

Narrative restrictions and proxies

Raffaella Giacomini
Toru Kitagawa
Matthew Read

The Institute for Fiscal Studies
Department of Economics, UCL

cemmap working paper CWP09/22



Narrative Restrictions and Proxies*

Raffaella Giacomini[†], Toru Kitagawa[‡] and Matthew Read[§]

April 6, 2022

Abstract

We compare two approaches to using information about the signs of structural shocks at specific dates within a structural vector autoregression (SVAR): imposing ‘narrative restrictions’ (NR) on the shock signs in an otherwise set-identified SVAR; and casting the information about the shock signs as a discrete-valued ‘narrative proxy’ (NP) to point-identify the impulse responses. The NP is likely to be ‘weak’ given that the sign of the shock is typically known in a small number of periods, in which case the weak-proxy robust confidence intervals in Montiel-Olea et al. (2021) are the natural approach to conducting inference. However, we show both theoretically and via Monte Carlo simulations that these confidence intervals have distorted coverage – which may be higher or lower than the nominal level – unless the sign of the shock is known in a large number of periods. Regarding the NR approach, we show that the prior-robust Bayesian credible intervals from Giacomini et al. (2021a) deliver coverage exceeding the nominal level, but which converges towards the nominal level as the number of NR increases.

*This paper was prepared for the *Journal of Business and Economic Statistics* Session of the 2022 North American Winter Meeting of the Econometric Society. We thank Lutz Killian, Mikkel Plagborg-Møller and Juan Rubio-Ramírez for their illuminating discussions of our paper, and Chris Hansen for organising the session. We also thank James Stock, Christian Wolf and seminar participants at various venues for discussions that motivated this work. We gratefully acknowledge financial support from ERC grants (numbers 536284 and 715940) and the ESRC Centre for Microdata Methods and Practice (CeMMAP) (grant number RES-589-28-0001). The views in this paper are the authors’ and do not reflect the views of the Federal Reserve Bank of Chicago, the Federal Reserve System or the Reserve Bank of Australia.

[†]University College London, Department of Economics and Federal Reserve Bank of Chicago. Email: r.giacomini@ucl.ac.uk

[‡]Brown University, Department of Economics and University College London, Department of Economics. Email: toru_kitagawa@brown.edu

[§]University College London, Department of Economics and Reserve Bank of Australia. Email: readm@rba.gov.au

1 Introduction

Structural vector autoregressions (SVARs) are used in macroeconomics to estimate the dynamic causal effects of structural shocks. A common approach to identifying the effects of these shocks is to impose a set of sign and/or zero restrictions on functions of the SVAR’s structural parameters that together result in set-identification of the parameters of interest (e.g. sign restrictions on impulse responses, as in Uhlig (2005)). A growing number of papers augment these ‘traditional’ set-identifying restrictions with restrictions that involve the values of the structural shocks in specific periods. For instance, Antolín-Díaz and Rubio-Ramírez (2018) propose restricting the signs of structural shocks based on the historical narrative about the nature of the shocks hitting the economy in particular episodes. An example is the restriction that there was a positive monetary policy shock in the United States in October 1979, which is the year in which the Federal Reserve dramatically raised the federal funds rate following Paul Volcker becoming chairman.¹

In this paper, we compare two alternative approaches to using information about the signs of structural shocks in specific periods in an SVAR framework. The first approach is to impose the information as restrictions on the signs of the structural shocks in an otherwise set-identified SVAR, as in Antolín-Díaz and Rubio-Ramírez (2018); we refer to this as ‘shock-sign narrative restrictions’ (NR). The second approach follows a suggestion in Plagborg-Møller and Wolf (2021a) that the information about the shock signs could be recast as an external instrument or ‘proxy’ for use in a proxy SVAR (e.g. Mertens and Ravn 2013; Stock and Watson 2018).² We refer to this as the ‘narrative-proxy’ (NP) approach.

Following Antolín-Díaz and Rubio-Ramírez (2018), the literature that makes use of NR typically imposes these restrictions within a Bayesian framework.³ This involves specifying a uniform-normal-inverse-Wishart prior over the orthogonal reduced-form parameterisation of the SVAR, which is the standard approach in set-identified SVARs

¹Other papers that impose NR include Ben Zeev (2018), Ludvigson, Ma and Ng (2018, 2021), Furlanetto and Robstad (2019), Cheng and Yang (2020), Kilian and Zhou (2020a, 2020b), Laumer (2020), Redl (2020), Zhou (2020), Caggiano et al. (2021), Maffei-Faccioli and Vella (2021), Berger et al. (2022) and Inoue and Kilian (in press).

²More precisely, the possibility of using narrative information to construct a proxy variable is discussed in the Supplement to Plagborg-Møller and Wolf (2021a); see Plagborg-Møller and Wolf (2021b).

³An exception is Ludvigson et al. (2018, 2021), who use a bootstrap to conduct inference.

(e.g. Arias et al. 2018). However, Giacomini et al. (2021a) – henceforth, GKR – point out some undesirable features of this approach. Under shock-sign restrictions, the likelihood function possesses flat regions, which implies that a component of the prior is never updated by the data. Furthermore, as discussed by Baumeister and Hamilton (2015) in the context of standard sign restrictions, the prior is informative about parameters of interest, such as impulse responses. Given that the prior is primarily chosen for computational convenience and does not reflect credible prior information about the parameters, these features of standard Bayesian inference under NR raise the concern that posterior inference may be mainly driven by the choice of prior.⁴

To address the issue of posterior sensitivity to the choice of prior under NR, GKR propose applying a variant of the ‘robust’ (multiple-prior) Bayesian approach to inference for set-identified models developed in Giacomini and Kitagawa (2021). This involves replacing the unrevisable component of the prior with a class of priors that are consistent with the identifying restrictions. The class of priors generates a class of posteriors, which can be summarised in various ways. For example, rather than having a single posterior mean, there is a *set* of posterior means, which is an interval containing every possible posterior mean that could be obtained under the class of priors. One can also report a ‘robust credible interval’, which is an interval that receives at least a given posterior probability under every posterior in the class of posteriors. This approach eliminates the source of posterior sensitivity arising due to the unrevisable component of the prior. See Giacomini et al. (2021b) for a review of robust Bayesian approaches to inference in set-identified econometric models.

As noted in Plagborg-Møller and Wolf (2021a), an alternative way to use information about the signs of a particular structural shock in specific periods is to recast this information as a discrete-valued proxy for the structural shock. Specifically, construct a variable that is equal to one in periods where the shock is known to be positive, minus one in periods where it is known to be negative, and zero in all other periods. This variable is positively correlated with the target structural shock (i.e. ‘relevant’)

⁴An additional problem arises under more-general classes of NR, such as restrictions on the historical decomposition. In constructing the posterior, Antolín-Díaz and Rubio-Ramírez (2018) use the *conditional* likelihood, which is the likelihood function conditional on the NR holding. GKR show that the use of this likelihood can distort the posterior distribution towards parameter values that result in a lower *ex ante* probability that the NR are satisfied. To address this problem, GKR propose using the *unconditional* likelihood to construct the posterior. This problem does not arise under shock-sign restrictions, because the conditional and unconditional likelihoods are identical up to a multiplicative constant.

and uncorrelated with all other structural shocks (i.e. ‘exogenous’). The variable is therefore a valid proxy for the structural shock and can be used to point-identify impulse responses to that shock in a proxy SVAR (Mertens and Ravn 2013; Stock and Watson 2018).⁵

Montiel-Olea et al. (2021) – henceforth, MSW – show that frequentist inference about impulse responses is nonstandard when the proxy variable is only weakly correlated with the target structural shock (see also Lunsford (2015)). They propose a weak-proxy robust approach to inference, which – under some conditions – has asymptotically valid coverage of the true impulse response regardless of the strength of the correlation between the proxy and shock. We show that the narrative proxy is likely to be weakly correlated with the target structural shock when there are only a small number of periods in which the sign of the shock is known, which tends to be the case empirically; for example, the applications considered in Antolín-Díaz and Rubio-Ramírez (2018) impose NR in no more than a handful of periods. Accordingly, the weak-proxy robust approach to inference is the natural approach to frequentist inference when using NP.

We discuss the frequentist properties of the inferential procedures described above when the parameter of interest is the impulse response to a unit structural shock (e.g. a monetary policy shock that raises the federal funds rate by 100 basis points).⁶ When the number of NR is fixed with the sample size, we show that the robust credible region has asymptotic frequentist coverage of the true impulse response that is weakly greater than the nominal level. When the number of NR increases proportionally with the sample size, the coverage probability converges to the nominal level asymptotically.

We then show that the assumptions required for the asymptotic validity of the weak-proxy robust confidence intervals in MSW are violated under the NP approach. Specifically, the covariance between the NP and the reduced-form VAR innovations is not \sqrt{T} -asymptotically normal (where T is the sample size) when the NP is nonzero in a fixed number of periods. The MSW confidence intervals are constructed following the Anderson-Rubin approach, which involves inverting a particular Wald test. The

⁵The proxy is also uncorrelated with leads and lags of the structural shocks, so it could alternatively be used to point-identify the impulse responses to that shock in an instrumental-variables local projection, which does not require assuming that the structural shock is invertible (Stock and Watson 2018).

⁶While most studies using set-identified SVARs focus on impulse responses to a standard-deviation shock as the parameters of interest, the impulse responses to a unit shock are naturally of greater interest to policymakers (e.g. Stock and Watson 2018).

Wald statistic for this test is asymptotically $\chi^2(1)$ under the assumptions in MSW. We show that if the NP is nonzero in a fixed number of periods, the null distribution of this test statistic is not $\chi^2(1)$, no matter how large the sample size T is. Consequently, the size of the test with $\chi^2(1)$ critical values is generally not equal to the nominal level and the confidence intervals obtained from inverting the test are not guaranteed to attain correct coverage even when T is large. Thus, weak-proxy robust confidence intervals based on the standard Anderson-Rubin approach are not recommended when the NP is nonzero in only a small number of periods.

We provide Monte Carlo evidence in support of our theoretical results. The coverage probability of the weak-proxy robust confidence intervals differs from the nominal level unless the sign of the shock is known in a relatively large number of periods. Whether the intervals under- or over-cover depends on the confidence level. In contrast, the robust credible intervals always over-cover, but the coverage probability converges towards the nominal level as the number of NR increases.

We focus our discussion on a bivariate, static SVAR. The main reason for this simplification is that it allows us to derive an analytical expression for the ‘conditional identified set’ (i.e. the set of parameter values that are consistent with the NR given the reduced-form parameters), which is useful for both illustration and simulation of the NR approach. The theoretical findings do not rely on this simplification, and we provide some suggestions in the paper on how the analysis could be generalized to incorporate dynamics or additional variables.

To the best of our knowledge, this is the first paper to explore whether information about structural shocks should be used to impose NR in an otherwise set-identified SVAR or to construct a proxy for use in a proxy SVAR. Boer and Lütkepohl (2021) compare the efficiency of proxy SVAR estimators that use only the signs of structural shocks on particular dates against proxies that also use information about the size of the shock (i.e. standard proxies). Based on Monte Carlo simulations, they conclude that the estimator that uses information only about the sign of the shock is nearly as (or, in some cases, more) efficient than the estimator based on the quantitative information. Budnik and Rünstler (2020) consider identification of impulse responses in Bayesian proxy SVARs when the proxy represents the sign of a certain structural shock in particular periods. Their approach to identification departs from the standard proxy SVAR setting and is implemented using linear discriminant analysis or a non-parametric sign-concordance criterion.

The remainder of the paper is structured as follows. Section 2 describes the two approaches to using information about shock signs within the context of a stylised SVAR. Section 3 examines the theoretical frequentist properties of the robust Bayesian approach to inference under NR and the weak-proxy robust approach to inference under NP. Section 4 provides results from Monte Carlo simulations supporting the theoretical results. Section 5 concludes.

2 Frameworks For Using Information about Shock Signs

This section introduces a stylised bivariate SVAR and discusses the two approaches to using information about the sign of a structural shock.

2.1 Stylised SVAR

Consider the SVAR(0) for the 2×1 vector of endogenous variables $\mathbf{y}_t = (y_{1t}, y_{2t})'$:

$$\mathbf{A}_0 \mathbf{y}_t = \boldsymbol{\varepsilon}_t, \quad t = 1, \dots, T, \quad (1)$$

where $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \varepsilon_{2t})'$ with $\boldsymbol{\varepsilon}_t \stackrel{iid}{\sim} \mathcal{N}(\mathbf{0}_{2 \times 1}, \mathbf{I}_2)$, where \mathbf{I}_2 is the 2×2 identity matrix. The orthogonal reduced form of the model reparameterizes \mathbf{A}_0^{-1} as $\boldsymbol{\Sigma}_{tr} \mathbf{Q}$, where $\boldsymbol{\Sigma}_{tr}$ is the lower-triangular Cholesky factor (with positive diagonal elements) of $\boldsymbol{\Sigma} = \mathbb{E}(\mathbf{y}_t \mathbf{y}_t')$ = $\mathbf{A}_0^{-1} (\mathbf{A}_0^{-1})'$. We parameterize $\boldsymbol{\Sigma}_{tr}$ as

$$\boldsymbol{\Sigma}_{tr} = \begin{bmatrix} \sigma_{11} & 0 \\ \sigma_{21} & \sigma_{22} \end{bmatrix} \quad (\sigma_{11}, \sigma_{22} > 0), \quad (2)$$

and denote the vector of reduced-form parameters as $\boldsymbol{\phi} = \text{vech}(\boldsymbol{\Sigma}_{tr})$. \mathbf{Q} is an orthonormal matrix in the space of 2×2 orthonormal matrices, $\mathcal{O}(2)$:

$$\mathbf{Q} \in \mathcal{O}(2) = \left\{ \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} : \theta \in [-\pi, \pi] \right\} \cup \left\{ \begin{bmatrix} \cos \theta & \sin \theta \\ \sin \theta & -\cos \theta \end{bmatrix} : \theta \in [-\pi, \pi] \right\}, \quad (3)$$

where the first set is the set of ‘rotation’ matrices and the second set is the set of ‘reflection’ matrices.

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

$$\mathbf{A}_0^{-1} \in \left\{ \begin{aligned} & \left[\begin{array}{cc} \sigma_{11} \cos \theta & -\sigma_{11} \sin \theta \\ \sigma_{21} \cos \theta + \sigma_{22} \sin \theta & \sigma_{22} \cos \theta - \sigma_{21} \sin \theta \end{array} \right] : \theta \in [-\pi, \pi] \\ \cup & \left[\begin{array}{cc} \sigma_{11} \cos \theta & \sigma_{11} \sin \theta \\ \sigma_{21} \cos \theta + \sigma_{22} \sin \theta & \sigma_{21} \sin \theta - \sigma_{22} \cos \theta \end{array} \right] : \theta \in [-\pi, \pi] \end{aligned} \right\}. \quad (4)$$

The identified set for \mathbf{A}_0 is then

$$\mathbf{A}_0 \in \left\{ \begin{aligned} & \frac{1}{\sigma_{11}\sigma_{22}} \left[\begin{array}{cc} \sigma_{22} \cos \theta - \sigma_{21} \sin \theta & \sigma_{11} \sin \theta \\ -\sigma_{21} \cos \theta - \sigma_{22} \sin \theta & \sigma_{11} \cos \theta \end{array} \right] : \theta \in [-\pi, \pi] \\ \cup & \left[\begin{array}{cc} \sigma_{22} \cos \theta - \sigma_{21} \sin \theta & \sigma_{11} \sin \theta \\ \sigma_{22} \sin \theta + \sigma_{21} \cos \theta & -\sigma_{11} \cos \theta \end{array} \right] : \theta \in [-\pi, \pi] \end{aligned} \right\}. \quad (5)$$

Henceforth, we leave implicit that $\theta \in [-\pi, \pi]$ and we impose the sign normalisation $\text{diag}(\mathbf{A}_0) \geq \mathbf{0}_{2 \times 1}$, which is a normalisation on the signs of the structural shocks.

The impulse response of the second variable to a shock to the first variable that raises the first variable by one unit is

$$\eta_{21}(\phi, \theta) = \frac{\sigma_{21} \cos \theta + \sigma_{22} \sin \theta}{\sigma_{11} \cos \theta} = \frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}}{\sigma_{11}} \tan \theta. \quad (6)$$

In what follows, we assume that this is the parameter of interest.

2.2 Narrative Restrictions and Robust Bayesian Inference

Consider the NR that $\varepsilon_{1t} \geq 0$ for some $t \in \{1, \dots, T\}$, which is one of the restrictions proposed by Antolín-Díaz and Rubio-Ramírez (2018) and Ludvigson et al. (2018):

$$\varepsilon_{1t} = \mathbf{e}'_1 \mathbf{A}_0 \mathbf{y}_t = (\sigma_{11}\sigma_{22})^{-1} (\sigma_{22}y_{1t} \cos \theta + (\sigma_{11}y_{2t} - \sigma_{21}y_{1t}) \sin \theta) \geq 0, \quad (7)$$

where \mathbf{e}_1 is the first column of the 2×2 identity matrix, \mathbf{I}_2 . Under the sign normalization and the NR, θ is restricted to the set

$$\theta \in \{\theta : \sigma_{21} \sin \theta \leq \sigma_{22} \cos \theta, \cos \theta \geq 0, \sigma_{22}y_{1t} \cos \theta \geq (\sigma_{21}y_{1t} - \sigma_{11}y_{2t}) \sin \theta\}$$

$$\cup \{\theta : \sigma_{21} \sin \theta \leq \sigma_{22} \cos \theta, \cos \theta \leq 0, \sigma_{22} y_{1t} \cos \theta \geq (\sigma_{21} y_{1t} - \sigma_{11} y_{2t}) \sin \theta\}. \quad (8)$$

Since y_{1t} and y_{2t} enter the inequalities characterising this set, the NR induces a set-valued mapping from ϕ to θ – and thus to η_{21} – that depends on the realization of \mathbf{y}_t ; GKR refer to this mapping as the ‘conditional identified set’.⁷ For example, in the case where $\sigma_{21} < 0$, $h(\phi, \mathbf{y}_t) = \sigma_{21} y_{1t} - \sigma_{11} y_{2t} > 0$ and $C(\phi, \mathbf{y}_t) = \sigma_{22} y_{1t} / h(\phi, \mathbf{y}_t) > \sigma_{22} / \sigma_{21}$, the conditional identified set for θ is

$$CIS_{\theta}(\phi | \mathbf{y}_t) = \left[\arctan\left(\frac{\sigma_{22}}{\sigma_{21}}\right), \arctan(C(\phi, \mathbf{y}_t)) \right], \quad (9)$$

and the conditional identified set for η_{21} is

$$CIS_{\eta}(\phi | \mathbf{y}_t) = \left[\frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}^2}{\sigma_{21} \sigma_{11}}, \frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}}{\sigma_{11}} C(\phi, \mathbf{y}_t) \right]. \quad (10)$$

The conditional identified set for η_{21} may be unbounded; for example, when $\sigma_{21} < 0$ and $h(\phi, \mathbf{y}) < 0$, the conditional identified set for θ includes $\pi/2$ and the conditional identified set for η_{21} is $(-\infty, \infty)$, which implies that the NR is completely uninformative about the impulse response.⁸

As in GKR, we do not model a mechanism for when the shock-sign restrictions become available. Our analysis assumes that the sign of ε_{1t} is revealed in the first K periods, but our results would also hold if we assumed that the sign of ε_{1t} was revealed in K random periods. The set of restrictions can be represented as $\{N(\theta, \phi, \mathbf{y}_t) \geq 0\}_{t=1}^K$. The ‘unconditional’ likelihood based on the sample $\mathbf{y}^T = (\mathbf{y}'_1, \dots, \mathbf{y}'_T)'$ and the NR is

$$p(\mathbf{y}^T, \{N(\theta, \phi, \mathbf{y}_t) \geq 0\}_{t=1}^K | \theta, \phi) = \prod_{t=1}^T (2\pi)^{-1} |\Sigma|^{-\frac{1}{2}} \exp\left(-\frac{1}{2} \mathbf{y}'_t \Sigma^{-1} \mathbf{y}_t\right) \prod_{t=1}^K 1(N(\theta, \phi, \mathbf{y}_t) \geq 0), \quad (11)$$

⁷The conditional identified set is the set of values of a particular parameter that are consistent with the reduced-form parameters and the imposed restrictions (i.e. the NR and any other identifying restrictions, such as sign restrictions on impulse responses). The concept of a conditional identified set differs from that of a standard identified set. The latter is defined by a (set-valued) mapping from reduced-form to structural parameters, while the former additionally depends on the realisation of the data in particular periods. See GKR for further details.

⁸Appendix A characterises the conditional identified set for θ when $\sigma_{21} < 0$ under the restriction that $\varepsilon_{1t} \geq 0$ or $\varepsilon_{1t} \leq 0$. It also explains how to obtain the conditional identified set for η_{21} . We draw on these analytical expressions when conducting Monte Carlo exercises in Section 4.

where $1(N(\theta, \phi, \mathbf{y}_t) \geq 0)$ is the indicator function equal to one when $N(\theta, \phi, \mathbf{y}_t) \geq 0$ and equal to zero otherwise.⁹ The dependence of the likelihood on θ is only through the indicator functions. The NR truncate the likelihood so that it is flat for any θ satisfying the NR and is zero otherwise, with the points of truncation depending on the realisation of $\{\mathbf{y}_t\}_{t=1}^K$.

The standard approach to Bayesian inference in this setting (e.g. Antolín-Díaz and Rubio-Ramírez 2018) would involve specifying a uniform prior over θ , which corresponds to the usual uniform or ‘Haar’ prior over the orthonormal matrix \mathbf{Q} (e.g. Baumeister and Hamilton 2015; Arias et al. 2018). However, GKR argue that the flat likelihood function raises concerns about the sensitivity of posterior inference to the choice of prior, which is primarily chosen for computational convenience and does not reflect actual subjective beliefs about the parameters; the posterior for θ will be proportional to the prior whenever the likelihood is nonzero. This is similar to the case under standard sign restrictions (e.g. Baumeister and Hamilton 2015). GKR therefore propose extending the robust Bayesian approach of Giacomini and Kitagawa (2021) to this setting.

The key feature of the robust Bayesian approach to inference under NR is that, rather than imposing a single conditional prior over θ , we consider the class of all priors that are consistent with any ‘traditional’ set-identifying restrictions (in this case, the sign normalisations). Combining the class of priors with the unconditional likelihood generates a class of posteriors for θ – and thus for η_{21} – that are consistent with the NR. Importantly, this approach eliminates the source of posterior sensitivity to the unrevisable component of the prior (i.e. the conditional prior for θ given ϕ).

In the current context, the robust Bayesian approach to conducting inference about η_{21} proceeds as follows. First, impose a prior over ϕ and obtain draws of ϕ from its posterior, $\pi_{\phi|\mathbf{Y}^T}$. Second, for each draw of ϕ , compute the conditional identified set for η_{21} given the NR:¹⁰

$$CIS_{\eta}(\phi | \{\mathbf{y}_t\}_{t=1}^K) = \left\{ \eta_{21}(\phi, \theta) : \{N(\theta, \phi, \mathbf{y}_t) \geq 0\}_{t=1}^K \right\} \quad (12)$$

$$= \left[\ell(\phi, \{\mathbf{y}_t\}_{t=1}^K), u(\phi, \{\mathbf{y}_t\}_{t=1}^K) \right]. \quad (13)$$

⁹In this case the ‘conditional’ likelihood used by Antolín-Díaz and Rubio-Ramírez (2018) is equal to the unconditional likelihood up to a multiplicative constant that does not depend on θ ; the conditional likelihood is obtained from the unconditional likelihood by dividing by the probability that the structural shocks satisfy the shock-sign restrictions, which is $(1/2)^K$.

¹⁰In defining the conditional identified sets, we leave implicit the sign normalisation on \mathbf{A}_0 .

Finally, using the posterior draws of the bounds of the conditional identified set, construct a robust credible region with credibility α as an interval C_α satisfying

$$\pi_{\phi|\mathbf{Y}^T} \left(\left[\ell \left(\phi, \{\mathbf{y}_t\}_{t=1}^K \right), u \left(\phi, \{\mathbf{y}_t\}_{t=1}^K \right) \right] \subset C_\alpha \right) = \alpha. \quad (14)$$

This interval receives posterior probability at least α under all posterior distributions within the class of posteriors for η_{21} . The robust credible interval with credibility α will be unbounded whenever $\pi_{\phi|\mathbf{Y}^T} \left(CIS_\eta \left(\phi | \{\mathbf{y}_t\}_{t=1}^K \right) = \mathbb{R} \right) \geq 1 - \alpha$.

2.3 Narrative Proxies and Weak-proxy Robust Inference

Plagborg-Møller and Wolf (2021a) suggest that the information about the signs of structural shocks imposed under NR can alternatively be used to construct a proxy variable, which can be used to point-identify impulse responses in a proxy SVAR or instrumental-variable local projection.

Solving the equation for ε_{2t} in (1) for y_{2t} and plugging in (6), we obtain, for $t = 1, \dots, T$,

$$y_{2t} = \eta_{21}y_{1t} + w_t, \quad (15)$$

where w_t depends on ε_{2t} , σ_{11} and θ . To identify the impulse response of interest η_{21} , consider using some variable z_t satisfying $\mathbb{E}(z_t\varepsilon_{1t}) \neq 0$ (relevance) and $\mathbb{E}(z_t\varepsilon_{2t}) = 0$ (exogeneity) as an instrument for y_{1t} .¹¹ Since $y_{1t} = \mathbf{e}'_1 \mathbf{A}_0^{-1} \boldsymbol{\varepsilon}_t$, relevance of z_t implies $\mathbb{E}(z_t y_{1t}) \neq 0$. Hence, by the standard identification approach using the linear instrumental-variable method, we can identify η_{21} by the Wald estimand (Stock and Watson (2018)):

$$\eta_{21} = \frac{\mathbb{E}(z_t y_{2t})}{\mathbb{E}(z_t y_{1t})}. \quad (16)$$

As above, we assume that the sign of the shock is observable in the first K periods. The proposal from Plagborg-Møller and Wolf (2021a) is to use the information about the shock signs to construct z_t as:

$$z_t = \begin{cases} \text{sign}(\varepsilon_{1t}) & \text{if } t = 1, \dots, K \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

¹¹In a model with dynamics or additional variables, impulse responses would be identified from the instrumental-variables regression of one reduced-form VAR innovation on another; see Stock and Watson (2016, 2018) or MSW.

$\text{sign}(\varepsilon_{1t})$ is clearly (positively) correlated with ε_{1t} itself and is uncorrelated with ε_{2t} , because the two shocks are independent by assumption. This proxy will therefore satisfy the relevance and exogeneity requirements, and, therefore, at any $t \in \{1, \dots, K\}$, the Wald estimand (16) identifies η_{21} . Since the expectations in (16) are common for all $t = 1, \dots, K$, a sample analogue estimator for η_{21} can be formed as

$$\hat{\eta}_{21} = \frac{\frac{1}{K} \sum_{t=1}^K z_t y_{2t}}{\frac{1}{K} \sum_{t=1}^K z_t y_{1t}} = \frac{\frac{1}{T} \sum_{t=1}^T z_t y_{2t}}{\frac{1}{T} \sum_{t=1}^T z_t y_{1t}} = \frac{\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) y_{2t}}{\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) y_{1t}}, \quad (18)$$

which is identical to the two-stage least square estimator using z_t as an instrumental variable for y_{1t} .

In the case where $K = 1$, $\hat{\eta}_{21} = y_{2,1}/y_{1,1}$. The point estimator of the impulse response will therefore be sensitive to the realisation of the structural shocks in a single period (i.e. the period in which the sign of the shock is known). For example, if $\varepsilon_{2t} \neq 0$ and $\varepsilon_{1t} \approx 0$, $\hat{\theta} \approx \arctan(-\cot \theta_0)$ and $\hat{\eta}_{21}$ is approximately equal to

$$\frac{\sigma_{21}}{\sigma_{11}} + \frac{\sigma_{22}}{\sigma_{11}} \tan(\arctan(-\cot(\theta_0))) = \frac{\sigma_{21}}{\sigma_{11}} - \frac{\sigma_{22}}{\sigma_{11}} \cot \theta_0, \quad (19)$$

which, rather than the true value of η_{21} , equals the true impulse response of the second variable to the *second* shock such that the first variable increases by one unit.¹² More generally, when $K > 1$, the estimator of the impulse response under NP may be sensitive to the realisations of the non-target shocks in the periods in which the sign of the target shock is known.

As discussed in MSW, standard approaches to inference (e.g. based on asymptotic normality of the reduced-form parameters and the delta-method) in proxy SVARs are invalid when the proxy is only weakly correlated with the target structural shock. This is likely to be the case in applications using NP. For example, under the assumption that the structural shocks are normally distributed and the sign of the shock is known in K periods, the expected value of the sample covariance between the proxy variable and the target structural shock is $(K/T)\sqrt{2/\pi}$, where T is the number of observations and expectations are taken over alternative realisations of a time series of length T . For K small relative to T , this expected covariance will be close to zero. Furthermore, for K fixed as T approaches infinity, the expected covariance will converge to zero at

¹²The estimator under the NP when $K = 1$ coincides with the estimator that would be obtained under the ‘narrative zero restriction’ $\varepsilon_{2,1} = 0$.

rate T , which is faster than the \sqrt{T} rate of convergence typically considered under weak-instrument asymptotics. The scenario where K is small (or fixed) relative to T appears to be the empirically relevant case given that most studies that impose NR do so in at most a handful of periods.¹³

Given that the NP is likely to be weakly correlated with the target structural shock in typical empirical applications, it seems natural to turn to the weak-proxy robust approach to inference developed in MSW. Building on the Anderson-Rubin approach to weak-instrument robust inference in microeconometrics, they develop weak-proxy robust confidence intervals for the impulse responses in a proxy SVAR. MSW show that their weak-proxy robust confidence intervals have correct coverage of the true impulse response under the assumption that the reduced-form parameters are asymptotically normally distributed. Importantly, this result does not depend on the strength of the covariance between the NP and the target structural shock.

The Anderson-Rubin confidence intervals are obtained as follows. Define the vector of average covariances between the NP and the observations by

$$\mathbf{\Gamma}_T \equiv \begin{bmatrix} \gamma_{1T} \\ \gamma_{2T} \end{bmatrix} = \frac{1}{T} \sum_{t=1}^T \mathbb{E}(z_t \mathbf{y}_t). \quad (20)$$

As noted above, we assume that the sign of ε_{1t} is known in the first K periods. This implies that the proxy $\{z_t : t = 1 \dots T\}$ is not weakly stationary and the covariance $\mathbb{E}(z_t \mathbf{y}_t)$ varies with t :

$$\mathbb{E}(z_t \mathbf{y}_t) = \begin{cases} \mathbb{E}(\text{sign}(\varepsilon_{1t}) \mathbf{y}_t) & \text{for } 1 \leq t \leq K, \\ 0, & \text{for } t > K, \end{cases}$$

so $\mathbf{\Gamma}_T = \frac{1}{T} \sum_{t=1}^K \mathbb{E}(\text{sign}(\varepsilon_{1t}) \mathbf{y}_t)$.

Let $\hat{\mathbf{\Gamma}}_T$ denote an analogue estimator of $\mathbf{\Gamma}_T$. Because we know the signs of the first structural shock in the first K periods only,

$$\hat{\mathbf{\Gamma}}_T = \begin{bmatrix} \hat{\gamma}_{1T} \\ \hat{\gamma}_{2T} \end{bmatrix} = \frac{1}{T} \sum_{t=1}^T z_t \mathbf{y}_t = \frac{1}{T} \sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \mathbf{y}_t. \quad (21)$$

¹³For example, when estimating the effects of US monetary policy, Antolín-Díaz and Rubio-Ramírez (2018) impose NR in a single period in their main specification and in eight periods in a robustness exercise.

The Anderson-Rubin confidence intervals in MSW are constructed by inverting tests based on the Wald statistic

$$W_T(\eta_{21}) = \frac{T(\hat{\gamma}_{2T} - \eta_{21}\hat{\gamma}_{1T})^2}{\begin{pmatrix} -\eta_{21} & 1 \end{pmatrix} \hat{\Omega}_T \begin{pmatrix} -\eta_{21} & 1 \end{pmatrix}'}, \quad (22)$$

where $\hat{\Omega}_T$ is the sample variance-covariance matrix of $(z_t \mathbf{y}_t : t = 1, \dots, T)$. Under the assumptions in MSW, $W_T(\eta_{21})$ asymptotically follows the $\chi^2(1)$ distribution under the null. Constructing an α -level confidence interval for η_{21} requires finding the set of values of η_{21} such that $W_T(\eta_{21}) \leq \chi_\alpha^2(1)$, where $\chi_\alpha^2(1)$ is the α quantile of the $\chi^2(1)$ distribution. This inequality is quadratic in η_{21} and has an analytical solution.

3 Frequentist Properties of Approaches to Inference

This section discusses the asymptotic frequentist properties of the robust Bayesian approach to inference under NR and the weak-proxy robust approach to inference under NP.

3.1 Robust Bayesian Inference Under NR

GKR show that their robust Bayesian approach to inference under NR is asymptotically valid from a frequentist perspective when the number of NR is fixed with the sample size. In particular, the robust credible interval has at least the nominal level of coverage for the true impulse response to a standard-deviation shock asymptotically under some high-level assumptions. These assumptions are that the conditional identified set for the impulse response is closed and convex, and has lower and upper bounds that are differentiable in ϕ at $\phi = \phi_0$ with nonzero derivatives.¹⁴

In the current context, we are interested in the impulse response to a unit shock rather than a standard-deviation shock. As discussed in Read (forthcoming), this shift in the parameter of interest may give rise to an unbounded conditional identified set, depending on ϕ and the realization of the data. This means that the asymptotic

¹⁴Propositions B.1–B.3 in GKR provide primitive conditions under which these assumptions hold when there are shock-sign restrictions.

validity claim from GKR does not immediately apply when the parameter of interest is the impulse response to a unit shock. However, the following informal proof shows that asymptotic validity of the robust credible interval also holds in the NR setting considered in this paper.

Conditional on the realization of the data entering the NR, $\{\mathbf{y}_t\}_{t=1}^K$, if the conditional identified set is bounded at any ϕ in an open neighbourhood of ϕ_0 and has lower and upper bounds differentiable in ϕ (with nonzero derivatives), the robust credible interval will asymptotically contain the true conditional identified set, $CIS_\eta(\phi_0|\{\mathbf{y}_t\}_{t=1}^K)$, with probability α under the sampling distribution of the remaining data (i.e. $\{\mathbf{y}_t\}_{t=K+1}^T$). Assuming the NR are correctly specified, the true impulse response will lie within the conditional identified set, so the robust credible interval will also asymptotically contain the true impulse response with probability at least α . If the conditional identified set is unbounded at $\phi = \phi_0$ and its neighborhood, the robust credible interval is also unbounded with a high probability in terms of the sampling distribution of the remaining data, i.e., the robust credible interval contains the true impulse response with probability approaching one under the sampling distribution of the remaining data. Thus, regardless of whether the conditional identified set is bounded around $\phi = \phi_0$, the unconditional coverage probability – obtained by averaging over the possible realisations of $\{\mathbf{y}_t\}_{t=1}^K$ – is greater than the nominal coverage probability.

GKR also show that when the number of NR increases with the sample size (i.e. $K \rightarrow \infty$ in proportion to $T \rightarrow \infty$), the posterior distributions of the lower and upper bounds of the conditional identified set for θ are \sqrt{T} -consistent for θ_0 . The posterior consistency property will also carry over to the impulse response η_{21} , in the sense that the posterior distributions of the lower and upper bounds of the conditional identified set for η_{21} converge to the true η_{21} so long as $\theta_0 \neq \pm\pi/2$, which are the points of singularity for $\eta_{21}(\phi, \theta)$. Furthermore, the posteriors of the lower and upper bounds of the conditional identified set for η_{21} are asymptotically normal, and the posterior credible intervals constructed upon them as well as the robust Bayes credible interval for η_{21} asymptotically attain the correct frequentist coverage.

3.2 Weak-proxy Robust Inference Under NP

Assumption 2 of MSW requires that the reduced-form covariance estimator $\hat{\mathbf{\Gamma}}_T$ is normally distributed asymptotically. However, in the NP setting with fixed K ,

$$\begin{aligned} \sqrt{T}(\hat{\mathbf{\Gamma}}_T - \mathbf{\Gamma}_T) &= \sqrt{T} \left(\frac{1}{T} \sum_{t=1}^T [z_t \mathbf{y}_t - \mathbb{E}(z_t \mathbf{y}_t)] \right) \\ &= \frac{1}{\sqrt{T}} \mathbf{A}_0^{-1} \sum_{t=1}^K [\text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t - \mathbb{E}(\text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t)] \\ &= \frac{1}{\sqrt{T}} \mathbf{A}_0^{-1} \sum_{t=1}^K \left[\text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t - \frac{K}{T} \sqrt{\frac{2}{\pi}} \mathbf{e}_1 \right] \\ &\rightarrow_p \mathbf{0}_{2 \times 1} \quad \text{as } T \rightarrow \infty. \end{aligned}$$

The reduced-form covariance between the NP and the data therefore has a degenerate distribution asymptotically after rescaling by \sqrt{T} . Consequently, Assumption 2 in MSW is not satisfied, and it is unclear whether the weak-proxy robust confidence intervals will have valid coverage of the true impulse response in the NP setting.

To examine the coverage properties of the Anderson-Rubin confidence intervals in this setting, express the parameter of interest η_{21} in terms of $\mathbf{\Gamma}_T$. η_{21} is identified by the ratio of the second and first elements of $\mathbf{\Gamma}_T$:

$$\eta_{21} = \frac{T^{-1} \sum_{t=1}^T \mathbb{E}(z_t y_{2t})}{T^{-1} \sum_{t=1}^T \mathbb{E}(z_t y_{1t})} = \frac{T^{-1} \sum_{t=1}^K \mathbb{E}(\text{sign}(\varepsilon_{1t}) y_{2t})}{T^{-1} \sum_{t=1}^K \mathbb{E}(\text{sign}(\varepsilon_{1t}) y_{1t})}. \quad (23)$$

Following MSW, we view the Anderson-Rubin confidence intervals as providing inference for the ratio of the two means based on their unbiased statistics $\hat{\mathbf{\Gamma}}_T = \frac{1}{T} \sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \mathbf{y}_t$. As discussed above, with fixed K , $\sqrt{T}(\hat{\mathbf{\Gamma}}_T - \mathbf{\Gamma}_T)$ is not asymptotically normal. In addition, even under the assumption that the structural shocks are Gaussian, the distribution of $\hat{\mathbf{\Gamma}}_T$ is not jointly normal due to the multiplier term $\text{sign}(\varepsilon_{1t})$.

As described in Section 2.3, the standard Anderson-Rubin confidence intervals are constructed by inverting a Wald test. Let η_{21} be the true impulse response specified under the null. If we naively apply the Anderson-Rubin confidence intervals, the Wald statistic for the test to be inverted is given by Equation (22). $\hat{\mathbf{\Omega}}_T$, which appears in

the denominator of the Wald statistic, can be expressed as

$$\begin{aligned}\hat{\Omega}_T &= \frac{1}{T} \sum_{t=1}^T (z_t \mathbf{y}_t - \hat{\Gamma}_T)(z_t \mathbf{y}_t - \hat{\Gamma}_T)' \\ &= \mathbf{A}_0^{-1} \left[\frac{1}{T} \sum_{t=1}^K \boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t' - \left(\frac{1}{T} \sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right) \left(\frac{1}{T} \sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right)' \right] (\mathbf{A}_0^{-1})'. \quad (24)\end{aligned}$$

Multiplying the numerator of the Wald statistic by T yields

$$\begin{aligned}T^2 (\hat{\gamma}_{2T} - \eta_{21} \hat{\gamma}_{1T})^2 &= \\ &= \begin{pmatrix} -\eta_{21} & 1 \end{pmatrix} \mathbf{A}_0^{-1} \left[\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right] \left[\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right]' (\mathbf{A}_0^{-1})' \begin{pmatrix} -\eta_{21} & 1 \end{pmatrix}'. \quad (25)\end{aligned}$$

The Wald statistic (22) can then be expressed as

$$\begin{aligned}W_T(\eta_{21}) &= \\ &= \frac{\begin{pmatrix} -\eta_{21} & 1 \end{pmatrix} \mathbf{A}_0^{-1} \left[\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right] \left[\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right]' (\mathbf{A}_0^{-1})' \begin{pmatrix} -\eta_{21} & 1 \end{pmatrix}'}{\begin{pmatrix} -\eta_{21} & 1 \end{pmatrix} \mathbf{A}_0^{-1} \left[\sum_{t=1}^K \boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t' - \left(\frac{1}{\sqrt{T}} \sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right) \left(\frac{1}{\sqrt{T}} \sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right)' \right] (\mathbf{A}_0^{-1})' \begin{pmatrix} -\eta_{21} & 1 \end{pmatrix}'}. \quad (26)\end{aligned}$$

Under the assumptions in MSW, the Wald statistic has the $\chi^2(1)$ distribution asymptotically. However, under NP, the expression above for the Wald statistic shows that $W_T(\eta_{21})$ does not follow the $\chi^2(1)$ distribution under the null even asymptotically as $T \rightarrow \infty$ with fixed K . First, the numerator is not a squared Gaussian random variable when $K \geq 2$. Second, the denominator remains random even asymptotically and does not converge in probability to the variance-covariance matrix of the numerator random variable under the null. In the special case where $K = 1$, the fact that $(\text{sign}(\varepsilon_{1t}))^2 = 1$ means that the numerator is a squared Gaussian random variable. However, in this case the denominator is $\frac{T-1}{T}$ times the numerator, so the Wald statistic is constant and equal to $\frac{T}{T-1}$ under the null. Performing the Wald test using the statistic $W_T(\eta_{21})$ and critical values from the $\chi^2(1)$ distribution therefore does not generally control the type I error of the test, because the actual distribution of $W_T(\eta_{21})$ is not $\chi^2(1)$. Consequently, a naive application of the Anderson-Rubin confidence intervals is not guaranteed to have the nominal coverage level of the true

impulse response.


To explore the severity of the distortion of the null distribution of the Wald statistic under NP, Figure 1 plots the null distribution of the Wald statistics under the distributional specification of $\boldsymbol{\varepsilon}_t \stackrel{iid}{\sim} \mathcal{N}(\mathbf{0}_{2 \times 1}, \mathbf{I}_2)$ for different values of K , alongside the $\chi^2(1)$ distribution. When $K = 1$, the Wald statistic is constant and equal to $T/(T - 1)$, so the cumulative distribution function of the Wald statistic is a step function. Since $T/(T - 1) \approx 1$ lies below the 95th percentile of the $\chi^2(1)$ distribution (3.84), the 95 per cent Anderson-Rubin confidence intervals span the entire real line and have a coverage probability equal to one. In contrast, since the 68th percentile of the $\chi^2(1)$ distribution (0.99) is less than $T/(T - 1)$, the 68 per cent confidence interval excludes any value of η_{21} except for $\eta_{21} = \hat{\gamma}_{2T}/\hat{\gamma}_{1T}$ and the intervals have coverage probability equal to zero. As K increases, the null distribution of the Wald statistic converges towards the $\chi^2(1)$ distribution. However, the 95th percentile of the null distribution continues to lie below the 95th percentile of the $\chi^2(1)$ distribution, so the 95 per cent Anderson-Rubin confidence interval will have coverage probability greater than the nominal level. The 68th percentile of the null distribution lies above the 68th percentile of the $\chi^2(1)$ distribution, so the 68 per cent confidence interval will have coverage lower than the nominal level. For moderately large values of K (e.g. 20), the null distribution of the Wald statistic is well-approximated by the $\chi^2(1)$ distribution and the coverage probabilities of the confidence intervals should be close to the nominal level.

Note that, under the assumption that the distribution of $\boldsymbol{\varepsilon}_t$ is known (i.e., that it is Gaussian), one approach to performing valid inference about the impulse response under NP would be to use $T^2 (\hat{\gamma}_{2T} - \eta_{21} \hat{\gamma}_{1T})^2$ as a test statistic and obtain the null distribution of this statistic by simulating a feasible version of the right-hand side of (25). Let $\hat{\boldsymbol{\Sigma}}_{tr}$ be a consistent estimator for $\boldsymbol{\Sigma}_{tr}$ (e.g. the maximum likelihood estimator). Given η_{21} at the null and $\hat{\boldsymbol{\Sigma}}_{tr}$, let $\hat{\mathbf{Q}}$ be the orthonormal matrix pinned down by the value of θ solving (6) with $\hat{\boldsymbol{\Sigma}}_{tr}$ plugged in and under the sign normalizations. We accordingly define $\hat{\mathbf{A}}_0^{-1} \equiv \hat{\boldsymbol{\Sigma}}_{tr} \hat{\mathbf{Q}}$. A feasible version of the right-hand side of (25) is then

$$\begin{pmatrix} -\eta_{21} & 1 \end{pmatrix} \hat{\mathbf{A}}_0^{-1} \left[\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right] \left[\sum_{t=1}^K \text{sign}(\varepsilon_{1t}) \boldsymbol{\varepsilon}_t \right]' (\hat{\mathbf{A}}_0^{-1})' \begin{pmatrix} -\eta_{21} & 1 \end{pmatrix}', \quad (27)$$

where $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_2)$, $t = 1, \dots, K$. In addition to relying on a strong distributional

Figure 1: Cumulative (Null) Distribution Function of Wald Statistic Under Narrative Proxy



Note: Dashed lines represent 68th and 95th percentiles; data-generating process assumes $T = 1,000$, $\eta_{21} = 0.4$ and $\text{vech}(\Sigma_{tr}) \approx (0.7, -1.1, 1.4)'$; null distribution approximated using 100,000 Monte Carlo replications.

assumption, a practical drawback of this approach is that it is computationally burdensome; the null distribution of the statistic is not pivotal, which means that the critical value of the test needs to be simulated at every null value of η_{21} .

Finally, note that arguments similar to those derived in this section for the static bivariate case also apply to the general case with additional variables and dynamics.¹⁵

¹⁵For instance, consider conducting inference about η_{21} in an n -variable SVAR with dynamics. Analogously to Equation (20), Γ_T is replaced by the $n \times 1$ vector of average covariances between the NP and the $n \times 1$ vector of reduced-form VAR innovations, \mathbf{u}_t , and $\hat{\Gamma}_T$ is the corresponding sample analogue. The Wald statistic is still given by $W_T(\eta_{21}) = T(\hat{\gamma}_{2T} - \eta_{21}\hat{\gamma}_{2T})^2/\hat{\omega}_T$, where $\hat{\omega}_T$ is the estimator of the asymptotic variance of $\sqrt{T}(\hat{\gamma}_{2T} - \eta_{21}\hat{\gamma}_{2T})$. \mathbf{y}_t is replaced with \mathbf{u}_t in Equation (24) and the vector $(-\eta_{21}, 1)$ in Equations (25) and (26) is replaced with $(-\eta_{21}, 1, \mathbf{0}_{1 \times (n-2)})$.

4 Monte Carlo

This section describes a Monte Carlo simulation designed to compare the properties of the weak-proxy robust confidence intervals in the NP setting against that of the robust Bayesian credible intervals under a corresponding set of shock-sign NR.

4.1 Design

For consistency with the theoretical results obtained above and to leverage analytical results about the conditional identified set under NR, we maintain the bivariate SVAR(0) described in Section 2. We set

$$\mathbf{A}_0^{-1} = \begin{bmatrix} 0.5 & -0.5 \\ 0.2 & 1.8 \end{bmatrix} \Rightarrow \boldsymbol{\Sigma} = \begin{bmatrix} 0.5 & -0.8 \\ -0.8 & 3.28 \end{bmatrix}. \quad (28)$$

The parameter of interest is the impulse response of y_{2t} to a unit shock in y_{1t} , which is $0.2/0.5 = 0.4$.

We proceed under the assumption that $\boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}_{2 \times 1}, \mathbf{I}_2)$. We assume that the sign of ε_{1t} is revealed in the first K periods (although the results would be identical if we assumed that the sign of the shock was revealed in K random periods). We fix the sample length $T = 1,000$ and compare the two approaches to identification and inference under different assumptions about the number of periods in which the sign of the shock is known.

Under the NP approach, the proxy z_t is generated according to $z_t = \text{sign}(\varepsilon_{1t})$ for $t = 1, \dots, K$ and $z_t = 0$ otherwise. For each of 10,000 Monte Carlo samples, we compute weak-proxy robust confidence intervals by applying the replication code from MSW.¹⁶

Under the NR approach, we impose the restriction $\text{sign}(\varepsilon_{1t})\varepsilon_{1t} \geq 0$ for $t = 1, \dots, K$ along with the sign normalisation $\mathbf{e}'_1 \mathbf{A}_0 \mathbf{e}_1 \geq 0$. We conduct inference using the robust Bayesian approach from GKR. We assume an improper Jeffreys' prior over $\boldsymbol{\Sigma}$, which implies that the posterior is inverse-Wishart. In each of 1,000 Monte Carlo samples, we obtain 1,000 draws of $\boldsymbol{\Sigma}_{tr}$ and compute the lower and upper bounds of the conditional identified set for θ by first computing the lower and upper bounds

¹⁶The replication code was obtained from José Luis Montiel Olea's website: <http://www.joseluismontielolea.com>. We modify the code to omit a constant term from the VAR, since no constant appears in the data-generating process.

that would be obtained under each separate shock-sign restriction (for which we have analytical expressions; see Appendix for details). We then intersect the bounds obtained under each separate shock-sign restriction to obtain the bounds for the conditional identified set under the joint set of restrictions. The conditional identified set for θ is then projected into impulse-response space to obtain the conditional identified set for η_{21} . We construct a robust credible interval with credibility α by computing the $(1 - \alpha)/2$ quantile of the posterior distribution of the lower bound and the $(1 + \alpha)/2$ quantile of the posterior distribution of the upper bound.

4.2 Results

To summarise the performance of the weak-proxy robust approach to inference in the NP setting, we compute five statistics of interest. The first is the coverage probability of the weak-proxy robust confidence interval, which is the share of Monte Carlo replications in which the confidence interval includes the true value of the impulse response. The second is the proportion of Monte Carlo samples in which the confidence interval is unbounded. We also present the median width of the confidence interval across the Monte Carlo replications. Finally, we present the average value of the Wald statistic for testing the null hypothesis that the covariance between the proxy and y_{1t} is zero (i.e. $\hat{\gamma}_{1T} = 0$), which is a measure of the strength of the proxy. The α -level weak-proxy robust confidence intervals are bounded if and only if this statistic exceeds the α quantile of the $\chi^2(1)$ distribution.

The 95 per cent weak-proxy robust confidence interval is always unbounded when $K = 1, 2, 3$, so the coverage probability is trivially equal to one (Table 1). As K increases, the confidence intervals are bounded with higher probability. Coverage probabilities remain higher than the nominal level for small-to-moderate values of K . Only for larger values of K do the confidence intervals possess approximately correct coverage. The improving coverage properties of the confidence intervals as K increases reflects the fact that null distribution of W_T is better approximated by the $\chi^2(1)$ distribution (as discussed in the previous section). As K increases, the average value of the Wald statistic for testing the null hypothesis that $\hat{\gamma}_{1T} = 0$ increases, which indicates that the NP becomes a stronger (or more relevant) proxy as it becomes less sparse.

Table 2 repeats this exercise when $\alpha = 0.68$. Here, the key results are essentially

Table 1: Weak-proxy Robust Inference – Monte Carlo Results for $\alpha = 0.95$

K	Coverage prob.	Unbounded	Median width	Average W
1	1.000	1.000	∞	1.001
2	1.000	1.000	∞	1.194
3	1.000	1.000	∞	1.452
4	0.996	0.985	∞	1.761
5	0.978	0.887	∞	2.043
10	0.958	0.549	∞	3.623
20	0.952	0.180	8.222	6.800
30	0.951	0.047	5.067	10.009
40	0.954	0.011	3.975	13.256
50	0.951	0.001	3.371	16.550
100	0.947	0.000	2.133	33.362
500	0.950	0.000	0.894	189.693
1000	0.950	0.000	0.625	468.572

Notes: ‘Coverage prob.’ is the coverage probability of the 95 per cent weak-proxy robust confidence interval; ‘Unbounded’ is the proportion of Monte Carlo samples in which the confidence interval is unbounded; W is the Wald statistic for testing the null hypothesis that the covariance between the proxy and y_{1t} is zero.

the reverse of the case where $\alpha = 0.95$. Consistent with the analysis in the previous section, the 68 per cent confidence intervals have a coverage probability equal to zero when $K = 1$; in this case, the confidence interval always has zero width – it is a singleton equal to the point estimate. The confidence intervals possess coverage probabilities lower than the nominal level for small values of K . Again, the coverage probabilities approach the nominal level only for larger values of K .

Tables 3 and 4 present analogous statistics for the robust credible intervals obtained under the shock-sign NR. When the restrictions are imposed in only a handful of periods, the robust credible intervals are unbounded in a large proportion of the Monte Carlo replications (this occurs in any particular sample when the conditional identified set is unbounded with high posterior probability). The robust credible interval contains the true impulse response with probability greater than the nominal level at all values of K , which is consistent with the arguments in Section 3.1. Even when $K = 1,000$ – so that the sign of the shock is known in every period – the coverage probabilities are a bit higher than the nominal level (e.g. 95.5 per cent for a nominal level of 95 per cent and 71.8 per cent for a nominal level of 68 per cent).

These results suggest a trade-off when deciding whether to use the NP or NR

Table 2: Weak-proxy Robust Inference – Monte Carlo Results for $\alpha = 0.68$

K	Coverage prob.	Unbounded	Median width	Average W
1	0.000	0.000	0.000	1.001
2	0.493	0.376	7.824	1.194
3	0.581	0.360	8.549	1.452
4	0.603	0.311	7.008	1.761
5	0.627	0.276	6.075	2.043
10	0.654	0.109	3.634	3.623
20	0.654	0.020	2.386	6.800
30	0.667	0.003	1.895	10.009
40	0.674	0.000	1.640	13.256
50	0.674	0.000	1.457	16.550
100	0.682	0.000	1.008	33.362
500	0.687	0.000	0.448	189.693
1000	0.683	0.000	0.316	468.572

Notes: ‘Coverage prob.’ is the coverage probability of the 68 per cent weak-proxy robust confidence interval; ‘Unbounded’ is the proportion of Monte Carlo samples in which the confidence interval is unbounded; W is the Wald statistic for testing the null hypothesis that the covariance between the proxy and y_{1t} is zero.

approach. In particular, when the sign of the shock is known in a small number of periods, the weak-proxy robust approach to inference under NP yields confidence intervals with coverage that may differ markedly from the nominal level; whether the confidence intervals under- or over-cover will depend on the confidence level. In contrast, the robust Bayesian approach to inference under NR generates credible intervals with a coverage probability that always exceeds the nominal level. In both cases, the coverage probabilities converge towards the nominal level as K increases, but our Monte Carlo results indicate that the rate of convergence is slower under the NR approach than under the NP approach.

One limitation of this exercise is that it abstracts from the common practice of imposing additional identifying restrictions alongside shock-sign NR. For example, it is common to also impose sign restrictions on functions of the structural parameters, such as impulse responses, as well as other types of NR, such as restrictions on the historical decomposition (e.g. Antolín-Díaz and Rubio-Ramírez 2018). A natural sign restriction in this context – which we do not impose – is that the impulse response of the first variable to the first shock is nonnegative on impact. We have deliberately made this choice so that we are using the same information under both approaches

Table 3: Robust Bayesian Inference – Monte Carlo Results for $\alpha = 0.95$

K	Coverage prob.	Unbounded	Median width
1	1.000	0.766	∞
2	1.000	0.596	∞
3	1.000	0.461	28.343
4	1.000	0.345	8.725
5	0.999	0.271	6.083
10	0.998	0.075	2.742
20	0.996	0.005	1.571
30	0.994	0.000	1.226
40	0.993	0.000	1.034
50	0.993	0.000	0.932
100	0.986	0.000	0.717
500	0.966	0.000	0.545
1000	0.955	0.000	0.522

Notes: ‘Coverage prob.’ is the coverage probability of the 95 per cent robust credible interval; ‘Unbounded’ is the proportion of Monte Carlo samples in which the robust credible interval is unbounded.

to identification and inference. The imposition of additional restrictions under the NR approach can tighten the robust credible intervals and yield frequentist coverage probabilities closer to, but still exceeding, the nominal level.

5 Conclusion

Although it has become increasingly common to impose ‘narrative restrictions’ in SVARs, there has been little formal work exploring econometric methods for imposing these restrictions. Using the information underlying narrative restrictions about the signs of structural shocks in particular periods to construct a proxy variable – as proposed by Plagborg-Møller and Wolf (2021a) – is a potentially useful alternative approach that in principle can be implemented using (relatively) standard inferential procedures.

We show that the ‘narrative proxy’ is likely to be weak when the sign of the shock is known in only a small number of periods, which tends to be the case empirically. Furthermore, under this proxy approach, the weak-proxy robust confidence intervals developed in MSW are not asymptotically valid when the sign of the shock is known in a relatively small number of periods. Monte Carlo simulations suggest that

Table 4: Robust Bayesian Inference – Monte Carlo Results for $\alpha = 0.68$

K	Coverage prob.	Unbounded	Median width
1	1.000	0.757	∞
2	0.998	0.581	∞
3	0.997	0.442	19.893
4	0.995	0.330	7.615
5	0.994	0.256	5.422
10	0.985	0.066	2.436
20	0.970	0.004	1.302
30	0.955	0.000	0.969
40	0.947	0.000	0.778
50	0.932	0.000	0.678
100	0.882	0.000	0.466
500	0.757	0.000	0.296
1000	0.718	0.000	0.275

Notes: ‘Coverage prob.’ is the coverage probability of the 68 per cent robust credible interval; ‘Unbounded’ is the proportion of Monte Carlo samples in which the robust credible interval is unbounded.

the coverage probability of these intervals may be higher or lower than the nominal level depending on the confidence level. The coverage probability of these intervals approaches the desired level when the sign of the shock is known in a fairly large number of periods. In contrast, the robust Bayesian credible intervals from GKR – which are applicable when shock-sign restrictions are imposed in an otherwise set-identified model – demonstrate coverage probabilities exceeding the nominal level but that converge to the nominal level as the number of shock-sign restrictions increases.

References

- ANTOLÍN-DÍAZ, J. AND J. F. RUBIO-RAMÍREZ (2018): “Narrative Sign Restrictions for SVARs,” *American Economic Review*, 108, 2802–29.
- ARIAS, J. E., J. F. RUBIO-RAMÍREZ, AND D. F. WAGGONER (2018): “Inference Based on Structural Vector Autoregressions Identified with Sign and Zero Restrictions: Theory and Applications,” *Econometrica*, 86, 685720.
- BAUMEISTER, C. AND J. D. HAMILTON (2015): “Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information,” *Econometrica*, 83, 1963–1999.
- BEN ZEEV, N. (2018): “What Can We Learn About News Shocks from the Late 1990s and Early 2000s Boom-bust Period?” *Journal of Economic Dynamics and Control*, 87, 94–105.
- BERGER, T., J. RICHTER, AND B. WONG (2022): “A Unified Approach for Jointly Estimating the Business and Financial Cycle, and the Role of Financial Factors,” *Journal of Economic Dynamics and Control*, 104315.
- BOER, L. AND H. LÜTKEPOHL (2021): “Qualitative versus Quantitative External Information for Proxy Vector Autoregressive Analysis,” *Journal of Economic Dynamics and Control*, 127, 104118.
- BUDNIK, K. AND G. RÜNSTLER (2020): “Identifying SVARs from Sparse Narrative Instruments: Dynamic Effects of U.S. Macroprudential Policies,” European Central Bank Working Paper No. 2353.
- CAGGIANO, G., E. CASTELNUOVO, S. DELRIO, AND R. KIMA (2021): “Financial Uncertainty and Real Activity: The Good, the Bad, and the Ugly,” *European Economic Review*, 136, 103750.
- CHENG, K. AND Y. YANG (2020): “Revisiting the Effects of Monetary Policy Shocks: Evidence from SVAR with Narrative Sign Restrictions,” *Economics Letters*, 196, 109598.
- FURLANETTO, F. AND Ø. ROBSTAD (2019): “Immigration and the Macroeconomy: Some New Empirical Evidence,” *Review of Economic Dynamics*, 34, 1–19.

- GIACOMINI, R. AND T. KITAGAWA (2021): “Robust Bayesian Inference for Set-identified Models,” *Econometrica*, 89, 1519–1556.
- GIACOMINI, R., T. KITAGAWA, AND M. READ (2021a): “Identification and Inference Under Narrative Restrictions,” arXiv: 2102.06456 [econ.EM].
- (2021b): “Robust Bayesian Analysis for Econometrics,” Centre for Economic Policy Research Discussion Paper DP16488.
- INOUE, A. AND L. KILIAN (in press): “Joint Bayesian Inference about Impulse Responses in VAR Models,” *Journal of Econometrics*.
- KILIAN, L. AND X. ZHOU (2020a): “Does Drawing Down the US Strategic Petroleum Reserve Help Stabilize Oil Prices?” *Journal of Applied Econometrics*, 35, 673–691.
- (2020b): “Oil Prices, Exchange Rates and Interest Rates,” Center for Financial Studies Working Paper Series No. 646.
- LAUMER, S. (2020): “Government Spending and Heterogeneous Consumption Dynamics,” *Journal of Economic Dynamics and Control*, 114, 103868.
- LUDVIGSON, S. C., S. MA, AND S. NG (2018): “Shock Restricted Structural Vector-Autoregressions,” National Bureau of Economic Research Working Paper No. 23225.
- (2021): “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?” *American Economic Journal: Macroeconomics*, 13, 369–410.
- LUNSFORD, K. G. (2015): “Identifying Structural VARs with a Proxy Variable and a Test for a Weak Proxy,” Federal Reserve Bank of Cleveland Working Paper 15-28.
- MAFFEI-FACCIOLI, N. AND E. VELLA (2021): “Does Immigration Grow the Pie? Asymmetric Evidence from Germany,” *European Economic Review*, 138, 103846.
- MERTENS, K. AND M. O. RAVN (2013): “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States,” *American Economic Review*, 103, 1212–47.

- MONTIEL-OLEA, J. L., J. H. STOCK, AND M. W. WATSON (2021): “Inference in Structural Vector Autoregressions Identified with an External Instrument,” *Journal of Econometrics*, 225, 74–87.
- PLAGBORG-MØLLER, M. AND C. K. WOLF (2021a): “Local Projections and VARs Estimate the Same Impulse Responses,” *Econometrica*, 89, 955–980.
- (2021b): “Supplement to “Local Projections and VARs Estimate the Same Impulse Responses”,” *Econometrica*, 89.
- READ, M. (forthcoming): “The Unit-effect Normalisation in Set-identified Structural Vector Autoregressions,” Reserve Bank of Australia Research Discussion Paper.
- REDL, C. (2020): “Uncertainty Matters: Evidence from Close Elections,” *Journal of International Economics*, 103296.
- STOCK, J. H. AND M. W. WATSON (2016): “Chapter 8 – Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, 415–525.
- (2018): “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *The Economic Journal*, 128, 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, 381–419.
- ZHOU, X. (2020): “Refining the Workhorse Oil Market Model,” *Journal of Applied Econometrics*, 35, 130–140.

Appendices

A Conditional Identified Sets

We first present analytical expressions for the conditional identified set for θ obtained under a single shock-sign NR. We consider the case where the restriction is that

$\varepsilon_{1t} \geq 0$, which is the case considered in GKR, as well as the case where $\varepsilon_{1t} \leq 0$. These analytical expressions are used to implement our Monte Carlo exercise; computing the bounds of the conditional identified set using existing numerical methods would be computationally demanding given the large number of Monte Carlo replications. After presenting expressions for the conditional identified set for θ under a single shock-sign NR, we explain how to use these to obtain the conditional identified set for η_{21} under a set of shock-sign NR.

The expressions for the conditional identified set differ depending on the sign of σ_{21} . The data-generating process satisfies $\sigma_{21} < 0$. Since the sample length T is sufficiently large in our Monte Carlo exercise, $\sigma_{21} < 0$ holds with posterior probability one almost-surely under the sampling distribution of the data. Consequently, for our exercise, it suffices to consider only the case where $\sigma_{21} < 0$.

In what follows, let $h(\boldsymbol{\phi}, \mathbf{y}_t) = \sigma_{21}y_{1t} - \sigma_{11}y_{2t}$ and $C(\boldsymbol{\phi}, \mathbf{y}_t) = \sigma_{22}y_{1t}/h(\boldsymbol{\phi}, \mathbf{y}_t)$.

Case 1: Consider the shock-sign restriction

$$\varepsilon_{1t} = \mathbf{e}'_1 \mathbf{A}_0 \mathbf{y}_t = (\sigma_{11}\sigma_{22})^{-1} (\sigma_{22}y_{1t} \cos \theta + (\sigma_{11}y_{2t} - \sigma_{21}y_{1t}) \sin \theta) \geq 0. \quad (29)$$

Under the sign normalization and the shock-sign restriction, θ is restricted to the set

$$\begin{aligned} \theta \in \{ & \theta : \sigma_{21} \sin \theta \leq \sigma_{22} \cos \theta, \cos \theta \geq 0, \sigma_{22}y_{1t} \cos \theta \geq (\sigma_{21}y_{1t} - \sigma_{11}y_{2t}) \sin \theta \} \\ & \cup \{ \theta : \sigma_{21} \sin \theta \leq \sigma_{22} \cos \theta, \cos \theta \leq 0, \sigma_{22}y_{1t} \cos \theta \geq (\sigma_{21}y_{1t} - \sigma_{11}y_{2t}) \sin \theta \}. \end{aligned} \quad (30)$$

First, consider the case where $h(\boldsymbol{\phi}, \mathbf{y}_t) < 0$. In this case, the conditional identified set for θ is

$$\theta \in \left[\arctan \left(\max \left\{ \frac{\sigma_{22}}{\sigma_{21}}, C(\boldsymbol{\phi}, \mathbf{y}_t) \right\} \right), \pi + \arctan \left(\min \left\{ \frac{\sigma_{22}}{\sigma_{21}}, C(\boldsymbol{\phi}, \mathbf{y}_t) \right\} \right) \right]. \quad (31)$$

Next, assume that $h(\boldsymbol{\phi}, \mathbf{y}_t) > 0$, in which case

$$\theta \in \begin{cases} \left[\arctan \left(\frac{\sigma_{22}}{\sigma_{21}} \right), \arctan(C(\boldsymbol{\phi}, \mathbf{y}_t)) \right], & \text{if } \frac{\sigma_{22}}{\sigma_{21}} < C(\boldsymbol{\phi}, \mathbf{y}_t) \\ \left[\pi + \arctan(C(\boldsymbol{\phi}, \mathbf{y}_t)), \pi + \arctan \left(\frac{\sigma_{22}}{\sigma_{21}} \right) \right], & \frac{\sigma_{22}}{\sigma_{21}} > C(\boldsymbol{\phi}, \mathbf{y}_t). \end{cases} \quad (32)$$

Case 2: Under the sign normalization and the shock-sign restriction that $\varepsilon_{1t} \leq 0$, θ

is restricted to the set

$$\begin{aligned} \theta \in & \{ \theta : \sigma_{21} \sin \theta \leq \sigma_{22} \cos \theta, \cos \theta \geq 0, \sigma_{22} y_{1t} \cos \theta \leq (\sigma_{21} y_{1t} - \sigma_{11} y_{2t}) \sin \theta \} \\ & \cup \{ \theta : \sigma_{21} \sin \theta \leq \sigma_{22} \cos \theta, \cos \theta \leq 0, \sigma_{22} y_{1t} \cos \theta \leq (\sigma_{21} y_{1t} - \sigma_{11} y_{2t}) \sin \theta \}. \end{aligned} \quad (33)$$

First, assume that $h(\boldsymbol{\phi}, \mathbf{y}_t) < 0$, in which case the conditional identified set for θ takes the form

$$\theta \in \begin{cases} \left[\arctan \left(\frac{\sigma_{22}}{\sigma_{21}} \right), \arctan (C(\boldsymbol{\phi}, \mathbf{y}_t)) \right], & \text{if } \frac{\sigma_{22}}{\sigma_{21}} < C(\boldsymbol{\phi}, \mathbf{y}_t) \\ \left[\pi + \arctan (C(\boldsymbol{\phi}, \mathbf{y}_t)), \pi + \arctan \left(\frac{\sigma_{22}}{\sigma_{21}} \right) \right], & \frac{\sigma_{22}}{\sigma_{21}} > C(\boldsymbol{\phi}, \mathbf{y}_t). \end{cases} \quad (34)$$

Next, assume that $h(\boldsymbol{\phi}, \mathbf{y}_t) > 0$, in which case

$$\theta \in \left[\arctan \left(\max \left\{ \frac{\sigma_{22}}{\sigma_{21}}, C(\boldsymbol{\phi}, \mathbf{y}_t) \right\} \right), \pi + \arctan \left(\min \left\{ \frac{\sigma_{22}}{\sigma_{21}}, C(\boldsymbol{\phi}, \mathbf{y}_t) \right\} \right) \right]. \quad (35)$$

Obtaining the Conditional Identified Set for η_{21} . Let $\ell(\boldsymbol{\phi}, \mathbf{y}_t)$ and $u(\boldsymbol{\phi}, \mathbf{y}_t)$ be the lower and upper bounds, respectively, of the conditional identified set for θ given a shock-sign restriction imposed in period t . The conditional identified set for θ given a set of shock-sign restrictions imposed in the first K periods is $[\max_{t \in \{1, \dots, K\}} \ell(\boldsymbol{\phi}, \mathbf{y}_t), \min_{t \in \{1, \dots, K\}} u(\boldsymbol{\phi}, \mathbf{y}_t)]$. The conditional identified set for η_{21} can be obtained by projecting the conditional identified set for θ into impulse-response space using Equation (6).

If the conditional identified set for θ does not include $-\pi/2$ and $\pi/2$, then $\eta_{21}(\boldsymbol{\phi}, \theta)$ is increasing in θ . Consequently, in this case the lower and upper bounds of the conditional identified set for η_{21} can be obtained by plugging in the lower and upper bounds of the conditional identified set for θ into Equation (6). If the conditional identified set for θ includes $-\pi/2$ or $\pi/2$, the conditional identified set for η_{21} will be unbounded. For example, consider a case where the conditional identified set for θ is such that $\ell(\boldsymbol{\phi}, \mathbf{y}_t) < \pi/2 < u(\boldsymbol{\phi}, \mathbf{y}_t)$. For $\theta \nearrow \pi/2$, $\tan \theta \rightarrow \infty$, so $\eta_{21} \rightarrow \infty$. For $\theta \searrow \pi/2$, $\tan \theta \rightarrow -\infty$, so $\eta_{21} \rightarrow -\infty$. Hence, the conditional identified set for η_{21} is $(-\infty, \infty)$.

A Note on Non-emptiness of the Conditional Identified Set. In principle, the conditional identified set may be empty at particular values of $\boldsymbol{\phi}$ and realisations of the data. In empirical applications, it is therefore necessary to check whether the

conditional identified set is nonempty before attempting to compute (or approximate) the bounds of the conditional identified set. However, in our Monte Carlo exercise, the sample size is sufficiently large such that the posterior probability that the conditional identified set is nonempty is one almost-surely under the sampling distribution of the data; intuitively, if the shock-sign NR are correctly specified, there must exist a value of θ satisfying the restrictions when the values of the reduced-form parameters are sufficiently close to their true values. This means that it is not necessary to check whether the conditional identified set is nonempty at each draw of ϕ from its posterior, which greatly reduces the computational demands of the Monte Carlo exercise.