under some small lower bound on the FEVD will also be sensitive to the choice of this lower bound when the identified set is unbounded in the absence of this restriction (see Appendix A.3 for an analysis of this case in the context of the bivariate model).

6.3 Priors Over the Structural Parameterisation

Rather than placing a prior over the orthogonal reduced-form parameterisation of the SVAR, Baumeister and Hamilton (2015, 2018, 2019) advocate imposing an explicit prior over the SVAR's structural parameterisation. They argue that the structural parameters (or functions of these parameters, such as impulse responses) have economic interpretations that facilitate prior elicitation. By placing zero prior probability on the hypothesis that the impact response of the federal funds rate to a standard-deviation monetary policy shock is zero, a prior over the structural parameters can rule out the possibility that the impulse responses to a unit shock are unbounded. Alternatively, a prior on the impulse responses to a unit shock could be used to directly rule out this possibility.

As in the case where the prior is over the orthogonal reduced form, one issue with this approach is that a component of the prior will never be updated by the data. For example, in the case where the prior over the structural parameters $\pi_{\boldsymbol{\theta}}$ embeds dogmatic sign or zero restrictions (as in Baumeister and Hamilton (2015)), a prior for the structural parameters can be decomposed into a prior over the reduced-form parameters and a conditional prior for the orthonormal matrix given the reduced-form parameters, $\pi_{\boldsymbol{\theta}} = \pi_{\boldsymbol{\phi}}\pi_{\boldsymbol{Q}|\boldsymbol{\phi}}$. The prior over the structural parameters therefore implicitly assumes a conditional prior over the orthonormal matrix, which is never revised by the data. Consequently, the problem of posterior sensitivity to the choice of prior remains relevant when the prior is over the structural parameterisation.

Although the possibility of posterior sensitivity to the choice of prior is arguably not problematic if the prior reflects credible subjective information about the structural parameters, some features of the prior are typically chosen for convenience rather than as a reflection of any particular prior knowledge. For example, conjugate prior distributions are often used so that (conditional) posterior distributions can be obtained in closed form. Additionally, prior knowledge about parameter values often comes in the form of ranges from previous studies, and this information is not necessarily sufficient to pin down the shape of a prior distribution (i.e. a particular family of distributions). Practitioners may therefore be concerned that somewhat arbitrary assumptions about the prior distribution are driving posterior inference. Hence, it remains valuable to consider the robustness of posterior inferences to the choice of prior even when there exists (partially) credible prior information about the impulse responses.[22] As in the case where the prior is placed over the orthogonal reduced form, the identifying restrictions embedded within a particular prior over the structural parameters may fail to rule out the possibility that the identified sets for the impulse responses to a unit shock are unbounded, in which case the issues raised in this paper are again relevant.

7. Conclusion

In SVARs that are set-identified using sign and/or zero restrictions, the identified set for the impulse responses to a unit shock may be unbounded. This raises complications when conducting robust Bayesian (or frequentist) inference about these impulse responses. However, it may still be possible to draw useful inferences about impulse responses to a unit shock when the identified set is unbounded with positive probability. I develop an easily verifiable sufficient condition for assessing whether the identified set for the impulse response to a unit shock is unbounded. When the sufficient condition is not satisfied, I describe how to numerically check whether the identified set is unbounded.

The empirical exercise in this paper demonstrates the importance of these issues. Under the identifying restrictions considered in Arias *et al* (2019), the identified set for the impulse response to a 100 basis point monetary policy shock is unbounded at all horizons and for all values of the reduced-form parameters. The identifying restrictions are therefore extremely uninformative about the magnitude of these impulse responses, and standard Bayesian inference will be very misleading about the information contained in the data and identifying restrictions. After adding the sign restrictions on impulse responses from Uhlig (2005), the identified set is bounded with high posterior probability. Additionally adding the narrative restrictions from Antolín-Díaz and Rubio-Ramírez (2018) yields a bounded identified set in 100 per cent of draws from the reduced-form posterior. The results under the latter two sets of restrictions are broadly consistent with the effects of US monetary policy being at the smaller end of the range of existing estimates.

---

22 Giacomini, Kitagawa and Uhlig (2019) discuss these issues and propose an approach that can be used to assess posterior sensitivity to the choice of prior when the researcher possesses a partially credible prior for set-identified parameters. See also Giacomini *et al* (2021b) for a discussion of this and other approaches to robust Bayesian inference.

Appendix A: Derivations for Bivariate SVAR

A.1 Sign Restrictions on Impulse Responses

In this appendix, I derive the identified sets for the impulse responses to a unit shock under the sign restrictions on impulse responses presented in Section 3.

In the absence of any identifying restrictions, the identified set for $\boldsymbol{A}_0^{-1}$ (the matrix of impact impulse responses) is

$$\boldsymbol{A}_0^{-1} \in \left\{\begin{bmatrix} \sigma_{11}\cos\theta & -\sigma_{11}\sin\theta \\ \sigma_{21}\cos\theta + \sigma_{22}\sin\theta & \sigma_{22}\cos\theta - \sigma_{21}\sin\theta \end{bmatrix}\right\} \cup \left\{\begin{bmatrix} \sigma_{11}\cos\theta & \sigma_{11}\sin\theta \\ \sigma_{21}\cos\theta + \sigma_{22}\sin\theta & \sigma_{21}\sin\theta - \sigma_{22}\cos\theta \end{bmatrix}\right\},$$

and the identified set for $\boldsymbol{A}_0$ is

$$\boldsymbol{A}_0 \in \left\{\frac{1}{\sigma_{11}\sigma_{22}}\begin{bmatrix} \sigma_{22}\cos\theta - \sigma_{21}\sin\theta & \sigma_{11}\sin\theta \\ -\sigma_{21}\cos\theta - \sigma_{22}\sin\theta & \sigma_{11}\cos\theta \end{bmatrix}\right\} \cup \left\{\frac{1}{\sigma_{11}\sigma_{22}}\begin{bmatrix} \sigma_{22}\cos\theta - \sigma_{21}\sin\theta & \sigma_{11}\sin\theta \\ \sigma_{21}\cos\theta + \sigma_{22}\sin\theta & -\sigma_{11}\cos\theta \end{bmatrix}\right\}.$$

The identified set for the impact response of the second variable to a shock that raises the first variable by one unit is

$$\tilde{\eta}_{2,1,0} = \frac{\eta_{2,1,0}}{\eta_{1,1,0}} = \frac{\sigma_{21}\cos\theta + \sigma_{22}\sin\theta}{\sigma_{11}\cos\theta} = \frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}}{\sigma_{11}}\tan\theta.$$

Under the sign restrictions on impulse responses described in Section 3 (including the sign normalisation), the parameter θ is restricted to lie within the following set:

$$\theta \in \{\theta: \sigma_{11}\cos\theta \geq 0, \sigma_{21}\cos\theta \leq -\sigma_{22}\sin\theta, \sigma_{22}\cos\theta \geq \sigma_{21}\sin\theta\}$$

$$\cup\, \{\theta: \sigma_{11}\cos\theta \geq 0, \sigma_{21}\cos\theta \leq -\sigma_{22}\sin\theta, \sigma_{22}\cos\theta \geq \sigma_{21}\sin\theta, -\sigma_{11}\cos\theta \geq 0\}.$$

There are two cases to consider depending on the sign of σ_{21}. If $\sigma_{21} < 0$, the second set is empty. The first set is equivalent to

$$\left\{\theta:\ \cos\theta > 0, \tan\theta \leq -\frac{\sigma_{21}}{\sigma_{22}}, \tan\theta \geq \frac{\sigma_{22}}{\sigma_{21}}\right\}.$$

This set of inequalities implies that the identified set for θ is

$$\theta \in \left[\arctan\left(\frac{\sigma_{22}}{\sigma_{21}}\right), \arctan\left(-\frac{\sigma_{21}}{\sigma_{22}}\right)\right].$$

The identified set for the impact response of the first variable to the first shock, $\eta_{11} = \sigma_{11}\cos\theta$, is then

$$\eta_{1,1,0} \in \left[\sigma_{11}\cos\left(\arctan\left(\min\left\{\frac{\sigma_{22}}{\sigma_{21}}, \frac{\sigma_{21}}{\sigma_{22}}\right\}\right)\right), \sigma_{11}\right].$$

The identified set for this impulse response is bounded away from zero. In this case, the identified set for $\tilde{\eta}_{2,1,0}$ is

$$\tilde{\eta}_{2,1,0} \in \left[\frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}^2}{\sigma_{11}\sigma_{21}}, 0\right],$$

which is bounded.

Similarly, if $\sigma_{21} > 0$, θ is restricted to lie in the set

$$\theta \in \left\{\theta \colon \cos\theta > 0, \tan\theta \leq -\frac{\sigma_{21}}{\sigma_{22}}, \tan\theta \leq \frac{\sigma_{22}}{\sigma_{21}}\right\} \cup \left\{-\frac{\pi}{2}\right\}.$$

The second inequality implies that $\tan\theta \leq 0$, so the last inequality never binds. The identified set for θ is therefore

$$\theta \in \left[-\frac{\pi}{2}, \arctan\left(-\frac{\sigma_{21}}{\sigma_{22}}\right)\right],$$

and the identified set for $\eta_{1,1,0}$ is

$$\eta_{1,1,0} \in \left[0, \sigma_{11}\cos\left(\arctan\left(-\frac{\sigma_{21}}{\sigma_{22}}\right)\right)\right].$$

If $\sigma_{21} = 0$, θ is restricted to the set

$$\theta \in \{\theta\colon\ \cos\theta \geq 0, 0 \leq -\sigma_{22}\sin\theta\} \cup \{\theta\colon 0 \leq -\sigma_{22}\sin\theta, \cos\theta \geq 0, -\sigma_{11}\cos\theta \geq 0\}.$$

The first set implies $\theta \in [-\pi/2, 0]$ and the second implies $\theta = -\pi/2$, so $\eta_{1,1,0} \in [0, \sigma_{11}]$. The expression for the identified set for $\eta_{1,1,0}$ when $\sigma_{21} > 0$ therefore also applies when $\sigma_{21} = 0$. $\tan\theta \to -\infty$ as θ approaches $-\pi/2$ from above. $\tan\theta$ is strictly increasing over the identified set for θ, so the upper bound for the identified set for $\tilde{\eta}_{2,1,0}$ is obtained by evaluating $\tilde{\eta}_{2,1,0}$ at the upper bound of the identified set for θ. Consequently, $\tilde{\eta}_{2,1,0} \in (-\infty, 0]$.

A.2 Magnitude Restrictions

In addition to the sign restrictions considered in the previous section, consider the restriction that $\eta_{1,1,0} \geq \lambda$ for some $\lambda > 0$. Under this set of restrictions, θ is restricted to lie within the following set:

$$\theta \in \{\theta\colon \sigma_{11}\cos\theta \geq \lambda, \sigma_{21}\cos\theta \leq -\sigma_{22}\sin\theta, \sigma_{22}\cos\theta \geq \sigma_{21}\sin\theta\}$$

$$\cup \{\theta\colon \sigma_{11}\cos\theta \geq \lambda, \sigma_{21}\cos\theta \leq -\sigma_{22}\sin\theta, \sigma_{22}\cos\theta \geq \sigma_{21}\sin\theta, -\sigma_{11}\cos\theta \geq 0\}.$$

The second set is always empty, since $\sigma_{11}\cos\theta \geq \lambda$ and $-\sigma_{11}\cos\theta \geq 0$ cannot hold simultaneously when $\lambda > 0$. The identified set for θ is empty if $\lambda > \sigma_{11}$, since $\cos\theta \leq 1$ for all θ.

If $\sigma_{21} \geq 0$, the first set is equivalent to

$$\theta \in \left\{\theta: \cos\theta \geq \frac{\lambda}{\sigma_{11}}, \tan\theta \leq -\frac{\sigma_{21}}{\sigma_{22}}, \tan\theta \leq \frac{\sigma_{22}}{\sigma_{21}}\right\}.$$

The last inequality never binds and the identified set for θ is

$$\theta \in \left[-\arccos\left(\frac{\lambda}{\sigma_{11}}\right), \arctan\left(-\frac{\sigma_{21}}{\sigma_{22}}\right)\right].$$

The identified set for $\tilde{\eta}_{2,1,0}$ is therefore

$$\tilde{\eta}_{2,1,0} \in \left[\frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}}{\sigma_{11}}\tan\left(-\arccos\left(\frac{\lambda}{\sigma_{11}}\right)\right), 0\right].$$

The lower bound of this identified set, $\ell(\boldsymbol{\phi}, \lambda)$, can be expressed as

$$\ell(\boldsymbol{\phi}, \lambda) = \frac{\sigma_{21}}{\sigma_{11}} - \frac{\sigma_{22}}{\lambda}\sqrt{\left(1 - \left(\frac{\lambda}{\sigma_{11}}\right)^2\right)},$$

which converges to $-\infty$ as λ approaches zero from above. The derivative of $\ell(\boldsymbol{\phi}, \lambda)$ with respect to λ is

$$\frac{\partial \ell(\boldsymbol{\phi}, \lambda)}{\partial \lambda} = \frac{\left(1 - \left(\frac{\lambda}{\sigma_{11}}\right)^2\right)^{-\frac{1}{2}}}{\lambda^2}.$$

In the limit as λ approaches zero from above, this derivative approaches ∞, which implies that the lower bound is extremely sensitive to small changes in λ when λ is close to zero.

A.3 Bounds on the FEVD

The FEV of of y_{1t} is σ_{11}^2 and the contribution of ε_{1t} to the FEV of y_{1t} is $\sigma_{11}^2 \cos^2\theta$. The FEVD of y_{1t} with respect to ε_{1t}, $FEVD_{\varepsilon_{1t}}^{y_{1t}}$, is therefore $\cos^2\theta$. Consider imposing the restriction that $FEVD_{\varepsilon_{1t}}^{y_{1t}} \geq \kappa$ for some $0 < \kappa \leq 1$ in addition to the sign restrictions considered in Section 3 and A.1. Under this set of restrictions, θ is restricted to lie within the following set:

$$\theta \in \{\theta: \sigma_{11}\cos\theta \geq 0, \sigma_{21}\cos\theta \leq -\sigma_{22}\sin\theta, \sigma_{22}\cos\theta \geq \sigma_{21}\sin\theta, \cos^2\theta \geq \kappa\}$$

$$\cup \{\theta: \sigma_{11}\cos\theta \geq 0, \sigma_{21}\cos\theta \leq -\sigma_{22}\sin\theta, \sigma_{22}\cos\theta \geq \sigma_{21}\sin\theta, -\sigma_{11}\cos\theta \geq 0, \cos^2\theta \geq \kappa\}.$$

When $\sigma_{21} \geq 0$, the first set is equivalent to

$$\theta \in \left\{\theta: \cos\theta > 0, \tan\theta \leq -\frac{\sigma_{21}}{\sigma_{22}}, \tan\theta \leq \frac{\sigma_{22}}{\sigma_{21}}, -\arccos\sqrt{\kappa} \leq \theta \leq \arccos\sqrt{\kappa}\right\}.$$

The inequalities $\tan\theta \leq \sigma_{22}/\sigma_{21}$ and $\theta \leq \arccos\sqrt{\kappa}$ never bind and the identified set for θ is

$$\theta \in \left[-\text{arccos}\sqrt{\kappa}, \text{arctan}\left(-\frac{\sigma_{21}}{\sigma_{22}}\right)\right].$$

The identified set for $\tilde{\eta}_{2,1,0}$ is therefore

$$\tilde{\eta}_{2,1,0} \in \left[\frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}}{\sigma_{11}}\tan\left(-\text{arccos}\left(\sqrt{\kappa}\right)\right), 0\right].$$

The lower bound of this identified set, $\ell(\boldsymbol{\phi}, \kappa)$, can be expressed as

$$l(\boldsymbol{\phi}, \kappa) = \frac{\sigma_{21}}{\sigma_{11}} - \frac{\sigma_{22}}{\sigma_{11}}\frac{\sqrt{1-\kappa}}{\sqrt{\kappa}}.$$

The lower bound converges to $-\infty$ as κ approaches zero from above. The derivative of $\ell(\boldsymbol{\phi}, \kappa)$ with respect to κ is

$$\frac{\partial \ell(\boldsymbol{\phi}, \kappa)}{\partial \kappa} = \frac{1}{2}\frac{\sigma_{22}}{\sigma_{11}}\kappa^{-\frac{3}{2}}(1-\kappa)^{-\frac{1}{2}}.$$

In the limit as κ approaches zero from above, this derivative approaches ∞, which implies that the lower bound is extremely sensitive to small changes in κ when κ is close to zero.

To summarise, under the additional restriction on the FEVD, the identified set is bounded; in the absence of this restriction (or as κ converges to zero from above), the identified set is $(-\infty, 0]$. However, as in the case where the normalising impulse response is directly bounded away from zero, the lower bound of the identified set is sensitive to the choice of κ, particularly for small values of κ; the derivative of the lower bound tends to ∞ as κ approaches zero from above. Setting κ to some small positive number to rule out an unbounded identified set for $\tilde{\eta}_{2,1,0}$ will therefore yield an identified set that is highly sensitive to the exact choice of κ.

Appendix B: Proofs of Propositions

Proof of Proposition 3.1. Assume $\eta_{1,1,0}(\boldsymbol{\phi}|S,F)$ does not include zero, so $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 > 0$ for any $\boldsymbol{Q} \in \mathcal{Q}(\boldsymbol{\phi}|S,F)$. There exists $\delta > 0$ such that $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 > \delta$ for all $\boldsymbol{Q} \in \mathcal{Q}(\boldsymbol{\phi}|S,F)$. Under the assumption that the reduced-form parameter space $\boldsymbol{\Phi}$ is such that the VMA(∞) representation of the VAR exists, $|\eta_{i,j,h}(\boldsymbol{\phi},\boldsymbol{Q})| < \infty$ for all $\boldsymbol{\phi} \in \boldsymbol{\Phi}$ and $\boldsymbol{Q} \in \mathcal{Q}(\boldsymbol{\phi}|S,F)$. There thus exists $\kappa < \infty$ such that $|\eta_{i,1,h}(\boldsymbol{\phi},\boldsymbol{Q})| < \kappa$ for all $\boldsymbol{\phi} \in \boldsymbol{\Phi}$ and $\boldsymbol{Q} \in \mathcal{Q}(\boldsymbol{\phi}|S,F)$. It follows that $|\tilde{\eta}_{i,1,h}(\boldsymbol{\phi},\boldsymbol{Q})| < \frac{\kappa}{\delta} < \infty$ for all $\boldsymbol{\phi} \in \boldsymbol{\Phi}$ and $\boldsymbol{Q} \in \mathcal{Q}(\boldsymbol{\phi}|S,F)$, so the lower and upper bounds of $\tilde{\eta}_{i,1,h}(\boldsymbol{\phi}|S,F)$ must be bounded.

Proof of Proposition 3.2. Assume that the sign restrictions are ordered such that the first row of $S(\boldsymbol{\phi})$ is $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}$. Since $\boldsymbol{\Sigma}_{tr}$ is lower triangular, the condition $\boldsymbol{e}'_{1,n}\boldsymbol{\Sigma}_{tr}\boldsymbol{q}_1 = 0$ is satisfied only for values of $\boldsymbol{q}_1$ such that $\boldsymbol{q}_1 = (0, \boldsymbol{q}'_{1,2:n})'$, where $\boldsymbol{q}_{1,2:n}$ is an $(n-1)$-dimensional vector.[23] For such a value of $\boldsymbol{q}_1$, the entries in the first columns of $F(\boldsymbol{\phi})$ and $S(\boldsymbol{\phi})$ do not enter the equality restrictions in $F(\boldsymbol{\phi})\boldsymbol{q}_1 = \boldsymbol{0}_{r\times 1}$ and the last $s-1$ inequalities in $S(\boldsymbol{\phi})\boldsymbol{q}_1 \geq \boldsymbol{0}_{s\times 1}$, respectively. Let $\ddot{F}(\boldsymbol{\phi})$ represent the matrix of coefficients in the zero restrictions after dropping the first column and let $\ddot{S}(\boldsymbol{\phi})$ represent the matrix of coefficients in the sign restrictions after dropping the first row and column. According to Proposition 3.1 of Read (forthcoming), the system of sign and zero restrictions in $\mathbb{R}^{n-1}$, $\ddot{F}(\boldsymbol{\phi})\boldsymbol{q}_{1,2:n} = \boldsymbol{0}_{f\times 1}$ and $\ddot{S}(\boldsymbol{\phi})\boldsymbol{q}_{1,2:n} \geq \boldsymbol{0}_{(s-1)\times 1}$, can be expressed as an equivalent system of sign restrictions in $\mathbb{R}^{n-f-1}$. Let $\check{S}(\boldsymbol{\phi})\check{\boldsymbol{q}} = \boldsymbol{0}_{(s-1)\times 1}$ represent the transformed system of sign restrictions, where $\check{\boldsymbol{q}} \in \mathbb{R}^{n-f-1}$ and $\check{S}(\boldsymbol{\phi})$ is obtained from $\ddot{S}(\boldsymbol{\phi})$ using the transformation described in Read (forthcoming). Corollary 3.1 of Read (forthcoming) implies that the identified set for $\eta_{1,1,0}$ will include zero if and only if there exists $\check{\boldsymbol{q}}$ satisfying $\check{S}(\boldsymbol{\phi})\check{\boldsymbol{q}} \geq \boldsymbol{0}_{(s-1)\times 1}$. In what follows, I show that such a vector always exists under the assumptions of the proposition.

Consider the case where $\mathrm{rank}\left(\check{S}(\boldsymbol{\phi})\right) = s-1$. By Gordan's Theorem (e.g. Mangasarian 1969; Border 2020), either $\check{S}(\boldsymbol{\phi})\boldsymbol{x} > \boldsymbol{0}_{(s-1)\times 1}$ for some $\boldsymbol{x} \in \mathbb{R}^{n-f-1}$ or $\check{S}(\boldsymbol{\phi})'\boldsymbol{y} = \boldsymbol{0}_{(n-r-1)\times 1}$ for some $\boldsymbol{y} \neq \boldsymbol{0}_{(s-1)\times 1}$. Since $\check{S}(\boldsymbol{\phi})$ has full rank, there cannot exist a $\boldsymbol{y} \neq \boldsymbol{0}_{(s-1)\times 1}$ such that $\check{S}(\boldsymbol{\phi})'\boldsymbol{y} = \boldsymbol{0}_{(n-f-1)\times 1}$, so there must exist $\check{\boldsymbol{q}}$ satisfying $\check{S}(\boldsymbol{\phi})\check{\boldsymbol{q}} > \boldsymbol{0}_{(s-1)\times 1}$. Any $\check{\boldsymbol{q}}$ satisfying $\check{S}(\boldsymbol{\phi})\check{\boldsymbol{q}} > \boldsymbol{0}_{(s-1)\times 1}$ also satisfies $\check{S}(\boldsymbol{\phi})\check{\boldsymbol{q}} \geq \boldsymbol{0}_{(s-1)\times 1}$. Next, consider the case where $\mathrm{rank}\left(\check{S}(\boldsymbol{\phi})\right) < s-1$ and let $N(\check{S}(\boldsymbol{\phi}))$ represent an orthonormal basis for the null space of $\check{S}(\boldsymbol{\phi})$. By the rank-nullity theorem, $N(\check{S}(\boldsymbol{\phi}))$ has dimension $(n-f-1) - \mathrm{rank}\left(\check{S}(\boldsymbol{\phi})\right)$. Since $\mathrm{rank}\left(\check{S}(\boldsymbol{\phi})\right) < s-1$ and $s+f \leq n$, $N(\check{S}(\boldsymbol{\phi}))$ will have dimension strictly greater than one. Thus, when $\mathrm{rank}\left(\check{S}(\boldsymbol{\phi})\right) < s-1$, it is always possible to construct a unit-length vector satisfying $\check{S}(\boldsymbol{\phi})\check{\boldsymbol{q}} = \boldsymbol{0}_{(s-1)\times 1}$ by taking any column of $N(\check{S}(\boldsymbol{\phi}))$. Such a vector clearly satisfies $\check{S}(\boldsymbol{\phi})\check{\boldsymbol{q}} \geq \boldsymbol{0}_{(s-1)\times 1}$. □

23 This assumes that the (1,1) element of $\boldsymbol{\Sigma}_{tr}$ is nonzero, which is guaranteed so long as $\boldsymbol{\Sigma}$ is nonsingular.

References

Amir-Ahmadi P and T Drautzburg (2021), 'Identification and Inference with Ranking Restrictions', *Quantitative Economics*, 12(1), 1–39.

Antolín-Díaz J and JF Rubio-Ramírez (2018), 'Narrative Sign Restrictions for SVARs', *American Economic Review*, 108(10), pp 2802–29.

Arias JE, D Caldara and JF Rubio-Ramírez (2019), 'The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure', *Journal of Monetary Economics*, 101, pp 1–13.

Arias JE, JF Rubio-Ramírez and DF Waggoner (2018), 'Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications, *Econometrica*, 86(2), pp 685–720.

Bacchiocchi E and T Kitagawa (2021), 'A Note on Global Identification in Structural Vector Autoregressions', cemmap working paper CWP03/21.

Baumeister C and JD Hamilton (2015), 'Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information', *Econometrica*, 83(5), 1963–1999.

Baumeister C and JD Hamilton (2018), 'Inference in Structural Vector Autoregression When the Identifying Assumptions Are Not Fully Believed: Re-evaluating the Role of Monetary Policy in Economic Fluctuations', 100, 48–65.

Baumeister C and JD Hamilton (2019), 'Structural Interpretation of Vector Autoregressions with Incomplete Identification: Revisiting the Role of Oil Supply and Demand Shocks', *American Economic Review*, 109(5), pp 1873–1910.

Baumeister C and JD Hamilton (2022), 'Advances in Using Vector Autoregressions to Estimate Structural Magnitudes', Unpublished Manuscript.

Border KC (2020), 'Alternative Linear Inequalities', Unpublished manuscript, California Institute of Technology, available at <https://healy.econ.ohio-state.edu/kcb/Notes/Alternative.pdf>.

Del Negro M and F Schorfheide (2011), 'Bayesian Macroeconometrics', in Geweke J, G Koop and H Van Dijk (eds), Oxford Handbook of Bayesian Econometrics, Oxford University Press, Oxford, pp 293–389.

Gafarov B, M Meier and JL Montiel Olea (2018), 'Delta-method Inference for a Class of Set-identified SVARs', *Journal of Econometrics*, 203(2), 316–237.

Giacomini R and T Kitagawa (2020), 'Robust Bayesian Inference for Set-Identified Models', cemmap working paper CWP12/20.

Giacomini R and T Kitagawa (2021), 'Robust Bayesian Inference for Set-Identified Models', *Econometrica*, 89(4), pp 1519–1556.

Giacomini R, T Kitagawa and M Read (2021a), 'Identification and Inference Under Narrative Restrictions', arXiv 2102:06456 [econ.EM].

Giacomini R, T Kitagawa and M Read (2021b), 'Robust Bayesian Analysis for Econometrics', Centre for Economic Policy Research Discussion Paper DP16488.

Giacomini R, T Kitagawa and M Read (2022), 'Robust Bayesian Inference in Proxy SVARs', *Journal of Econometrics*, 228(1), pp 107–126.

Giacomini R, T Kitagawa and H Uhlig (2019), 'Estimation Under Ambiguity', cemmap working paper CWP24/19.

Granziera E, HR Moon and F Schorfheide (2018), 'Inference for VARs Identified with Sign Restrictions', *Quantitative Economics*, 9(3), pp 1087–1121.

Hamilton JD (1994), *Time Series Analysis*, Princeton University Press, Princeton.

Kilian L and H Lütkepohl (2017), *Structural Vector Autoregressive Analysis*, Cambridge University Press, Cambridge.

Mangasarian OL (1969), '*Nonlinear Programming',* McGraw-Hill, New York.

Poirier DJ (1998), 'Revising Beliefs in Nonidentified Models', *Econometric Theory*, 14(4), pp 483–509.

Ramey VA (2016), 'Macroeconomic Shocks and Their Propagation', in JB Taylor and H Uhlig (eds), Handbook of Macroeconomics: Volume 2A, Handbooks in Economics, Elsevier, Amsterdam, pp 71–162.

Read M (forthcoming), 'Algorithms for Inference in SVARs Identified with Sign and Zero Restrictions', *The Econometrics Journal*.

Romer CD and DH Romer (1989), 'Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz', in OJ Blanchard and S Fischer (eds), NBER Macroeconomics Annual, Volume 4, MIT Press, Cambridge, pp 121–84.

Rubio-Ramírez JF, DF Waggoner and T Zha (2010), 'Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference', *Review of Economic Studies*, 77(2), pp 665–696.

Stock JH and MW Watson (2016), 'Dynamic Factor Models, Factor-augmented Vector Autoregressions and Structural Vector Autoregressions in Macroeconomics', in Taylor JB and H Uhlig (eds), Handbook of Macroeconomics: Volume 2A, Handbooks in Economics, Elsevier, Amsterdam, pp 415–525.

Stock JH and MW Watson (2018), 'Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments', *The Economic Journal*, 128(610), 917–948.

Uhlig H (2005), 'What Are the Effects of Monetary Policy On Output? Results From an Agnostic Identification Procedure', *Journal of Monetary Economics*, 52(2), pp 381–419.

Uhlig H (2017), 'Shocks, Sign Restrictions, and Identification', In B Honore, A Pakes, M Piazzesi and L Samuelson (eds), *Advances in Economics and Econometrics: Eleventh World Congress*, Econometric Society Monographs, Cambridge University Press, Cambridge, pp 95–127.

Volpicella A (forthcoming), 'SVARs Identification Through Bounds on the Forecast Error Variance', *Journal of Business and Economic Statistics*.