

Inference for VARs identified with sign restrictions

ELEONORA GRANZIERA

Monetary Policy and Research Department, Bank of Finland

HYUNGSIK ROGER MOON

Department of Economics, University of Southern California and School of Economics, Yonsei University

FRANK SCHORFHEIDE

Department of Economics, University of Pennsylvania

There is a fast growing literature that set-identifies structural vector autoregressions (SVARs) by imposing sign restrictions on the responses of a subset of the endogenous variables to a particular structural shock (sign-restricted SVARs). Most methods that have been used to construct pointwise coverage bands for impulse responses of sign-restricted SVARs are justified only from a Bayesian perspective. This paper demonstrates how to formulate the inference problem for sign-restricted SVARs within a moment-inequality framework. In particular, it develops methods of constructing confidence bands for impulse response functions of sign-restricted SVARs that are valid from a frequentist perspective. The paper also provides a comparison of frequentist and Bayesian coverage bands in the context of an empirical application—the former can be substantially wider than the latter.

KEYWORDS. Bayesian inference, frequentist inference, set-identified models, sign restrictions, structural VARs.

JEL CLASSIFICATION. C1, C32.

1. INTRODUCTION

During the three decades following Sims (1980) “Macroeconomics and Reality,” structural vector autoregressions (SVARs) have become an important tool in empirical macroeconomics. They have been used for macroeconomic forecasting and policy analysis, as well as to investigate the sources of business cycle fluctuations and to provide a

Eleonora Granziera: eleonora.granziera@bof.fi

Hyungsik Roger Moon: moonr@usc.edu

Frank Schorfheide: schorf@ssc.upenn.edu

We thank Andres Santos (coeditor), Fabio Canova, Eric Renault, Paul Sangrey, James Stock, several anonymous referees, as well as participants at various conferences and seminars for helpful comments. We also thank Mihye Lee for her contributions to the first draft of this paper and Minchul Shin for research assistance. Schorfheide gratefully acknowledges financial support from the National Science Foundation under Grants SES 1061725 and 1424843. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Finland. The Online Technical Appendix as well as data and software to replicate the empirical analysis are available in the Supplementary Material (Granziera, Moon, and Schorfheide (2018)).

benchmark against which modern dynamic macroeconomic theories are evaluated. The most controversial step in the specification of a structural VAR is the mapping between reduced form one-step-ahead forecast errors and orthogonalized, economically interpretable structural innovations. Traditionally, SVARs have been constructed by imposing sufficiently many restrictions such that the relationship between structural innovations and forecast errors is one-to-one.

Over the past decade, starting with Faust (1998), Canova and De Nicolò (2002), and Uhlig (2005), empirical researchers have used more agnostic approaches that generate bounds on structural impulse response functions by restricting the sign of certain responses. We refer to this class of models as sign-restricted SVARs. They have been employed, for instance, to measure the effects of monetary policy shocks (e.g., Faust (1998), Canova and De Nicolò (2002), Uhlig (2005)), technology shocks (e.g., Dedola and Neri (2007), Peersman and Straub (2009)), government spending shocks (Mountford and Uhlig (2009), Pappa (2009)), oil price shocks (e.g., Baumeister and Peersman (2013), Kilian and Murphy (2012)), and financial shocks (e.g., Hristov, Hülsewig, and Wollmershäuser (2012), Gambetti and Musso (2017)).

Because impulse responses in sign-restricted SVARs can only be restricted to a bounded set, they belong to the class of set-identified econometric models, using the terminology of Manski (2003).¹ With the exception of Faust, Rogers, Swanson, and Wright (2003) and Faust, Swanson, and Wright (2004) (since the two papers are methodologically equivalent, we are using the abbreviation FRSW to refer to both of them), researchers have exclusively reported Bayesian credible bands for sign-restricted VARs, and a general method for constructing uniformly asymptotically valid frequentist confidence intervals was absent from the literature when the first draft of this paper was written; see Moon, Schorfheide, Granziera, and Lee (2011). As shown in detail in Moon and Schorfheide (2012), the large-sample numerical equivalence of frequentist confidence sets and Bayesian credible sets breaks down in set-identified models, which means that Bayesian credible bands should not be interpreted as approximate frequentist confidence bands.²

The goal of this paper is to provide researchers with an easy-to-use tool to construct valid frequentist confidence bands for impulse responses and other measures of the dynamic effects of structural shocks (e.g., variance decompositions) of sign-restricted SVARs. The specific contributions are the following: First, we formulate the problem of analyzing set-identified sign-restricted SVAR models in a moment-inequality-based minimum distance framework. Second, we find an easily interpretable sufficient condition for the nonemptiness of the identified set of sign-restricted structural impulse responses, and we propose a consistent estimator of the identified set that is straightforward to compute. Third, using our minimum distance framework, we formally analyze

¹The microeconometrics literature uses the terms *set* and *partially* identified model interchangeably. In the VAR literature, a *partially* identified structural VAR is one in which the researcher tries to identify only a subset of the structural shocks. To avoid confusion, we shall use the term *set* identified because we are focusing on models in which impulse responses can only be bounded.

²Treatments of Bayesian inference in sign-restricted SVARs can be found, for instance, in Uhlig (2005), Rubio-Ramirez, Waggoner, and Zha (2010), Baumeister and Hamilton (2015), Kilian and Lütkepohl (2017), and the references cited therein.

Bonferroni confidence sets. Fourth, we provide step-by-step recipes for practitioners on how to compute these confidence sets.³

At an abstract level, our inference problem is characterized by a vector of point-identified reduced-form parameters ϕ , a vector of structural parameters (impulse responses or variance decompositions) θ , and a vector of nuisance parameters q . The sign restrictions generate an identified set for q , $F^q(\phi)$. Conditional on q and ϕ the vector θ is point identified, but because q is set-identified, so is θ and we denote its identified set by $F^\theta(\phi)$. To obtain a confidence set for θ , we pursue a Bonferroni approach: We construct a confidence interval for the set-identified nuisance parameter q and then take the union of standard Wald confidence sets for θ that are generated conditional on all q in the first-stage confidence set. The Bonferroni inequality is used to ensure the desired coverage probability of the resulting confidence set for θ . We also show that the plug-in estimator $F^\theta(\hat{\phi})$ delivers a consistent estimate of the identified set for θ , denoted by $F^\theta(\phi)$.

In the first draft of this paper, we used a projection approach instead; see Moon et al. (2011).⁴ We constructed a joint confidence interval for the set-identified pair (q, θ) and projected it onto the θ ordinate. Here, θ is the response of a particular variable at a particular horizon to a particular shock. In order to generate point-wise confidence bands for impulse response functions, we repeated the computations for different definitions of θ . In subsequent research, we compared the projection approach and the Bonferroni approach and found that there is no clear ranking of the two types of confidence sets. However, the Bonferroni approach has a clear computational advantage: the irregular confidence set for the nuisance parameter q only has to be computed once. Conditional on $F^q(\phi)$, one can easily generate standard confidence sets for impulse responses of different variables at different horizons, of vectors of responses, and of variance decompositions and then take unions over q . Thus, we decided to focus on the Bonferroni set in the current version of the paper.

The Bonferroni approach has a long history in the time series literature. For instance, Cavanagh, Elliott, and Stock (1995) and Campbell and Yogo (2006) used it to eliminate nuisance parameters that characterize the persistence of error terms or regressors. In the context of structural VARs, the Bonferroni approach has been used by FRSW. However, FRSW's setup is quite different from ours. In their framework, the set identification of q arises from a rank deficiency in equality restrictions, which depend on estimated parameters. While FRSW restrict q further by imposing inequality conditions, these inequality conditions do not depend on estimated parameters. Our analysis, on the other hand, focuses on inequality restrictions for q that may bind and depend on estimated parameters. This generalization is essential to cover the wide range of empirical applications referenced above. In the Monte Carlo analysis, we explore ideas by Campbell and Yogo (2006) and McCloskey (2017) to tighten the Bonferroni sets.

³The contribution of this paper is meant to be positive. We do not criticize the use of Bayesian inference methods as long as it is understood that their output needs to be interpreted from a Bayesian perspective. We provide applied researchers who are interested in impulse response confidence bands that are valid from a frequentist perspective with econometric tools to compute such bands.

⁴Inference procedures for subvectors have been further developed by Chaudhuri and Zivot (2011), Kaido, Molinari, and Stoye (2016), Andrews (2017), and Bugni, Canay, and Shi (2017).

Building upon recent advances in the moment-inequality literature in microeconomics, in particular, [Chernozhukov, Hong, and Tamer \(2007\)](#), [Rosen \(2008\)](#), [Andrews and Guggenberger \(2009\)](#), and [Andrews and Soares \(2010\)](#), we provide an asymptotic analysis of the Bonferroni approach to constructing confidence sets for the dynamic effects of shocks in sign-restricted SVARs. Because the number of linearly independent moment conditions is a function of the nuisance parameter q , we need to modify some of the existing microeconomic theory. As is common in the literature, we use a point-wise testing procedure to obtain a confidence set for q . We use Andrews and Soares' (2010) moment selection procedure to tighten the critical values for the point-wise testing procedures.⁵ We adapt their theory to account for the q -dependent rank of the set of moment inequalities in our model and prove that the proposed confidence sets are uniformly asymptotically valid. Our results on the nonemptiness of identified sets for impulse responses complement the equality-restriction-based VAR identification results reported in [Rubio-Ramirez, Waggoner, and Zha \(2010\)](#).

Since the working paper versions of this paper have been written, a series of alternative approaches for the construction of frequentist confidence bands of sign-restricted SVARs have been proposed. [Gafarov, Meier, and Montiel \(2016\)](#) developed an alternative projection-based approach that starts from a Wald confidence ellipsoid for the reduced-form VAR parameters and takes unions of the identified sets $F^\theta(\phi)$. Because it relies on the regular behavior of the estimator of the reduced-form coefficients, this method also has the interpretation of delivering Bayesian credible sets for the identified set $F^\theta(\phi)$. The downside is that it is quite conservative. [Gafarov, Meier, and Montiel \(2018\)](#) propose a δ -method confidence interval for sign-restricted SVARs which relies on a closed-form characterization of the endpoints of the identified set. While the resulting intervals are less conservative, their drawbacks are that they are only pointwise but not uniformly consistent, and that they can only be applied to scalar θ 's. Finally, [Giacomini and Kitagawa \(2015\)](#) construct robust Bayesian credible sets for impulse response functions in set-identified SVARs that have good frequentist properties.

The remainder of the paper is organized as follows. Section 2 develops the notation used in this paper and provides a simple example of a sign-restricted SVAR. We describe how set-identification arises in this model and sketch the Bonferroni approach to the construction of confidence intervals for the dynamic effects of structural shocks. Section 3 is geared toward practitioners and discusses the step-by-step implementation and computational aspects of the proposed inference method in the context of a general SVAR. Technical assumptions and large sample results are presented in Section 4. Some extensions are discussed in Section 5. To illustrate the proposed methods, we conduct a Monte Carlo study in Section 6 and generate confidence bands for output, inflation, interest rate, and money responses to a monetary policy shock in an empirical application in Section 7. Finally, Section 8 concludes. Proofs and detailed derivations as well as further information about the Monte Carlo experiments and the empirical analysis are

⁵A recent survey of the moment-inequality literature is provided by [Canay and Shaikh \(2017\)](#). Alternative procedures include [Andrews and Barwick \(2012\)](#) and [Romano, Shaikh, and Wolf \(2014\)](#). Because of a potentially large number of inequality restrictions, the refinements to the moment selection proposed in [Andrews and Barwick \(2012\)](#) did not seem practical.

relegated to the online Appendix in the Supplementary Material (Granziera, Moon, and Schorfheide (2018)).

We use the following notation throughout the remainder of the paper: $\mathcal{I}\{x \geq a\}$ is the indicator function that is one if $x \geq a$ and zero otherwise. $0_{n \times m}$ is an $n \times m$ matrix of zeros and I_n is the $n \times n$ identity matrix. We use $[AB]_{(i)}$ to denote the i th row of the matrix $(A \cdot B)$ and $[AC]_i$ to denote the i th element of the vector $(A \cdot C)$. \otimes is the Kronecker product, $\text{vec}(\cdot)$ stacks the columns of a matrix, and $\text{vech}(\cdot)$ vectorizes the lower triangular part of a square matrix. We use $\text{diag}(A_1, \dots, A_k)$ to denote a quasi-diagonal matrix with submatrices A_1, \dots, A_k on its diagonal and zeros elsewhere. If A is an $n \times m$ matrix, then $\|A\|_W = \sqrt{\text{tr}[W A' A]}$. In the special case of a vector, our definition implies that $\|A\|_W = \sqrt{A' W A}$. If the weight matrix is the identity matrix, we omit the subscript. We write $x \gg 0$ to mean that all elements of the vector x are strictly greater than zero; we write $x > 0$ to mean that all elements of x are greater than or equal to zero but not all equal to zero, that is, $x \neq 0$; finally, we write $x \geq 0$ to mean that all elements of x are greater than or equal to zero. We use \propto to indicate proportionality and “ \xrightarrow{p} ” and “ \implies ” to indicate convergence in probability and convergence in distribution, respectively, as $T \rightarrow \infty$. A multivariate normal distribution is denoted by $N(\mu, \Sigma)$. We use χ_m^2 to denote a χ^2 distribution with m degrees of freedom.

2. GENERAL SETUP AND ILLUSTRATIVE EXAMPLE

Throughout this paper, we consider an n -dimensional VAR with p lags, which takes the form

$$y_t = A_1 y_{t-1} + \dots + A_p y_{t-p} + u_t, \quad \mathbb{E}[u_t | \mathcal{F}_{t-1}] = 0, \quad \mathbb{E}[u_t u_t' | \mathcal{F}_{t-1}] = \Sigma_u. \quad (1)$$

Here, y_t is an $n \times 1$ vector and the information set $\mathcal{F}_{t-1} = \{y_{t-1}, y_{t-2}, \dots\}$ is composed of the lags of y_t . Constants and deterministic trend terms are omitted because they are irrelevant for the subsequent discussion. The one-step-ahead forecast errors (reduced-form shocks) u_t are linear functions of a vector of fundamental innovations (structural shocks) ε_t :

$$u_t = A_\varepsilon \varepsilon_t = \Sigma_{tr} \Omega_\varepsilon \varepsilon_t, \quad \mathbb{E}[\varepsilon_t | \mathcal{F}_{t-1}] = 0, \quad \mathbb{E}[\varepsilon_t \varepsilon_t' | \mathcal{F}_{t-1}] = I_n, \quad (2)$$

where Σ_{tr} is the lower triangular Cholesky factor of Σ_u and Ω_ε is an arbitrary orthogonal matrix. Assuming that the lag polynomial associated with the VAR in (1) is invertible, one can express y_t as the following infinite-order vector moving average (VMA) process:

$$y_t = \sum_{h=0}^{\infty} C_h(A_1, \dots, A_p) \Sigma_{tr} \Omega_\varepsilon \varepsilon_{t-h}. \quad (3)$$

We assume that the object of interest is the propagation of the first shock, $\varepsilon_{1,t}$, and denote the first column of the matrix Ω_ε by q , where q is a unit-length vector. The domain of q is the n -dimensional unit sphere $\mathbb{S}^n = \{q \in \mathbb{R}^n \mid \|q\| = 1\}$. In Section 2.1, we discuss how imposing sign restrictions on some impulse responses generates set identification

of the dynamic effects of $\varepsilon_{1,t}$. Section 2.2 provides a bivariate VAR(0) illustration. Section 2.3 introduces some important notation. We present the construction of a Bonferroni confidence set for the dynamic effects of $\varepsilon_{1,t}$ in Section 2.4. The Bonferroni set is based on a confidence set for q , which is described in Section 2.5.

2.1 Sign restrictions and set identification

The SVAR identification problem arises because the one-step-ahead forecast error covariance matrix Σ_u is invariant to the orthogonal matrix Ω_ε , which implies that Ω_ε and its first column q are not identifiable from the data. Point identification could be achieved by selecting a particular q as a function of $(A_1, \dots, A_p, \Sigma_{\text{tr}})$. The recent SVAR literature has pursued a more agnostic approach and restricted the set of admissible q by a collection of sign restrictions on impulse responses.⁶ The impulse response of variable $y_{i,t}$ to $\varepsilon_{1,t}$ at horizon h is given by

$$\text{IRF}(i, h \mid \varepsilon_{1,t} = 1) = [C_h(A_1, \dots, A_p)\Sigma_{\text{tr}}]_{(i)} q. \quad (4)$$

We define the $n \times 1$ vector

$$\phi_j = [C_h(A_1, \dots, A_p)\Sigma_{\text{tr}}]_{(i)}$$

as the responses of a variable i at horizon h to the vector of reduced-form innovations u_t and summarize the sign-restricted impulse responses as

$$\Phi'_q q \geq 0, \quad \text{where } \Phi_q = [\phi_1, \dots, \phi_r], \quad (5)$$

r is the number of restrictions, and Φ_q is a $n \times r$ matrix. The vectors ϕ_j are possibly multiplied by -1 to restrict the response to be weakly negative rather than weakly positive. Moreover, the notation in (5) is general enough to accommodate sign restrictions on cumulative impulse responses over \bar{h} periods, obtained from $\sum_{h=0}^{\bar{h}} C_h(A_1, \dots, A_p)\Sigma_{\text{tr}} q$.

The object of inference is a k -dimensional parameter defined as

$$\theta = f(\Phi_\theta(A_1, \dots, A_p, \Sigma_{\text{tr}}), q) \in \Theta \subset \mathbb{R}^k. \quad (6)$$

The parameter set Θ is chosen to be consistent with potential sign restrictions for elements of θ implied by (5). Our leading example of $f(\cdot)$ is a vector of impulse responses, which can be expressed as a linear function of the reduced-form impulse responses

$$f(\Phi_\theta(A_1, \dots, A_p, \Sigma_{\text{tr}}), q) = \Phi'_\theta q, \quad (7)$$

where the definition of Φ_θ is similar to the definition of Φ_q in (5). In addition to impulse responses, researchers often report variance decompositions. For instance, the fraction

⁶We assume that these sign restrictions do not encode equality restrictions (e.g., by representing $a = 0$ as $a \leq 0$ and $a \geq 0$). The extension to models that combine sign-restrictions and equality restrictions is deferred to Section 5.

of the one-step-ahead forecast error variance of variable $y_{1,t}$ explained by shock $\varepsilon_{1,t}$ is given by

$$f(\Phi_\theta(A_1, \dots, A_p, \Sigma_{tr}), q) = \frac{\iota_1' \Sigma_{tr} q q' \Sigma_{tr}' \iota_1}{\iota_1' \Sigma_{tr} \Sigma_{tr}' \iota_1},$$

where $\iota_1 = [1, 0, \dots, 0]'$ is a $n \times 1$ vector.

While Φ_q and Φ_θ can be consistently estimated, the vector q as well as the object of interest θ are only set identified. We use $F^q(\Phi_q)$ and $F^\theta(\Phi_q, \Phi_\theta)$ to denote the identified sets of q and θ , respectively. Formally, they are defined as

$$F^q(\Phi_q) = \{q \in \mathbb{S}^n \mid \Phi_q' q \geq 0\}, \tag{8}$$

$$F^\theta(\Phi_q, \Phi_\theta) = \{\theta \in \Theta \mid \exists q \in F^q(\Phi_q) \text{ s.t. } \theta = f(\Phi_\theta, q)\}. \tag{9}$$

The goal is to construct a confidence set for θ . As an intermediate step in a Bonferroni approach, we will also construct a confidence set for q .

2.2 Identified sets $F^q(\cdot)$ and $F^\theta(\cdot)$ in a bivariate VAR(0)

For concreteness, consider the following example. Suppose the vector y_t is composed of inflation and output growth and the one-step-ahead forecast errors are linear functions of structural demand and supply shocks, stacked in the vector $\varepsilon_t = [\varepsilon_{D,t}, \varepsilon_{S,t}]'$. In order to obtain bounds for the effects of a demand shock, we impose the sign restriction that, contemporaneously, a demand shock moves prices and output in the same direction and the normalization restriction that a positive demand shock increases prices:

$$\phi_1' q = [\Sigma_{11}^{tr} \quad 0] \begin{bmatrix} q_1 \\ q_2 \end{bmatrix} \geq 0, \quad \phi_2' q = [\Sigma_{21}^{tr} \quad \Sigma_{22}^{tr}] \begin{bmatrix} q_1 \\ q_2 \end{bmatrix} \geq 0, \quad \Phi_q = [\phi_1, \phi_2]. \tag{10}$$

Suppose that the object of interest, θ , is the contemporaneous inflation response to a demand shock $\varepsilon_{D,t}$:

$$\theta = \phi_1' q, \quad \Phi_\theta = \phi_1. \tag{11}$$

Figure 1 provides an illustration of the two inequality constraints and the resulting identified sets. To simplify the graphical illustration, we assume that $\Sigma_{11}^{tr} = 1$ which implies $\theta = q_1$. Because the Cholesky factorization of the covariance matrix Σ_u is normalized such that Σ_{11}^{tr} and Σ_{22}^{tr} are nonnegative, (10) implies that

$$q_1 \geq 0 \quad \text{and} \quad q_2 \geq -\left(\frac{\Sigma_{21}^{tr}}{\Sigma_{22}^{tr}}\right) q_1. \tag{12}$$

The x -axis of both panels corresponds to $\theta = q_1$ and the y -axis represents q_2 . Each panel depicts a unit circle as well as the locus $\phi_2' q = 0$. In the left panel $\Sigma_{21}^{tr} < 0$, whereas in the right panel $\Sigma_{21}^{tr} > 0$. The identified set $F^q(\Phi_q)$ is given by the arc that ranges from the intersection of the unit circle with the y -axis to the intersection with $\phi_2' q = 0$. The identified set $F^\theta(\Phi_q, \Phi_\theta)$ is given by the projection of the arc onto the x -axis. The identified set is a singleton only if $\Sigma_{22}^{tr} = 0$ and $\Sigma_{21}^{tr} < 0$, which means that the one-step-ahead




FIGURE 1. Identified Sets for VAR(0).

forecast error covariance matrix is singular. We will rule out this case because it is not empirically relevant in practice.

2.3 Rank reductions and notation

In order to develop a notation for the general inference problem we need to accommodate two types of rank reductions. First, consider the matrix Φ_q for the bivariate VAR(0):

$$\Phi_q = \begin{bmatrix} \Sigma_{11}^{tr} & \Sigma_{21}^{tr} \\ 0 & \Sigma_{22}^{tr} \end{bmatrix}.$$

To construct a confidence set for q , we will replace the unknown Φ_q by a sample estimate $\hat{\Phi}_q$ and start from the high-level assumption that $\hat{\Phi}_q$ has a normal limit distribution. Because of the zeros in the Cholesky factorization of Σ_u (and possibly other restrictions imposed on the reduced-form parameters), the covariance matrix of $\hat{\Phi}_q$ has a rank-reduction. We circumvent this issues as follows. Let

$$\bar{S}' = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

be the selection matrix that deletes the zero elements in $\text{vec}(\Phi_q)$ such that

$$\phi_q = \bar{S}' \text{vec}(\Phi_q) \quad \text{and} \quad \bar{S} \phi_q = \text{vec}(\Phi_q).$$

Then we can express the sign restrictions as

$$\tilde{S}(q) \phi_q = (I \otimes q') \bar{S} \phi_q \geq 0. \tag{13}$$

The second type of rank reduction arises as follows. After eliminating $\Sigma_{12}^{tr} = 0$, we obtain $\phi_q = [\Sigma_{11}^{tr}, \Sigma_{21}^{tr}, \Sigma_{22}^{tr}]'$ in the bivariate VAR(0). In turn, the matrix $\tilde{S}(q)$ takes the form:

$$\tilde{S}(q) = \begin{bmatrix} q_1 & 0 & 0 \\ 0 & q_1 & q_2 \end{bmatrix}.$$

The first row of $\tilde{S}(q)$ equals zero if $q = [q_1, q_2]' = [0, 1]$. As a consequence the covariance matrix of $\tilde{S}(q)\hat{\phi}_q$ is singular at $q = [0, 1]'$ and cannot be inverted to form a weight matrix for the estimation of q . In order to eliminate the rows of zeros in $\tilde{S}(q)$, we introduce the selection matrix $V(q)$ and define

$$S(q) = V(q)\tilde{S}(q). \tag{14}$$

The row dimension of $S(q)$ is $r(q)$. By construction, the matrix $S(q)$ has the full row rank for all q .

Inference about q will be based on the objective function

$$G(q; \phi_q, W(\cdot)) = \min_{\mu \geq 0} \|S(q)\phi_q - V(q)\mu\|_{W(q)}^2. \tag{15}$$

The vector μ captures the slackness in the inequalities generated by the sign restrictions in (13). $W(q)$ is a symmetric and positive-definite weight matrix with a dimension that adjusts to the dimension of $V(q)$:

$$W(q) = V(q)\tilde{W}(q)V(q)'$$

The matrix $\tilde{W}(q)$ is a weight matrix that conforms with $\tilde{S}(q)\phi_q$. An important example of a $\tilde{W}(q)$ is the inverse of the asymptotic covariance matrix of $\sqrt{T}\tilde{S}(q)(\hat{\phi}_q - \phi_q)$. It can be verified that

$$q \in F^q(\phi_q) \text{ if and only if } G(q; \phi_q, W(\cdot)) = 0, \tag{16}$$

where we now write $F^q(\phi_q)$ instead of $F(\Phi_q)$.

2.4 A Bonferroni confidence interval for θ

Let S'_θ be the selection matrix that deletes the zeros other pre-determined elements in $\text{vec}(\Phi_\theta)$ and define $\phi_\theta = S'_\theta \text{vec}(\Phi_\theta)$. Moreover, let $\phi = [\phi'_q, \phi'_\theta]'$. We will often write $F^\theta(\phi)$ to abbreviate $F^\theta(\Phi_q, \Phi_\theta)$. The goal is to obtain a confidence interval $CS^\theta(\hat{\phi})$ that satisfies the condition

$$\liminf_{T \rightarrow \infty} \inf_{\rho \in \mathcal{R}} \inf_{\theta \in F^\theta(\phi(\rho))} P_\rho \{ \theta \in CS^\theta(\hat{\phi}) \} \geq 1 - \alpha. \tag{17}$$

The vector ϕ may be a subvector of a larger reduced-form parameter vector ρ with domain \mathcal{R} , that characterizes the distribution of the data y_1, \dots, y_T , which is why we write $\phi(\rho)$. For instance, in a Gaussian VAR ρ comprises the elements of A_1, \dots, A_p , and Σ_u ; see (1). The parameter θ does not appear as an index of the probability distribution P ,

because conditional on ϕ the parameter θ does not affect the distribution of the estimator of the reduced-form parameters $\hat{\phi}$. The confidence interval is indexed by $\hat{\phi}$ because it is a sufficient statistic in our setup.

As mentioned in the [Introduction](#), in our application it is quite natural to use a Bonferroni approach to compute an asymptotic $1 - \alpha$ confidence set for θ . Under the Bonferroni approach, one first constructs a confidence set for q . This is a “nonstandard” object because q is a set-identified parameter. Conditional on q , however, inference for θ becomes “standard,” because θ is point-identified. Let $\alpha = \alpha_1 + \alpha_2$. Bonferroni confidence steps can be obtained in three steps:

1. Construct a $1 - \alpha_1$ confidence interval $CS^q(\hat{\phi}_q)$ for q with the property that

$$\liminf_{T \rightarrow \infty} \inf_{\rho \in \mathcal{R}} \inf_{q \in F^q(\phi_q)} P_\rho \{q \in CS^q(\hat{\phi}_q)\} \geq 1 - \alpha_1. \tag{18}$$

2. Generate a confidence set for θ conditional on q . This is a “regular” problem because conditional on q , the vector θ is point identified. For instance, if θ is scalar and defined as $\theta = \phi'_\theta q$, one can use a Wald confidence interval of the form:

$$CS^{\theta}_q(\hat{\phi}_\theta) = \Theta \cap [\hat{\phi}'_\theta q - z_{\alpha_2/2} \hat{\sigma}_\theta, \hat{\phi}'_\theta q + z_{\alpha_2/2} \hat{\sigma}_\theta]. \tag{19}$$

Here, $z_{\alpha_2/2}$ is the two-sided α_2 critical value associated with the $N(0, 1)$ distribution and $\hat{\sigma}_\theta$ is a standard error estimate for $\hat{\theta} = \hat{\phi}'_\theta q$. The intersection of the symmetric confidence interval for θ can be used, for instance, to truncate the symmetric interval at zero, if θ is a response that is assumed to be nonnegative.

3. Construct the confidence set for θ by taking the following union of $CS^{\theta}_q(\hat{\phi}_\theta)$ sets:

$$CS^\theta(\hat{\phi}) = \bigcup_{q \in CS^q(\hat{\phi}_q)} CS^{\theta}_q(\hat{\phi}_\theta). \tag{20}$$

2.5 A confidence set for q

The main contribution of this paper is to adapt an inference procedure from the moment inequality literature to obtain a confidence set for q in the first step of the Bonferroni procedure. The confidence set is generated as a level set based on the sample analogue of the objective function in (15):

$$CS^q(\hat{\phi}_q) = \{q \in \mathbb{S}^n \mid G(q; \hat{\phi}_q, W(\cdot)) \leq c^{\alpha_1}(q)\}. \tag{21}$$

Here, $c^{\alpha_1}(q)$ is a critical value that guarantees that the confidence set satisfies (18). In the remainder of this subsection, we outline the derivation of the critical value $c^{\alpha_1}(q)$ for the bivariate VAR(0) example.

For illustrative purposes, suppose that the estimates of the reduced-form parameters have an exact standard normal distribution:

$$\sqrt{T}(\hat{\phi}_q - \phi_q) \sim N(0, I_3).$$

We parameterize the slackness in the inequality restrictions as

$$q_1 \phi_{q,1} = \tilde{\mu}_1, \quad q_1 \phi_{q,1} + q_2 \phi_{q,2} = \tilde{\mu}_2$$

and use the weight matrix that standardizes the distribution of $\tilde{S}(q)\hat{\phi}_q$:

$$W(q) = T \begin{bmatrix} \frac{1}{q_1^2} & 0 \\ 0 & 1 \end{bmatrix}.$$

In turn, we can express the sample analogue of the objective function in (15) as

$$\begin{aligned} G(q; \hat{\phi}_q, W(\cdot)) &= \min_{\mu \geq 0} [\sqrt{T}(\hat{\phi}_{q,1} - \phi_{q,1}) + \sqrt{T}(\mu_1 - \tilde{\mu}_1)/|q_1|]^2 \mathcal{I}\{q \neq [0, 1]'\} \\ &\quad + [\sqrt{T}(\hat{\phi}_{q,2} - \phi_{q,2})q_1 + \sqrt{T}(\hat{\phi}_{q,3} - \phi_{q,3})q_2 + \sqrt{T}(\mu_2 - \tilde{\mu}_2)]^2 \\ &= \min_{\nu \geq -\sqrt{T}\tilde{\mu}} [Z_1 - \nu_1/|q_1|]^2 \mathcal{I}\{q \neq [0, 1]'\} + [Z_2 - \nu_2]^2, \end{aligned}$$

where $\mu = [\mu_1, \mu_2]'$, $\nu = [\nu_1, \nu_2]'$, and Z_1 and Z_2 are two independent $N(0, 1)$ random variables.

A conservative upper bound on the sample objective function can be obtained by assuming that both inequalities are binding, that is, $\tilde{\mu}_1 = \tilde{\mu}_2 = 0$:

$$G(q; \hat{\phi}_q, W(\cdot)) \leq Z_1^2 \mathcal{I}\{Z_1 \leq 0\} + Z_2^2 \mathcal{I}\{Z_2 \leq 0\}.$$

Critical values for the distribution of the bound can be easily obtained by simulation. A sharper bound and a smaller critical value that leads to a smaller confidence set can be obtained by realizing that at most one inequality is binding. Thus, we will use the moment selection approach of [Andrews and Soares \(2010\)](#) to eliminate nonbinding moment conditions constructing critical values for $G(q; \hat{\phi}_q, W(\cdot))$. This means that in our illustrative example, the critical value can be essentially reduced to the $100(1 - \alpha_1)$ quantile of the distribution of $Z_1^2 \mathcal{I}\{Z_1 \leq 0\}$.

3. IMPLEMENTATION

This section focuses on the implementation of the proposed inference methods. A formal statement of assumptions and a rigorous analysis of the large sample properties of the confidence set will follow in Section 4. The remainder of this section is organized as follows. Section 3.1 briefly discusses the estimation of ϕ . Section 3.2 describes how we construct the confidence set for q . The calculation of confidence sets for θ given q is reviewed in Section 3.3. Finally, we provide some additional details for the computation of confidence bands for impulse responses in Section 3.4. Throughout this section, we assume that the impulse responses are not restricted through equality conditions (e.g., the restriction that certain responses must be zero). Extensions of our approach to a setting in which some identifying information is extracted from equality conditions are straightforward but notationally cumbersome and discussed in Section 5.

3.1 Estimating the reduced-form coefficients ϕ

We start from the assumption that $\hat{\phi}_q$ and $\hat{\phi}_\theta$ have Gaussian limit distributions and that the asymptotic covariance matrices can be estimated consistently:

$$\begin{aligned}\sqrt{T}(\hat{\phi}_q - \phi_q) &\implies N(0, \Lambda_{qq}) \quad \text{and} \quad \hat{\Lambda}_{qq} \xrightarrow{p} \Lambda_{qq} > 0, \\ \sqrt{T}(\hat{\phi}_\theta - \phi_\theta) &\implies N(0, \Lambda_{\theta\theta}) \quad \text{and} \quad \hat{\Lambda}_{\theta\theta} \xrightarrow{p} \Lambda_{\theta\theta} > 0.\end{aligned}\tag{22}$$

This assumption requires that all roots of the characteristic polynomial associated with the difference equation (1) lie outside of the unit circle. Throughout the paper, we rule out the presence of unit roots and assume that y_t is trend stationary.

We will also assume that Λ_{qq} and $\Lambda_{\theta\theta}$ are full rank. Because most impulse response function confidence bands are pointwise, the dimension of θ is typically one, which immediately leads to $\Lambda_{\theta\theta} > 0$. Whether or not $\Lambda_{qq} > 0$ is satisfied depends on the number of imposed sign restrictions. In a first-order approximation, the reduced-form responses stacked in ϕ_q are linear functions of the $n^2 p + n(n+1)/2$ coefficients in $(A_1, \dots, A_p, \Sigma_{\text{tr}})$. In order to restrict r structural responses, we need at least $nr - n(n-1)/2$ reduced form responses (recall that $n(n-1)/2$ responses upon impact are zero because Σ_{tr} is lower triangular). Thus, for $r > n(p+1)$ the matrix Λ_{qq} cannot be of full rank. In practice, most applications will satisfy the rank condition because we previously eliminated the $n(n-1)/2$ zero elements of the lower triangular matrix Σ_{tr} from the vector ϕ_q . The number of sign restrictions is usually small relative to the number of reduced-form VAR parameters.⁷

The VAR coefficient matrices A_1, \dots, A_p and Σ_u can be estimated by OLS. An estimate of Σ_{tr} is obtained by applying the Cholesky decomposition to $\hat{\Sigma}_u$. We then evaluate the functions $\Phi_q(\cdot)$ and $\Phi_\theta(\cdot)$ at $\hat{A}_1, \dots, \hat{A}_p, \hat{\Sigma}_{\text{tr}}$ to obtain $\hat{\phi}_q$ and $\hat{\phi}_\theta$. We obtain $\hat{\Lambda}_{qq}$ and $\hat{\Lambda}_{\theta\theta}$ by using a parametric bootstrap procedure: conditional on $\hat{A}_1, \dots, \hat{A}_p, \hat{\Sigma}_u$ we simulate n_Λ bootstrap samples $Y_{1:T}^*$ from the VAR in (1). Innovations u_t^* can either be drawn by resampling the residuals \hat{u}_t or by i.i.d. sampling from a $N(0, \hat{\Sigma}_u)$ distribution. In Sections 6 and 7, we do the latter. From each bootstrap sample, we compute $\hat{\phi}_q^*$. Finally, we compute the bootstrap sample covariance matrix of $\hat{\phi}_q^*$ and scale it appropriately to obtain $\hat{\Lambda}_{qq}$. The same approach is used to compute $\hat{\Lambda}_{\theta\theta}$.

3.2 Confidence set for q

The confidence interval for q is obtained by verifying whether

$$G(q; \hat{\phi}_q, W(\cdot)) \leq c^{\alpha_1}(q)$$

⁷Consider a 4-variable VAR(4) and suppose that the responses of 3 of the 4 variables are restricted upon impact and for the subsequent 3 periods. The number of estimated reduced-form coefficients is $4 \cdot 16 + 10 = 74$. In order to construct the sign-restricted responses, the number of elements in the vector ϕ_q is bounded by $3 \cdot 4 \cdot 4 - 3 = 45$ (because the impact effect of the structural shocks depends on a minimum of 3 reduced-form responses that are zero).

for $q \in \mathcal{Q}$. This requires the selection of a grid \mathcal{Q} and the evaluation of the critical-value function $c^{\alpha_1}(q)$.

Generating a grid for q . We generate $n_{\mathcal{Q}}$ grid points for $q \in \mathbb{S}^n$ from a distribution that is uniform under rotations using a well-known result by James (1954). Let $Z^{(j)}$, $j = 1, \dots, n_{\mathcal{Q}}$, be a sequence of $n \times 1$ vectors of i.i.d. $N(0, I_n)$ random vectors and define $q^{(j)} = Z^{(j)} / \|Z^{(j)}\|$. Then, $q^{(j)}$ is uniformly distributed on the unit sphere \mathbb{S}^n . We define the grid as $\mathcal{Q} = \{q^{(1)}, \dots, q^{(n_{\mathcal{Q}})}\}$. For the confidence intervals to be asymptotically valid, the number of grid points has to expand faster than the sample size. Our theoretical analysis in Section 4 abstracts from the discretization of \mathbb{S}^n .

Weight matrix for sample objective function. The weight matrix $\hat{W}_q(q)$ is obtained as follows. We denote the asymptotic covariance matrix of $\sqrt{T}S(q)(\hat{\phi}_q - \phi_q)$ as $\Sigma(q) = S(q)\Lambda_{qq}S(q)'$. A consistent estimator is given by

$$\hat{\Sigma}(q) = S(q)\hat{\Lambda}_{qq}S'(q) = \hat{D}^{1/2}(q)\hat{\Omega}(q)\hat{D}^{1/2}(q),$$

where $\hat{\Omega}(q)$ is the correlation matrix associated with $\hat{\Sigma}(q)$ and $\hat{D}^{1/2}(q)$ is a diagonal matrix of standard deviations. We then let

$$\hat{W}(q) = T\hat{D}^{-1/2}(q)\hat{B}(q)\hat{D}^{-1/2}(q) \quad (23)$$

and focus on two particular choices of $\hat{B}(q)$: $\hat{B}(q) = \hat{\Omega}^{-1}(q)$ and $\hat{B}(q) = I$. The choice of $\hat{B}(q) = \hat{\Omega}^{-1}(q)$ clearly requires our assumption in (22) that $\Lambda_{qq} > 0$. In the case of $\hat{B}(q) = I$, one could in principle allow for a singular covariance matrix $\Lambda_{qq} > 0$. In the formal analysis in Section 4, one would have to replace Λ_{qq} by its singular value decomposition. We did not pursue this extension below because it would make the notation and exposition more cumbersome.

Overall, this leads to the sample objective function

$$G(q; \hat{\phi}_q, \hat{W}(\cdot)) = \min_{\mu \geq 0} T \left\| \hat{D}^{-1/2}(q)S(q)\hat{\phi}_q - \hat{D}^{-1/2}(q)V(q)\mu \right\|_{\hat{B}(q)}^2.$$

The function $G(q; \hat{\phi}_q, \hat{W}(\cdot))$ has the same structure as the objective functions considered in the literature on moment inequality models, for example, Chernozhukov, Hong, and Tamer (2007), Rosen (2008), Andrews and Guggenberger (2009), and Andrews and Soares (2010). The main difference in our setup for the VAR application is the dimension of $S(q)\hat{\phi}_q$, $\hat{D}(q)$, and $\hat{B}(q)$ varies with q and the limit $\hat{D}(q)$ as a function of q can be singular.

Critical values. In order to obtain the critical value function $c^{\alpha_1}(q)$, we apply the moment selection approach of Andrews and Soares (2010). The moment selection tries to eliminate clearly nonbinding inequality conditions in the weak limit of the objective function $G(q; \hat{\phi}_q, \hat{W}(\cdot))$ and compute the required critical value.⁸ An estimate of the

⁸One can show that in population at most $n - 1$ inequality conditions (recall that n is the dimension of y_t) can be binding; see also the example in Section 2.4. Thus, we experimented with an algorithm that orders the inequalities based on the strength of their violation to select a subset of $n - 1$ binding conditions in case the Andrews–Soares procedure classifies more than $n - 1$ inequalities as binding. In finite samples, the se-

slackness in inequality condition $j = 1, \dots, r(q)$ is provided by

$$\hat{\xi}_{j,T}(q) = \hat{D}_{jj}^{-1/2}(q)[S(q)]_{(j)}\sqrt{T}\hat{\phi}_q. \quad (24)$$

Inequality condition j is deemed nonbinding if

$$\hat{\xi}_{j,T}(q) \geq \kappa_T, \quad (25)$$

where κ_T is a sequence that diverges *slowly* to infinity, for example, $\kappa_T = 1.96 \ln(\ln T)$. Thus, estimates of the number of nonbinding and binding moment inequality constraints are given by

$$\hat{r}_2(q) = \sum_{j=1}^{r(q)} \mathcal{I}\{\hat{\xi}_{j,T}(q) \geq \kappa_T\} \quad \text{and} \quad \hat{r}_1(q) = r(q) - \hat{r}_2(q), \quad (26)$$

respectively.

Define the $(\hat{r}_1(q) \times r(q))$ selection matrix $M_{\hat{\xi}}(q)$ that deletes rows of $\hat{D}^{-1/2}(q)S(q)\hat{\phi}_q$ that correspond to nonbinding inequality conditions in the sense of (25). Moreover, let m be the dimension of the vector ϕ_q ,

$$Z_m \sim N(0, I_m) \quad \text{and} \quad \hat{A}'(q) = \hat{D}^{-1/2}(q)S(q)\hat{L}_{qq}, \quad \text{where} \quad \hat{A}_{qq} = \hat{L}_{qq}\hat{L}'_{qq}.$$

Conditional on $\hat{B}(q)$ and $M_{\hat{\xi}}(q)$, define the random function (the randomness is induced through Z_m)

$$\bar{\mathcal{G}}_{Z_m}(q; \hat{B}(q), M_{\hat{\xi}}(q)) = \min_{\nu \geq 0} \|M_{\hat{\xi}}(q)\hat{A}'(q)Z_m - \nu\|_{M_{\hat{\xi}}(q)\hat{B}(q)M_{\hat{\xi}}'(q)}, \quad (27)$$

where ν is a $\hat{r}_1(q) \times 1$ vector. We adopt the convention that $\bar{\mathcal{G}}_{Z_m}(q; \hat{B}(q), M_{\hat{\xi}}(q)) = 0$ if $\hat{r}_1(q) = 0$. The critical value $c^{\alpha_1}(q)$ is defined as

$$c^{\alpha_1}(q) = 1 - \alpha_1 \text{ quantile of } \bar{\mathcal{G}}_{Z_m}(q; \hat{B}(q), M_{\hat{\xi}}(q)) \quad (28)$$

and can be obtained from a simulation approximation of the limit objective function.⁹ If $\hat{B}(q) = I$, then the evaluation of $\bar{\mathcal{G}}_{Z_m}(\cdot)$ is fast because the quadratic programming problem has the following closed-form solution:

$$\bar{\mathcal{G}}_{Z_m}(q; \hat{B}(q), M_{\hat{\xi}}(q)) = \sum_{j=1}^{\hat{r}_1(q)} [M_{\hat{\xi}}(q)\hat{A}'(q)Z_m]_j^2 \mathcal{I}\{[M_{\hat{\xi}}(q)\hat{A}'(q)Z_m]_j < 0\}.$$

lection of no more than $n - 1$ restrictions could potentially sharpen the confidence set for q . However, in our experiments the gains (if any) were so small that they were essentially not noticeable when we constructed the confidence bands for θ . Thus, we did not explore this idea further in this paper.

⁹For $j = 1, \dots, n_Z$ generate random vectors $Z_m^{(j)}$ and compute $\bar{\mathcal{G}}_{Z_m}^{(j)}(q; \hat{B}(q), M_{\hat{\xi}}(q))$. Then compute the $1 - \alpha_1$ percentile of the empirical distribution of the simulated limit objective functions.

3.3 Confidence set for θ conditional on q

Conditional on q , the dynamic effects of the shock $\varepsilon_{1,t}$ on y_t are point-identified and the inference about impulse responses and variance decompositions is regular. Methods on how to construct confidence intervals for these objects date back to [Runkle \(1987\)](#), who proposed to use either the δ -method in combination with numerical derivatives of the mapping from reduced-form VAR coefficients into IRFs and variance decompositions or to use a residual-based bootstrap. [Lütkepohl \(1990\)](#) derived asymptotic distributions based on analytical derivatives for the δ -method and [Mittnik and Zdrozny \(1993\)](#) provided extensions to VARMA models. A recent survey of the literature on frequentist inference for IRFs and variance decompositions in point-identified settings is provided by [Kilian and Lütkepohl \(2017\)](#). Any of these methods can be embedded into our Bonferoni approach.

3.4 Special case: Confidence bands for IRFs

Confidence bands for impulse responses in the VAR literature predominantly depict pointwise confidence intervals, which means that we can express the scalar parameter θ as $\tilde{S}_\theta(q)\phi_\theta$, where ϕ_θ summarizes the reduced-form impulse responses that are necessary to generate the structural response θ and $\tilde{S}_\theta(q)$ is defined similarly as $\tilde{S}(q)$ in Section 2.3. For this important special case, one can show that $F^\theta(\phi_q, \phi_\theta)$ is convex and bounded, which simplifies computations and reporting of results.

LEMMA 1. *Suppose that $F^q(\phi_q)$ is nonempty and not a singleton. Moreover, $\theta = S_\theta(q)\phi_\theta$ and $k = \dim(\theta) = 1$. Then $F^\theta(\phi_q, \phi_\theta)$ is convex and bounded.*

Approximating $F^q(\hat{\phi}_q)$ and $F^\theta(\hat{\phi})$. An estimate of the identified set for q can be obtained from

$$F^q(\hat{\phi}_q) \approx \hat{F}^q(\hat{\phi}_q) = \{q \in \mathcal{Q} \mid \tilde{S}(q)\hat{\phi}_q \geq 0\}.$$

Thus, for every $q \in \mathcal{Q}$ one checks whether $\tilde{S}(q)\hat{\phi}_q \geq 0$ and retains the q 's for which the condition is satisfied. Denote the elements of $\hat{F}^q(\hat{\phi}_q)$ by $q^{(j)}$, $j = 1, \dots, n_q$, where $n_q \leq n_{\mathcal{Q}}$. Compute $\theta^{(j)} = \hat{\phi}'_\theta q^{(j)}$. We show in the online Appendix that $F^\theta(\hat{\phi})$ is a bounded interval. Thus, we define the interval

$$\hat{F}^\theta(\hat{\phi}) = \left[\left(\min_{j=1, \dots, n_q} \theta^{(j)} \right), \left(\max_{j=1, \dots, n_q} \theta^{(j)} \right) \right].$$

Computing $CS^q(\hat{\phi}_q)$ and $CS^\theta(\hat{\phi})$. The computation of $CS^q(\hat{\phi}_q)$ follows the steps outlined in Section 3.2. Note that by construction $\hat{F}^q(\hat{\phi}_q) \subseteq CS^q(\hat{\phi}_q)$. Denote the elements of $CS^q(\hat{\phi}_q)$ by $q^{(j)}$, $j = 1, \dots, n_q$. Then, for each $q^{(j)}$, compute the Wald interval with bounds

$$\theta_l^{(j)} = \hat{\phi}'_\theta q^{(j)} - z_{\alpha_2/2} \sqrt{q^{(j)'} \hat{\Lambda}_{\theta\theta} q^{(j)} / T} \quad \text{and} \quad \theta_u^{(j)} = \hat{\phi}'_\theta q^{(j)} + z_{\alpha_2/2} \sqrt{q^{(j)'} \hat{\Lambda}_{\theta\theta} q^{(j)} / T},$$

and let

$$CS^\theta(\hat{\phi}) = \Theta \cap \left[\left(\min_{j=1, \dots, n_q} \theta_l^{(j)} \right), \left(\max_{j=1, \dots, n_q} \theta_u^{(j)} \right) \right].$$

The intersection with Θ can be used to restrict the confidence interval to values of θ that are consistent with the assumed sign restriction. While $CS^q(\hat{\phi}_q)$ has to be calculated only once, the computations for $CS^\theta(\hat{\phi})$ have to be repeated for every response $\theta = \partial y_{i,t+h} / \partial \varepsilon_{1,t}$ of interest. Here, i potentially ranges from $i = 1, \dots, n$ and $h = 0, 1, \dots, h_{\max}$.

4. LARGE SAMPLE ANALYSIS

This section formally establishes the consistency of the plug-in estimators $F^q(\hat{\phi}_q)$ and $F^\theta(\hat{\phi})$ and the asymptotic validity of the confidence sets $CS^q(\hat{\phi}_q)$ and $CS^\theta(\hat{\phi})$. As mentioned in Section 2.4, the vectors ϕ_q and ϕ are not always sufficient to characterize the sampling distribution of data and estimators. Thus, we again will use ρ to characterize the distribution of the data under the reduced-form VAR model (1). We denote this distribution by P_ρ . The statements about uniform asymptotic coverage probabilities will be made for $\rho \in \mathcal{R}$. Some of the regularity conditions will be required to hold for a slightly larger, δ -inflated open set

$$\mathcal{R}^\delta = \{\tilde{\rho} \in \bar{\mathcal{R}} \mid \exists \rho \in \mathcal{R} \text{ s.t. } \|\tilde{\rho} - \rho\| < \delta\}, \quad (29)$$

where $\bar{\mathcal{R}} \supset \mathcal{R}$ and $\delta > 0$.¹⁰ Asymptotic inference for q is discussed in Section 4.1 and Section 4.2 considers inference for θ .

4.1 Asymptotic inference for q

We begin by stating some high-level assumptions.

ASSUMPTION 1. *There exists a compact reduced-form parameter set \mathcal{R} and a δ -inflated superset \mathcal{R}^δ defined in (29) such that $\mathcal{R} \subset \mathcal{R}^\delta \subset \bar{\mathcal{R}}$ and:*

- (i) *For every $\rho \in \mathcal{R}^\delta$, there does not exist an $r \times 1$ vector $\lambda > 0$ such that*

$$\Phi_q \lambda = 0.$$

- (ii) *$\phi_q(\rho)$ is continuously differentiable for all $\rho \in \mathcal{R}^\delta$.*

(iii) *There exists an estimator $\hat{\phi}_q$ of $\phi_q(\rho_T)$ and a matrix $\Lambda_{qq}^{-1/2}(\rho_T)$ such that for each sequence $\{\rho_T\} \in \mathcal{R}$ (a) $\hat{\phi}_q - \phi_q(\rho_T) \xrightarrow{P} 0$; (b) $\sqrt{T} \Lambda_{qq}^{-1/2}(\rho_T)(\hat{\phi}_q - \phi_q(\rho_T)) \implies N(0, I)$.*

(iv) *For each $\rho \in \mathcal{R}$, the matrix $\Lambda_{qq}(\rho)$ is continuous, positive definite, and there exists a full-rank positive-definite matrix Λ_{\min} such that $\Lambda_{qq}(\rho) - \Lambda_{\min} \geq 0$ for all $\rho \in \mathcal{R}$.*

(v) *There exists an estimator $\hat{\Lambda}_{qq}$ of $\Lambda_{qq}(\rho_T)$ such that $\|\hat{\Lambda}_{qq} - \Lambda_{qq}(\rho_T)\| \xrightarrow{P} 0$ for any converging sequence $\{\rho_T\} \in \mathcal{R}$.*

Condition (i) of Assumption 1 states that the convex cone generated by the columns of the reduced-form impulse response matrix Φ_q does not contain the zero vector. This

¹⁰For instance, suppose ρ is an autocorrelation parameter for an AR(1) model. We could define $\mathcal{R} = [0, 0.999]$, $\mathcal{R}^\delta = [0, 1)$ for $\delta = 0.001$, and $\bar{\mathcal{R}} = [0, 1]$.

assumption is sufficient to ensure that the identified set $F^q(\phi_q(\rho))$ is nonempty and that the plug-in estimator $F^q(\hat{\phi}_q)$ is consistent whenever $\hat{\phi}_q \xrightarrow{p} \phi_q$ (see Theorem 1 below). Assumption 1(i) rules out, for instance, that equality conditions are coded as pairs of inequalities, and, more generally, that linear combinations of inequalities constrain impulse responses to be equal to zero. We discuss in Section 5 how our framework can be extended to allow for a mixture of inequality and equality restrictions on impulse responses.

Condition (i) is typically not satisfied for all values of the reduced-form parameter $\rho \in \bar{\mathcal{R}}$, which is why we only require it to hold on the set $\mathcal{R}^\delta \subset \bar{\mathcal{R}}$. For instance, consider a VAR(1) generalization of the bivariate VAR(0) in Section 2.2 with autoregressive coefficient matrix A_1 . As before, suppose y_t is composed of inflation and output growth and the investigator imposes the sign restriction that in response to a (positive) demand shock inflation and output responses are both nonnegative upon impact and one period after impact. In this case,

$$\Phi'_q = \begin{bmatrix} \Sigma_{\text{tr}} \\ A_1 \Sigma_{\text{tr}} \end{bmatrix}.$$

If $A_1 = \text{diag}(\rho_1, \rho_2)$ and $\rho_1, \rho_2 < 0$, then Condition (i) is violated. Conditional on these reduced-form parameters, the identified set is empty. Assumption 1 excludes these values of ρ from \mathcal{R}^δ . From a practitioner's perspective, an empty confidence set $CS^q(\hat{\phi}_q)$ provides evidence that the imposed sign restrictions are inconsistent with the estimated reduced-form parameters.

The continuity in Condition (ii) is with respect to the Euclidean norm. While ρ could in principle be infinite-dimensional if the distribution of the error terms is treated non-parametrically, the function $\phi_q(\cdot)$ only depends on the finite-dimensional subvector of ρ that contains the reduced-form parameters $A_1, \dots, A_p, \Sigma_u$; see Equation (4). In combination with the compactness of \mathcal{R} , Condition (ii) implies that the domain of ϕ_q , which is given by $\{\phi_q(\rho) : \rho \in \mathcal{R}\}$, is compact. Conditions (iii) and (v) require that $\hat{\phi}_q$ and $\hat{\Lambda}_{qq}$ converge uniformly for $\rho \in \mathcal{R}$. Note that the stated convergences in probability and in distribution are assumed to hold under the sequence of distributions P_{ρ_T} .¹¹ The uniform convergence of $\hat{\phi}_q$ to a Gaussian limit distribution also requires a restriction of the domain of ρ because it breaks down at the boundary of the stationary region in the VAR parameter space. For instance, in the context of an AR(1) model $y_t = \rho_T y_{t-1} + u_t$ with autoregressive coefficient $\rho_T = 1 - c/T$, an estimator of an impulse response at horizon $h = 1$, that is, $\phi_q(\rho) = \rho$, behaves according to

$$\sqrt{T}(1 - \rho_T^2)^{-1/2}(\hat{\phi}_T - \rho_T) = \frac{\frac{1}{T} \sum y_{t-1} u_t}{\sqrt{c(2 - c/T)} \frac{1}{T^2} \sum y_{t-1}^2} \not\Rightarrow N(0, 1).$$

¹¹For example, $\|\hat{\phi}_q - \phi_q(\rho_T)\| \xrightarrow{p} 0$ is shorthand for $P_{\rho_T}\{\|\hat{\phi}_q - \phi_q(\rho_T)\| > \varepsilon\} \rightarrow 0$ as $T \rightarrow \infty$ for any $\varepsilon > 0$.

Uniform convergence to a Gaussian limit distribution can be achieved if \mathcal{R} is restricted to the interval $[-1 + \varepsilon, 1 - \varepsilon]$ for some $\varepsilon > 0$.¹² From a practitioner's perspective, we are essentially assuming that the researcher has applied some stationarity-inducing transformations, for example, transformed prices into inflation rates. Because some authors, for example, Uhlig (2005), prefer to specify VARs in terms of variables that exhibit (near) nonstationary dynamics, our Monte Carlo experiments in Section 6 include designs in which the roots of the vector autoregressive lag polynomial are close to the unit circle.¹³

Our first theorem establishes that the identified set $F^q(\phi_q)$ is nonempty and not a singleton, that is, the dynamic effects of $\varepsilon_{1,t}$ are set-identified instead of point-identified. This result can be deduced from Assumption 1(i) using Gordan's alternative theorem (see, for instance, Border (2007)).

THEOREM 1. *Suppose Assumption 1(i) is satisfied. Then the identified set $F^q(\phi_q(\rho))$ is nonempty and is not a singleton for all $\rho \in \mathcal{R}^\delta$.*

The second theorem focuses on asymptotic inference. The first part establishes the consistency of the plug-in estimator $F^q(\hat{\phi}_q)$. The consistency is stated in terms of the Hausdorff distance. We denote the Hausdorff distance between two sets A and B by $d_H(A, B)$.¹⁴ The consistency relies on the compactness of $F^q(\phi_q)$ and the continuity of the correspondences with respect to ϕ_q . Unlike in some of the models studied by Chernozhukov, Hong, and Tamer (2007), it is not necessary to inflate the set $F^q(\hat{\phi}_q)$ by $\varepsilon_T \downarrow 0$ to achieve consistency.¹⁵ The second part of Theorem 2 establishes the asymptotic validity of the confidence set $CS^q(\hat{\phi}_q)$. A formal proof of the Theorem is provided in the online Appendix. The proof of the second part closely follows the proof of Theorem 1 in Andrews and Soares (2010). However, a number of nontrivial modifications are required to account for the potential rank reduction of $\tilde{S}(q)$ as a function of q .

THEOREM 2. *Suppose that Assumption 1 is satisfied.*

- (i) *Then $d_H(F^q(\hat{\phi}_q), F^q(\phi_q)) \xrightarrow{P} 0$.*
- (ii) *If $0 < \alpha < 1/2$, then the confidence set $CS^q(\hat{\phi}_q)$, defined in (21), is an asymptotically valid confidence set for q :*

$$\liminf_{T \rightarrow \infty} \inf_{\rho \in \mathcal{R}} \inf_{q \in F^q(\phi_q(\rho))} P_\rho \{q \in CS^q(\hat{\phi}_q)\} \geq 1 - \alpha.$$

¹²See Giraitis and Phillips (2004) for a more general discussion.

¹³An extension of our analysis to VARs with unit roots or cointegration restrictions is beyond the scope of this paper. The construction of uniformly valid confidence intervals for reduced-form parameters in itself is a very challenging task; see Mikusheva (2007).

¹⁴Formally, the Hausdorff distance is defined as $d(A, B) = \max\{d(A | B), d(B | A)\}$, where $d(A | B) = \sup_{a \in A} d(a, B)$ and $d(a, B) = \inf_{b \in B} \|a - b\|$. We set $d(A, B) = \infty$ if either A or B is empty.

¹⁵A result similar to ours in a general GMM setting is provided by Yildiz (2012). We prove the result directly based on Assumption 1.

4.2 Asymptotic inference for θ

As discussed in Section 3.3, we do not provide any new results on confidence intervals for impulse responses or variance decompositions *conditional on* the vector q . For these intervals, we rely on the existing literature. We use $CS_q^\theta(\hat{\phi}_\theta)$ to denote a confidence set for θ conditional on q . The following assumption is required for asymptotic inference about the parameter θ .

ASSUMPTION 2.

- (i) *The function $\theta = f(\Phi_\theta, q)$ is continuous in both its arguments.*
- (ii) *The set $CS_q^\theta(\hat{\phi}_\theta)$ satisfies*

$$\liminf_T \inf_{\rho \in \mathcal{R}} \inf_{(\theta, q) \in F^{\theta, q}(\phi(\rho))} P_\rho \{ \theta \in CS_q^\theta(\hat{\phi}_\theta) \} \geq 1 - \alpha_2,$$

where $F^{\theta, q}(\phi) = \{ \theta \in \Theta, q \in \mathbb{S}_n \mid q \in F^q(\phi_q), \theta = f(\Phi_\theta, q) \}$.

The first condition of Assumption 2(i) is quite weak and the two leading examples of θ in Section 2.1 satisfy this condition. The second condition of Assumption 2(ii) is a high-level condition that is needed for the asymptotic validity of the Bonferroni confidence set of θ . The condition requires that the pointwise confidence set $CS_q^\theta(\hat{\phi}_\theta)$ in q is uniformly valid. In the leading examples of θ , an impulse response of the form $\theta = \Phi'_\theta q$, the conditional confidence set CS_q^θ in (19) satisfies the uniformity condition because the asymptotic normality of $\hat{\phi}_\theta$ in (22) holds uniformly in ϕ . Combining the results of Theorem 2(ii) with Assumption 2 leads to the following theorem.

THEOREM 3. *Suppose that Assumption 1 is satisfied.*

- (i) *If $\hat{\phi}_\theta \xrightarrow{P} \phi_\theta$, then $d_H(F^\theta(\hat{\phi}), F^\theta(\phi)) \xrightarrow{P} 0$, where $\phi = [\phi'_q, \phi'_\theta]'$.*
- (ii) *Suppose that $0 < \alpha < 1/2$ and Assumption 2 is satisfied. Then the confidence set $CS^\theta(\hat{\phi})$, defined in (20), is an asymptotically valid confidence set for θ :*

$$\liminf_{T \rightarrow \infty} \inf_{\rho \in \mathcal{R}} \inf_{\theta \in F^\theta(\phi(\rho))} P_\rho \{ \theta \in CS^\theta(\hat{\phi}) \} \geq 1 - \alpha.$$

5. EXTENSIONS

We now discuss three extensions to the construction of $CS^q(\hat{\phi}_q)$: (i) models that use both sign restrictions and zero restrictions to identify structural impulse responses, (ii) the identification of multiple shocks, (iii) and the use of bootstrapped critical values instead of simulated asymptotic critical values.

Sign restrictions combined with equality restrictions. Assumption 1(i) rules out that opposing sign restrictions are used to represent equality restrictions on impulse responses. Nonetheless, it is straightforward to sharpen the identified set by combining sign restrictions with more traditional exclusion restrictions. In some applications, the restriction that certain responses are zero on impact (zero restrictions) can be translated

into a domain restriction for q that does not depend on any other reduced-form parameters. For instance, in the empirical analysis in Section 7.2 we will replace an unrestricted 4×1 vector $q \in \mathbb{S}^n$ by the restricted vector $q = [0_{1 \times 2}, q_2']'$, where q_2 is a 2×1 vector with $\|q_2\| = 1$. In this case, the previously developed methods can be applied without any modification.

If the equality restrictions imposed on the impulse responses lead to restrictions on q that depend on some of the reduced-form parameters, then they can be accommodated by generalizing the objective function $G(q; \phi, \tilde{W})$ in (15) as follows. Define

$$\tilde{G}(q; \phi, \tilde{W}) = \min_{\mu \geq 0} \left\| \begin{pmatrix} S_{eq}(q)\phi \\ S(q)\phi - V(q)\mu \end{pmatrix} \right\|_{\tilde{W}(q)}^2, \quad (30)$$

where $S_{eq}(q)\phi$ corresponds to the responses that are restricted to be zero. Following the arguments in Andrews and Soares (2010), it is straightforward albeit tedious to extend the proof of Theorem 2 to a mixture of equality and inequality conditions.¹⁶ The extension closely resembles the proof of Theorem 2(i) in the working paper version Moon, Schorfheide, and Granziera (2013) for the projection-based confidence set, which also involves a mix of equality and inequality conditions. From a practitioners perspective, the only other modification that is required, is to replace the limit objective function $\tilde{G}(\cdot)$ in (27) that is used to simulate the critical value $c^{\alpha_1}(q)$ by the limit expression of $\tilde{G}(q; \phi, \tilde{W})$ in (30).

Identifying multiple shocks. Some authors use sign-restricted SVARs to identify multiple shocks simultaneously. For instance, Peersman (2005) considers an $n = 4$ dimensional VAR, composed of oil price inflation, output growth, consumer price inflation, and nominal interest rates. He uses sign restrictions to identify an oil price shock, aggregate demand, and supply shocks, and a monetary policy shock. To identify n shocks, the unit vector q has to be replaced by an orthogonal matrix, and the restrictions will take the form

$$\tilde{S}(\Omega)\phi_q \geq 0$$

for a suitably defined function $\tilde{S}(\Omega)$. While all our results easily generalize to multiple shocks (just replace q by Ω), the implementation becomes computationally more difficult because the grid for the $n - 1$ dimensional vector q has to be replaced by a grid for orthogonal matrix Ω , which has $n(n - 1)/2$ degrees of freedom.

Bootstrapped critical values instead of asymptotic critical values. Our simulated critical values rely on the Gaussian limit distribution of $\sqrt{T}\hat{D}^{-1/2}(q)S(q)(\hat{\phi}_q - \phi_q)$, which is reflected in the vector $\hat{A}'(q)Z_m$ in the random function $\tilde{G}(\cdot)$ in (27). Alternatively, the critical values could be constructed by replacing draws from $\hat{A}'(q)Z_m$ with draws from the bootstrap approximation of $\sqrt{T}\hat{D}^{-1/2}(q)S(q)(\hat{\phi}_q - \phi_q)$. Bootstrap procedures for VAR impulse response functions are discussed, for instance, in Kilian (1998) and Kilian and Lütkepohl (2017).

¹⁶If we denote the matrix of zero-restricted orthogonalized responses by $\Phi_{q,eq}$, then the generalization of Assumption 1(i) is: there do not exist vectors $\lambda > 0$ and $\lambda_{eq} \geq 0$ such that $\Phi_q\lambda + \Phi_{q,eq}\lambda_{eq} = 0$. The generalized analysis would use Motzkin's transposition theorem; see Border (2007).

TABLE 1. Steps of Monte Carlo experiments.

1.	Generate a sample of size T from the data-generating process.
2.	Compute $\hat{\phi}_q$, $\hat{\phi}_\theta$, and the bounds of $F^\theta(\hat{\phi}_q, \hat{\phi}_\theta)$.
3.	Compute $\hat{\Lambda}_{qq}$ and $\hat{\Lambda}_{\theta\theta}$ using a parametric bootstrap approach.
4.	Compute the $1 - \alpha_1$ confidence set CS^q .
5.	For each definition of θ , compute the $1 - \alpha_2$ confidence sets CS^θ .
6.	For each definition of θ , compute the $1 - (\alpha_1 + \alpha_2)$ confidence sets CS^θ .

6. MONTE CARLO ILLUSTRATIONS

In this section, we conduct three Monte Carlo experiments to illustrate the properties of our proposed confidence sets. In these experiments, θ is a scalar impulse response. During preliminary computations, we noticed that the results for $\hat{B}(q) = I$ and $\hat{B}(q) = \hat{\Omega}^{-1}(q)$ were very similar. Thus, we decided to subsequently report results for $\hat{B}(q) = I$ because in this case the critical values can be computed much faster. We will drop the $\hat{\phi}$ arguments from the confidence sets and report coverage probabilities and average lengths for CS^q and CS^θ . Each Monte Carlo experiment involves the steps summarized in Table 1, which are repeated $n_{\text{sim}} = 5000$ times.

The three experiments differ with respect to the data generating process (DGP). Experiment 1 (Section 6.1) is based on the bivariate VAR(0) model in Section 2.2. Experiment 2 (Section 6.2) features a bivariate VAR(1). The simulation designs for the Experiments 1 and 2 are obtained by fitting a VAR(0) to data on U.S. inflation and GDP growth and fitting first-order VARs to inflation and either output growth or linearly detrended log GDP. Finally, Experiment 3 (Section 6.3) mimics the four-variable VAR(2) fitted to U.S. data on output, inflation, interest rates, and money balances in the empirical analysis of Section 7.

6.1 Experiment 1

Design. The parameterization of the DGP $y_t \sim \text{i.i.d. } N(0, \Sigma_u)$ is provided in Table 2 in the column labeled *Design 1*. We define θ as the response of $y_{1,t}$ to $\varepsilon_{1,t}$. Because $\Sigma_{21}^{\text{tr}} < 0$ in our design, the geometry of the Monte Carlo design corresponds to the left panel of Figure 1. Thus, the upper bounds (in polar coordinates) of F^q and CS^q are $\pi/2$ and the lower bounds of F^θ and CS^θ are zero, respectively. The identified set for θ is $F^\theta(\phi_{q,0}, \phi_{\theta,0}) = [0, 0.578]$. Below, we report coverage probabilities for the lower bound of F^q and the upper bound of F^θ because they are the least favorable parameter values in the respective identified sets. We consider sample sizes of $T = 100$ and $T = 500$. The grid \mathcal{Q} for q is obtained as follows: q is transformed into polar coordinates $[\cos(\varphi), \sin(\varphi)]'$ and we choose $n_{\mathcal{Q}} = 315$ equally spaced grid points for φ on the interval $(-\pi/2, \pi/2]$. The number of bootstrap repetitions to obtain $\hat{\Lambda}_{qq}$ and $\hat{\Lambda}_{\theta\theta}$ is $n_\Lambda = 1000$ and the number of simulations to obtain the critical value and $c^{\alpha_1}(q)$ is $n_Z = 500$. Further details on the implementation are provided in the online Appendix.

Results. Detailed results for the frequentist confidence intervals are summarized in Table 3. Recall that the nominal coverage probability for θ is 90%. For $T = 100$ the actual

TABLE 2. Monte Carlo design.

	Experiment 1		Experiment 2	
	Design 1 VAR(0)	Design 2 VAR(1)	Design 3 VAR(1)	Design 4 VAR(1)
Σ_{11}^{tr}	0.597	0.295	0.283	0.210
Σ_{21}^{tr}	-0.205	-0.092	-0.081	-0.043
Σ_{22}^{tr}	0.812	0.795	0.817	0.542
$A_{1,11}$		0.873	0.806	0.450
$A_{1,12}$		0.003	0.032	0.014
$A_{1,21}$		-0.229	-0.278	0.060
$A_{1,22}$		0.230	0.985	0.953
$\lambda_1(A_1)$		0.871	0.89 - 0.03i	0.955
$\lambda_2(A_1)$		0.231	0.89 + 0.03i	0.498

Note: Designs are obtained by estimating a VAR(0) or VAR(1) of the form $y_t = A_0 + A_1 y_{t-1} + u_t$, $\mathbb{E}[u_t u_t'] = \Sigma_{\text{tr}} \Sigma_{\text{tr}}'$ using OLS. $\lambda_i(A_1)$ is the i th eigenvalue of A_1 . $y_{1,t}$ is the log difference of the U.S. GDP deflator, scaled by 100 to convert into percentages. $y_{2,t}$ is either the log difference of U.S. GDP or deviations of log GDP from a linear trend, scaled by 100. Design 1: inflation and GDP growth, 1964:I to 2006:IV. Design 2: inflation and output deviations from trend, 1964:I to 2006:IV. Design 3: inflation and output growth, 1964:I to 2006:IV. Design 4: inflation and output deviations from trend, 1983:I to 2006:IV.

TABLE 3. Experiments 1 and 2: single-horizon sign restrictions.

	Experiment 1		Experiment 2					
	Design 1		Design 2		Design 3		Design 4	
	Coverage	Length	Coverage	Length	Coverage	Length	Coverage	Length
$F^q(\phi_q)$		$\frac{42}{100} \pi$		$\frac{36}{100} \pi$		$\frac{47}{100} \pi$		$\frac{51}{100} \pi$
$F^\theta(\phi)$		0.579		0.233		0.226		0.094
Sample Size $T = 100$								
CS^q	0.938	$\frac{47}{100} \pi$	0.936	$\frac{81}{100} \pi$	0.932	$\frac{57}{100} \pi$	0.940	$\frac{67}{100} \pi$
CS^θ	0.980	0.671	0.979	0.295	0.934	0.265	0.942	0.128
CS^{ϕ_q}	0.879		0.865		0.865		0.871	
Sample Size $T = 500$								
CS^q	0.930	$\frac{44}{100} \pi$	0.936	$\frac{44}{100} \pi$	0.932	$\frac{51}{100} \pi$	0.936	$\frac{56}{100} \pi$
CS^θ	0.990	0.622	0.991	0.265	0.963	0.244	0.958	0.110
CS^{ϕ_q}	0.909		0.894		0.901		0.904	

Note: *Length* refers to the average length of the confidence intervals across Monte Carlo repetitions. For $F^q(\phi_q)$ and CS^q , we report the arc length; see Figure 1. We let $\alpha_1 = \alpha_2 = 0.05$, which implies that the nominal coverage probabilities are 95% for CS^q and 90% for CS^θ and CS^{ϕ_q} . The confidence interval for ϕ_q has a nominal coverage probability of 90%.

coverage probability for the Bonferroni sets is 0.98. As we increase the sample size to $T = 500$, the length of the confidence intervals shrinks, while the actual coverage probabilities increases to 0.99. It is instructive to also examine the coverage probabilities of CS^q and the Wald confidence set for $\phi_q = \text{vech}(\Sigma_{\text{tr}})$, which we denote by CS^{ϕ_q} . The coverage probability for the reduced-form parameter vector ϕ_q is 88% for $T = 100$ and

approaches its nominal value of 90% as the sample size is increased to $T = 500$. This increase in coverage probability for ϕ_q mirrors the increase in coverage probability for θ . The Bonferroni intervals are computed based on $\alpha_1 = \alpha_2 = 0.05$, which implies that the nominal coverage probability of CS^q is 95%. The actual coverage probabilities for the nuisance parameter vector q are slightly smaller, namely, around 93%. Overall, the Bonferroni-type marginalization generates conservative confidence intervals for θ .

6.2 Experiment 2

Design. We now add first-order autoregressive terms to the simulation design to introduce persistence in the endogenous variables:

$$y_t = A_1 y_{t-1} + u_t, \quad u_t \sim \text{i.i.d. } N(0, \Sigma_u).$$

The choices for A_1 and Σ_u are summarized in Table 2 under the headings *Design 2*, *Design 3*, and *Design 4*. The designs differ with respect to the persistence of the vector autoregressive process. *Design 2* is the least persistent. The eigenvalues of A_1 are 0.871 and 0.231. *Design 4* is the most persistent with eigenvalues 0.955 and 0.498. We focus on responses at horizon $h = 1$, which can be obtained from $\phi_q = \text{vec}((A_1 \Sigma_{\text{tr}})')$. The structural parameter of interest, θ , is defined as $\partial y_{1,t+1} / \partial \varepsilon_{1,t}$. As in Experiment 1, we compute coverage probabilities for the lower bound of $F^q(\phi_q)$ and the upper bound of $F^\theta(\phi)$. The grid \mathcal{Q} for q is obtained as follows: q is transformed into polar coordinates $[\cos(\varphi), \sin(\varphi)]'$ and we choose $n_{\mathcal{Q}} = 629$ equally spaced grid points for φ on the interval $(-\pi, \pi]$. The remaining aspects of the design are the same as in Experiment 1.

Sign restrictions over a single horizon. We impose the sign restrictions that $\partial y_{1,t+1} / \partial \varepsilon_{1,t}$ and $\partial y_{2,t+1} / \partial \varepsilon_{1,t}$ are nonnegative:

$$\phi_{q,1} q_1 + \phi_{q,2} q_2 \geq 0 \quad \text{and} \quad \phi_{q,3} q_1 + \phi_{q,4} q_2 \geq 0.$$

For now, we do not impose sign restrictions on the responses at impact or at horizons greater than $h = 1$. The geometry of the identified sets $F^q(\phi_q)$ and $F^\theta(\phi_q, \phi_\theta)$ and its projections is similar to the geometry depicted in Figure 1. The main difference is that the second boundary of the identified set is given by the solution of $q_1 \phi_{q,3} + q_2 \phi_{q,4} = 0$ and $\|q\| = 1$, instead of $q = [0, 1]'$. Overall, the results for Experiment 2 reported in Table 3 are qualitatively similar to those for *Design 1*. The actual coverage probabilities of the confidence sets for q are around 0.94 and, therefore, close to the nominal coverage probability of $1 - \alpha_1 = 0.95$. The θ sets, on the other hand, are conservative. Their coverage probabilities range from 0.934 to 0.991, thereby exceeding the nominal level of 0.9.

Sign restrictions over multiple horizons. As before, we define θ as the contemporaneous impact of the shock on $y_{1,t}$: $\theta = \partial y_{1,t} / \partial \varepsilon_{1,t}$. However, we now restrict the signs of the impulse responses $\partial y_{i,t+h} / \partial \varepsilon_{1,t} \geq 0$ for both variables $i = 1, 2$ over multiple periods: $h = 0, 1, \dots, H$. This increases the number of inequality conditions. Monte Carlo results are presented in Table 4. The effect of adding sign restrictions differs across the three designs. In *Design 2*, the lengths of the identified sets $F^q(\cdot)$ and $F^\theta(\cdot)$ shrink drastically: from 0.35π and 0.265 for $H = 1$ to 0.006π and 0.007 , respectively, for $H = 4$. Under *Design 4*, the sizes of the two identified sets remain constant as H is increased from 1 to 4.

TABLE 4. Experiment 2: multiple-horizon sign restrictions, sample size $T = 100$.

	Design 2			Design 3			Design 4		
	Coverage	Length	Bind. Ineq.	Coverage	Length	Bind. Ineq.	Coverage	Length	Bind. Ineq.
Restrictions: $h = 0, 1$									
$F^q(\phi_q)$		$\frac{35}{100}\pi$			$\frac{43}{100}\pi$			$\frac{47}{100}\pi$	
CS^q	0.953	$\frac{48}{100}\pi$	1.29	0.976	$\frac{50}{100}\pi$	1.97	0.953	$\frac{53}{100}\pi$	2.06
$F^\theta(\phi)$		0.265			0.277			0.209	
CS^θ	0.982	0.307		0.977	0.317		0.949	0.236	
Restrictions: $h = 0, \dots, 4$									
$F^q(\phi_q)$		$\frac{6}{1000}\pi$			$\frac{37}{100}\pi$			$\frac{47}{100}\pi$	
CS^q	0.985	$\frac{40}{100}\pi$	7.78	0.989	$\frac{49}{100}\pi$	4.48	0.979	$\frac{56}{100}\pi$	7.83
$F^\theta(\phi)$		0.006			0.261			0.208	
CS^θ	1.000	0.277		0.993	0.315		0.949	0.235	

Note: *Length* either refers to the length of the population identified set or the average length of the confidence intervals across Monte Carlo repetitions. For $F^q(\phi_q)$ and CS^q , we report the arc length; see Figure 1. *Bind. Ineq.* is the average number of inequalities considered binding by the Andrews and Soares (2010) moment selection procedure. We let $\alpha_1 = \alpha_2 = 0.05$, which implies that the nominal coverage probabilities are 95% for CS^q and 90% for CS^θ .

Design 2 is an intermediate case. Restricting impulse responses at multiple horizons essentially adds rays to Figure 1. The location of the new rays relative to the $H = 1$ rays determines whether the identified sets shrink or not.

For *Designs 2* and *3*, the length of the confidence intervals for q and θ are decreasing in the number of inequality restrictions, but they do not shrink as quickly as the length of the identified sets. Simultaneously, the actual coverage probability for q increases with the number of sign restrictions. In *Design 2*, the coverage probability for $H = 1$ is 0.953, which is close to the nominal coverage probability of $1 - \alpha_1 = 0.95$. For $H = 4$, the coverage probability increases to 0.985. While in population for any H only one inequality is binding¹⁷ at the boundary of the identified set for q , the average number of moment conditions deemed binding by the Andrews and Soares (2010) selection rule rises from 1.29 to 7.78 in *Design 2*. Recall that in order to guarantee a uniform asymptotic coverage probability, the selection rule has to classify too many rather than too few moment conditions as binding. This inflates the critical value as well as the coverage probability and makes the q confidence set more conservative. We observe a similar pattern for *Design 3*. The Bonferroni sets for θ are generally conservative across all designs and maximum horizons H . Under *Designs 2* and *3*, the actual coverage probabilities tend to increase as more restrictions are added. Nonetheless, the average length decreases. This decrease is most pronounced for *Design 2*. Here, the length shrinks from 0.307 ($H = 1$) to 0.277 ($H = 4$). Under *Design 4*, the average length of the confidence interval and the coverage probability stay essentially constant as we vary the number of restrictions.

¹⁷An inspection of Figure 1 suggests that if more than one inequality condition is binding then it must be the case that the rays corresponding to the binding inequality conditions are identical.

6.3 Experiment 3

Design. Finally, we consider a four-variable VAR(2) that mimics the model for per capita GDP (in deviations from a linear trend), inflation, the federal funds rate, and real money balances used in the empirical application in Section 7:

$$y_t = c + A_1 y_{t-1} + A_2 y_{t-2} + u_t \quad \text{with } u_t \sim N(0, \Sigma). \quad (31)$$

The reduced-form parameters are set equal to the empirical point estimates (reported in the online Appendix). As in the application below, we consider the following sign restrictions:

$$\frac{\partial y_{2,t+h}}{\partial \varepsilon_{1,t}} \leq 0, \quad \frac{\partial y_{3,t+h}}{\partial \varepsilon_{1,t}} \geq 0, \quad \frac{\partial y_{4,t+h}}{\partial \varepsilon_{M,t}} \leq 0, \quad h = 0, 1.$$

The sample size for the simulated data sets is $T = 170$. We set the number of (randomly generated from a uniform distribution on the hypersphere) grid points for \mathcal{Q} to $n_{\mathcal{Q}} = 20,000$. The number of bootstrap repetitions to obtain $\hat{\Lambda}_{qq}$ and $\hat{\Lambda}_{\theta\theta}$ is $n_{\Lambda} = 1000$ and the number of simulations to obtain the critical value $c^{\alpha_1}(q)$ is $n_Z = 1000$. As in the previous experiments, we focus on $\hat{B} = I$. The implementation of the computations for $CS^{\theta}(I)$ follows the description in Section 3.

Benchmark results. Baseline results for $\alpha_1 = \alpha_2 = 0.05$ are plotted in Figure 2. The top panels depict the upper bounds and the lower bounds for pointwise identified sets and confidence sets for the responses of Variables 1 and Variables 2 at horizons $h = 0, 1, \dots, 23$. The bounds of the confidence sets are averaged across Monte Carlo repetitions. An economic interpretation of the responses will be provided in Section 7. For now, we focus on the widths of the confidence bands relative to the widths of the population identified sets and the coverage probabilities, which are depicted in the bottom panels. The identified sets have a considerable width that leaves the sign of the response of $y_{1,t}$ undetermined. The confidence bands are noticeably wider than the identified sets, which is a reflection of the sampling uncertainty associated with the estimators of the reduced-form parameters. The coverage probabilities in the bottom panels are computed for the upper bounds (dashed) and lower bounds (solid) of the pointwise identified sets. As in Experiments 1 and 2, the actual coverage probability is substantially larger than the nominal coverage probability of 90% (indicated by the solid horizontal line), making the confidence bands conservative.

Adjusting α_1 , keeping $\alpha = 0.1$ fixed. The baseline choice of $\alpha_1 = 0.05$ is arbitrary. Thus, it is worthwhile to explore what happens if we change α_1 , which determines the size of CS^q . Figure 1 provides some intuition for the potential outcomes of this experiment. In the left panel of the figure (labeled $\Sigma_{21}^{\text{tr}} < 0$), the upper bound for the identified set $F^{\theta}(\cdot)$ is determined by the lower bound of $F^q(\cdot)$. Thus, increasing α_1 and thereby decreasing the size of $F^q(\cdot)$ can potentially sharpen the confidence set for θ , provided that the decrease in $F^q(\cdot)$ exceeds the increase in the conditional confidence set CS_q^{θ} . Alternatively, if $\Sigma_{21}^{\text{tr}} > 0$ (depicted in the right panel of Figure 1), the upper bound of $F^{\theta}(\cdot)$ is determined by a value of q that lies strictly in the interior of $F^q(\cdot)$. Thus, in order to shrink the confidence set for θ one should lower α_1 and raise α_2 so that CS_q^{θ} shrinks.




FIGURE 2. Impulse responses bands and coverage probabilities. *Notes:* Top panels: population identified set $F^\theta(\phi)$ (lines with crosses), averaged upper (dashed) and lower (solid) bounds of pointwise confidence sets, and a zero line (thin solid). Bottom panels: actual coverage probabilities for 90% confidence sets at the lower (solid) and upper (dashed) bounds of the pointwise identified sets. The thin solid horizontal line indicates the nominal coverage probability. $\alpha_1 = \alpha_2 = 0.05$

Figure 3 depicts the coverage probabilities and the average interval width for three levels of α_1 : $\alpha_1 = 0.05$ which we used to generate the baseline results in Figure 2, $\alpha_1 = 0.01$, and $\alpha_1 = 0.09$. The figure also depicts the width of the population identified set. The differences between the widths of the confidence intervals and the identified set can be interpreted as the excess lengths of the confidence intervals. It turns out that in this particular Monte Carlo design it is advantageous to reduce α_1 . Setting $\alpha_1 = 0.01$ reduces the coverage probability of the intervals and shrinks the width of the intervals. However, the effect is modest at best. Relative to the overall width of the identified sets and the baseline confidence bands, the reduction is very small. The actual coverage probability remains above 95% except for the long horizon responses of $y_{2,t}$ which fall slightly below the nominal level of $\alpha = 0.9$ for $h \geq 18$. As we saw in Table 3, the confidence set for the reduced-form VAR parameters can have an actual coverage probability that is less than its nominal coverage probability, which in turn tightens the confidence interval for θ .

Adjusting α , keeping α_1 fixed. Several authors have devised methods to overcome the conservativeness of Bonferroni confidence intervals by raising the nominal level α

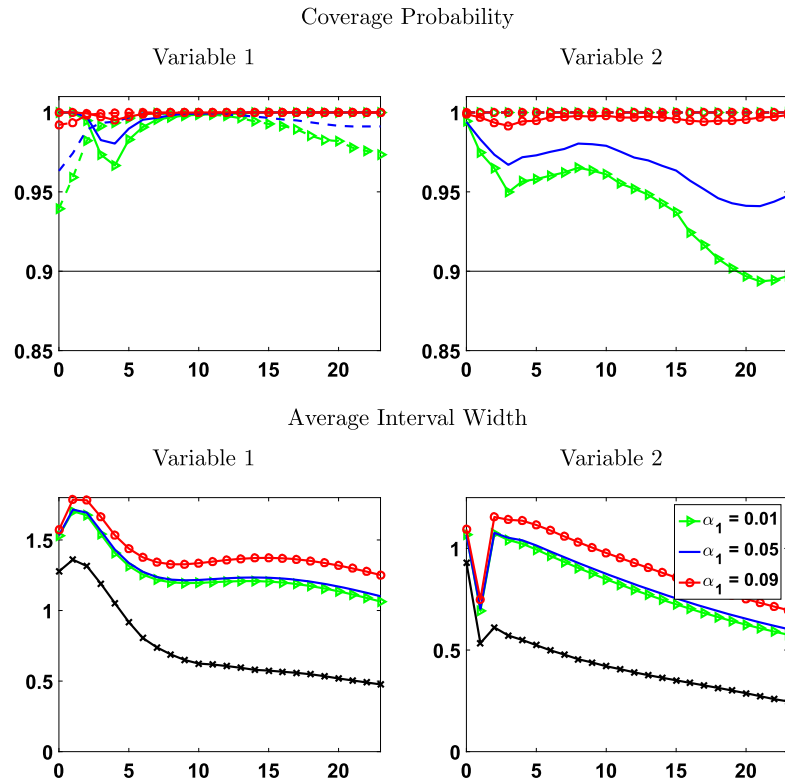


FIGURE 3. The effect of varying α_1 , fixed $\alpha = 0.10$. *Notes:* Top panels: Thin horizontal line indicates the nominal coverage probability of 90%, other lines represent actual coverage probabilities. Dashed lines refer to upper bounds and solid lines to lower bounds of pointwise identified sets; $\alpha_1 = 0.01$ is marked by triangles, $\alpha_1 = 0.05$ (baseline) has no line symbols, and $\alpha_1 = 0.09$ is marked by circles. Bottom panels: average width of confidence bands (triangles, no line symbol, circles) and width of population identified set (crosses).

to target a desired actual coverage probability $1 - \alpha_* = 0.90$, say. Examples of this approach are [Campbell and Yogo \(2006\)](#) and, most recently, [McCloskey \(2017\)](#). The former paper reduces the size of the first-stage confidence interval by raising α_1 , keeping the size of the second-stage intervals, CS_q^θ in our notation constant. The latter paper proposes to reduce the size of the second-stage intervals, keeping α_1 constant. In view of the results depicted in Figure 3, we informally follow [McCloskey's \(2017\)](#) approach by increasing α_2 from 0.05 (baseline) to 0.10 and 0.15, respectively. While, in principle, we could choose a different α_2 for each variable and each horizon, the results depicted in Figure 4 are generated using the same α_2 for each response. As expected, the actual coverage probability falls as we increase α_2 (and thereby $\alpha = \alpha_1 + \alpha_2$). For the $y_{1,t}$ responses the coverage probabilities remain above 90%, while for the long-horizon responses of $y_{2,t}$ the coverage probability drops substantially below 90% for $h \geq 15$ and $\alpha_2 = 0.10$. The attainable reduction in the width of the confidence band is larger than in the case of fixed α , but it remains small relative to the overall width of the bands.

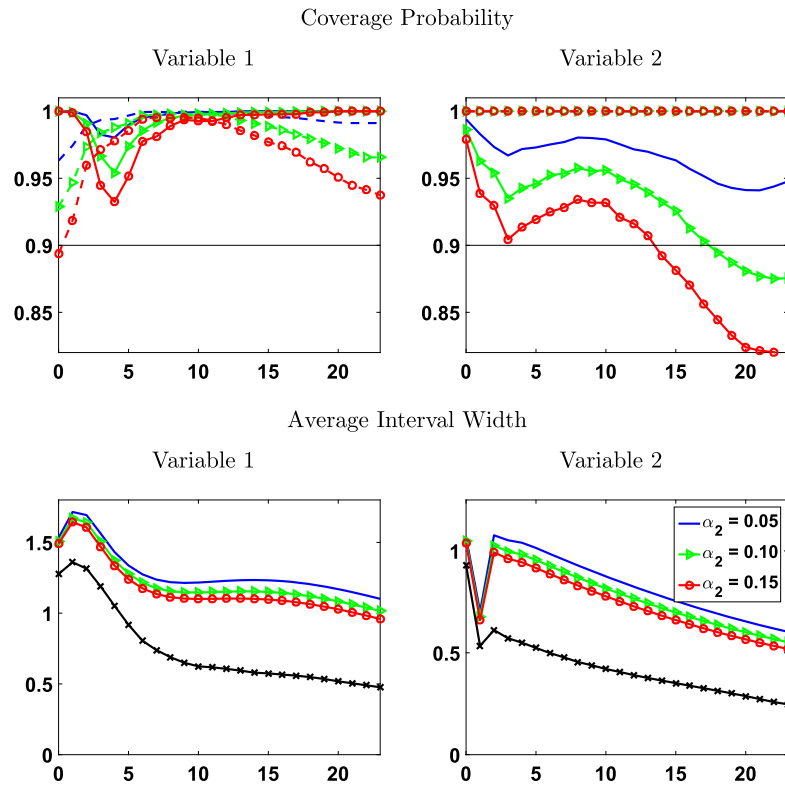


FIGURE 4. The effect of varying α_2 , fixed $\alpha_1 = 0.05$. *Notes:* Top panels: Thin horizontal line indicates the nominal coverage probability of 90%, other lines represent actual coverage probabilities. Dashed lines refer to upper bounds and solid lines to lower bounds of pointwise identified sets; $\alpha_2 = 0.05$ (baseline) has no line symbols, $\alpha_2 = 0.10$ is marked by triangles, and $\alpha_2 = 0.15$ is marked by circles. Bottom panels: average width of confidence bands (no line symbol, triangles, circles) and width of population identified set (crosses).

7. EMPIRICAL ILLUSTRATION

We now apply the previously developed methods to a four-variable VAR. The vector of observables consists of per capita real GDP (in deviations from a linear trend), inflation, the federal funds rate, and real money balances. We use quarterly U.S. data from 1965:I to 2006:IV, which excludes the 2007–09 recession and the subsequent period of zero nominal interest rates. A detailed description of the dataset is provided in the online Appendix. All VARs are estimated with $p = 2$ lags, which is the preferred lag length according to BIC. We will consider two set-identification schemes for monetary policy shocks. The first scheme involves only sign restrictions (Section 7.1), whereas the second identification is based on a combination of equality and sign restrictions (Section 7.2). As in Monte Carlo Experiment 3, we set $n_A = 1000$, $n_Q = 20,000$, and $n_Z = 1000$.

In addition to computing Bonferroni confidence bands, we also generate pointwise Bayesian credible intervals for the impulse responses, which have been widely used in empirical research. The Bayesian credible sets reported subsequently are based on the

VAR(p) given in (1) with Gaussian innovations $u_t \sim \text{i.i.d. } N(0, \Sigma_u)$. Let $A = [A_1, \dots, A_p]'$ and define the unnormalized vector \tilde{q} such that $q = \tilde{q}/\|\tilde{q}\|$. If $\tilde{q} \sim N(0, I_n)$, then q is uniformly distributed on the hypersphere. Following Uhlig (2005), we use an improper prior of the form

$$p(A, \Sigma, \tilde{q}) \propto |\Sigma|^{-(n+1)/2} \exp\{-\tilde{q}'\tilde{q}/2\} \mathcal{I}\left\{\frac{\tilde{q}}{\|\tilde{q}\|} \in F^q(\phi(A, \Sigma))\right\}. \quad (32)$$

We use the acceptance sampler described in Uhlig (2005) to generate 50,000 draws from the posterior distribution of (A, Σ, \tilde{q}) . These draws are then converted into impulse responses, and credible sets are computed from the impulse response draws.

7.1 Pure sign restrictions

In order to conduct inference about the effects of a contractionary monetary policy shock, we use the following sign restrictions to bound the identified set: in periods $h = 0, 1$ (i) the interest rate response is weakly positive; (ii) the inflation response is weakly negative; and (iii) real money balances do not rise above their steady-state level. These sign restrictions were also used in Monte Carlo Experiment 3 in Section 6.3.

Figure 5 depicts three bands: (pointwise) 90% Bonferroni confidence intervals (using a diagonal weight matrix) $CS^\theta(I)$, estimated sets $F^\theta(\hat{\phi})$, and (pointwise) 90% Bayesian




FIGURE 5. Pure sign restrictions over horizons $h = 0, 1$. Notes: The figure depicts 90% Bonferroni confidence bands $CS^\theta(I)$ (no line symbols) with $\alpha_1 = \alpha_2 = 0.05$; 90% Bayesian credible bands (circles); and the estimated sets $F^\theta(\hat{\phi})$ (crosses).

credible sets. The two most notable features of the bands are that the frequentist confidence bands (solid) are substantially wider than the Bayesian credible bands (short dashes) and that the Bayesian credible bands approximately coincide with the estimated set $F^\theta(\hat{\phi})$. As explained in detail in Moon and Schorfheide (2012), in a large sample, that is, a sample in which uncertainty about ϕ is small compared to the size of $F^\theta(\hat{\phi})$, the Bayesian intervals lie inside the estimated sets $F^\theta(\hat{\phi})$ because in the limit essentially all of the probability mass is concentrated on $F^\theta(\hat{\phi})$ and a 90% credible interval is always a subset of the support of the posterior distribution. The frequentist interval, on the other hand, has to extend beyond the boundaries of $F^\theta(\hat{\phi})$ because it has to have, say, 90% coverage probability for every element of the identified set $F^\theta(\phi)$, including the boundary points. From a substantive perspective, the use of sign restrictions leaves the direction of the output response undetermined.

The top panels of Figure 6 show output and inflation responses obtained by requiring the the sign-restrictions hold for periods $h = 0, 1, \dots, 8$, keeping $\alpha_1 = \alpha_2 = 0.05$. This modification increases the number of inequality restrictions from 6 to 27. As the num-

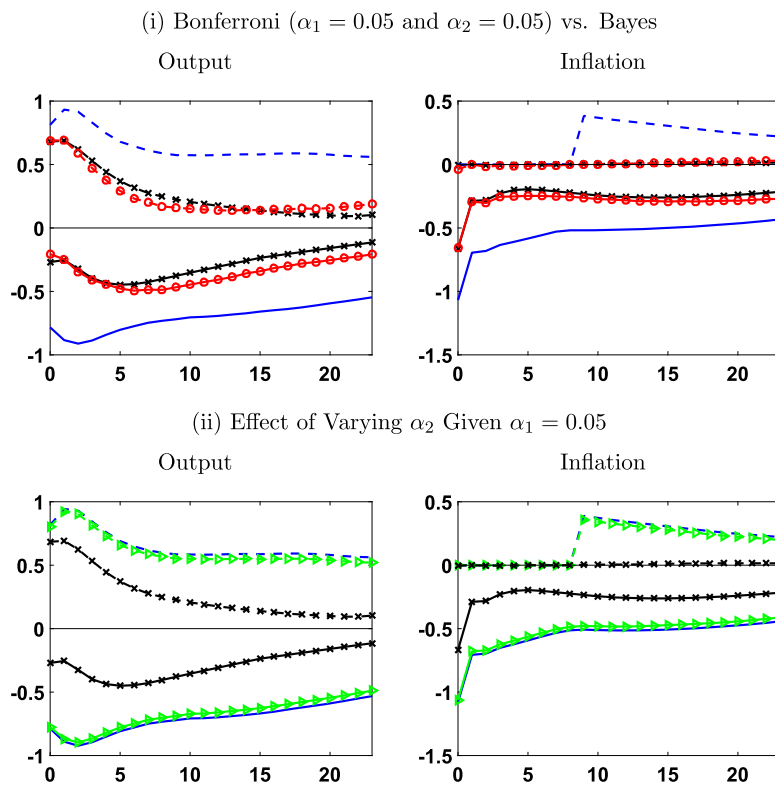


FIGURE 6. Pure sign restrictions over horizon $h = 0, 1, \dots, 8$. Notes: Top panels: 90% Bonferroni confidence bands $CS^\theta(I)$ (no line symbols); 90% Bayesian credible bands (circles); and the estimated sets $F^\theta(\hat{\phi}_q, \hat{\phi}_\theta)$ (crosses). Bottom panels: Bonferroni confidence bands $CS^\theta(I)$ (no line symbols) with $\alpha_2 = 0.05$; Bonferroni confidence bands $CS^\theta(I)$ (triangles) $\alpha_2 = 0.10$; and the estimated sets $F^\theta(\hat{\phi})$ (crosses).

ber of sign restrictions increases, the width of the identified sets decreases. As suggested by the Monte Carlo simulations, the width of the Bonferroni bands also decreases. The bottom panels of Figure 6 show the effect of raising α_2 from 0.05 to 0.10. According to the simulations in Section 6.3, this decrease in the nominal coverage probability brings the actual coverage probability closer to the desired coverage probability of 90%. As a result the width of the confidence bands shrinks, but not by much. In fact, in percentage terms, the width reduction is very small.

7.2 Combining sign restrictions and zero restrictions

A commonly used identification assumption for monetary policy shocks is that private-sector variables such as output and inflation cannot respond to changes in the federal funds rate within the period; see, for instance, [Christiano, Eichenbaum, and Evans \(1999\)](#). Because the initial impact of the monetary policy shock is given by $\Sigma_{\text{tr}}q$ and we ordered the elements of y_t such that output and inflation appear before interest rates and real money balances, the identification condition implies that the first two elements of the vector q have to equal zero. Thus, we can reduce the dimension of the vector q as follows: $q = [0, 0, \cos \varphi, \sin \varphi]'$, where $\varphi \in [0, 2\pi]$. This is more efficient than adding two equality conditions to the set of inequality conditions; see Section 5. The zero restriction on the instantaneous inflation response replaces the sign restriction used in Section 7.1. We maintain the other sign restrictions used previously, that is, the interest rate responses for $h = 0$ and $h = 1$ are weakly positive and the inflation response in period $h = 1$ as well as the real money balance responses in periods $h = 0$ and $h = 1$ are weakly negative.

Impulse response bands are depicted in Figure 7. A comparison of $F^\theta(\hat{\phi})$ in Figures 5 and 7 indicates that the use of zero restrictions reduces the size of the identified set drastically. For instance, if the zero restrictions are imposed, the inflation response is essentially point identified for horizons exceeding 8 quarters. As a consequence, for output as well as medium- and long-run inflation responses, the width of the frequentist and Bayesian coverage bands is now much more similar than under the pure-sign-restriction




FIGURE 7. Combining zero and sign restrictions over horizons $h = 0, 1$. *Notes:* The figure depicts 90% Bonferroni confidence bands $CS^\theta(I)$ (no line symbols) with $\alpha_1 = \alpha_2 = 0.05$; 90% Bayesian credible bands (circles); and the estimated sets $F^\theta(\hat{\phi})$ (crosses).

scenario. However, some differences remain with respect to the short-run inflation response. For the first two years, the frequentist intervals cover both positive and negative inflation responses, whereas the Bayesian credible intervals suggest that the inflation response is negative. With the zero restrictions imposed, the direction of the output response is no longer ambiguous—it is negative over the first two years.

8. CONCLUSION

With the exception of FRSW, the coverage bands for impulse responses of sign-restricted SVARs that have been reported in the literature thus far were only meaningful from a Bayesian perspective. The main contribution of our paper is to develop an easy-to-use frequentist method to construct Bonferroni confidence intervals for measures of the dynamic effects of structural shocks, such as impulse responses, in VARs that are set-identified based on sign restrictions. In the first stage, a confidence set for the vector of weights q on the reduced-form impulse responses is obtained by inverting a point-wise hypothesis test for the moment inequalities implied by the sign restrictions. We employ the Andrews and Soares (2010) moment selection procedure to obtain critical values for this test that are not diluted by nonbinding inequality conditions.

Our empirical application illustrates that in set-identified VARs, frequentist confidence bands can be substantially wider than Bayesian credible bands. As a by-product, we establish the consistency of the plug-in estimator $F^\theta(\hat{\phi})$ of the identified set of impulse responses. $F^\theta(\hat{\phi})$ is also useful from a Bayesian perspective. Because in a Bayesian analysis the prior distribution of the impulse response functions conditional on the reduced-form parameters is not updated, it is useful to report the identified set and the prior conditional on some estimate of ϕ , say, the posterior mean, so that the audience can judge whether the conditional prior distribution is highly concentrated in a particular area of the identified set.

REFERENCES

- Andrews, D. W. K. (2017), “Identification-robust subvector inference.” Cowles Foundation Discussion Papers, 3005. [1089]
- Andrews, D. W. K. and P. J. Barwick (2012), “Inference for parameters defined by moment inequalities: A recommended moment selection procedure.” *Econometrica*, 80, 2805–2826. [1090]
- Andrews, D. W. K. and P. Guggenberger (2009), “Validity of subsampling and ‘plug-in asymptotics’ inference for parameters defined by moment inequalities.” *Econometric Theory*, 25, 669–709. [1090, 1099]
- Andrews, D. W. K. and G. Soares (2010), “Inference for parameters defined by moment inequalities using generalized moment selection.” *Econometrica*, 78, 119–157. [1090, 1097, 1099, 1104, 1106, 1110, 1118]
- Baumeister, C. and J. Hamilton (2015), “Sign restrictions, structural vector autoregressions, and useful prior information.” *Econometrica*, 83 (5), 1963–1999. [1088]

- Baumeister, C. and G. Peersman (2013), “Time-varying effects of oil supply shocks on the U.S. economy.” *American Economic Journal: Macroeconomics*, 5 (4), 1–28. [1088]
- Border, K. C. (2007), “Alternative linear inequalities.” Manuscript, California Institute of Technology. [1104, 1106]
- Bugni, F. A., I. A. Canay, and X. Shi (2017), “Inference for subvectors and other functions of partially identified parameters in moment inequality models.” *Quantitative Economics*, 8 (1), 1–38. [1089]
- Campbell, J. Y. and M. Yogo (2006), “Efficient tests of stock return predictability.” *Journal of Financial Economics*, 81, 27–60. [1089, 1113]
- Canay, I. A. and A. M. Shaikh (2017), “Practical and theoretical advances in inference for partially-identified models.” In *Advances in Economics and Econometrics—Eleventh World Congress. Vol. II. Econom. Soc. Monogr.* (B. Honore, A. Pakes, M. Piazzesi, and L. Samuelson, eds.), 271–306, Cambridge University Press, Cambridge. [1090]
- Canova, F. and G. De Nicrolo (2002), “Monetary disturbances matter for business cycle fluctuations in the G-7.” *Journal of Monetary Economics*, 49, 1131–1159. [1088]
- Cavanagh, C. L., G. Elliott, and J. H. Stock (1995), “Inference in models with nearly integrated regressors.” *Econometric Theory*, 11, 1131–1147. [1089]
- Chaudhuri, S. and E. Zivot (2011), “A new method of projection-based inference in GMM with weakly identified nuisance parameters.” *Journal of Econometrics*, 164, 239–251. [1089]
- Chernozhukov, V., H. Hong, and E. Tamer (2007), “Estimation and confidence regions for parameter sets in econometric models.” *Econometrica*, 75, 1243–1284. [1090, 1099, 1104]
- Christiano, L., M. Eichenbaum, and C. Evans (1999), “Monetary policy shocks: What have we learned and to what end?” In *Handbook of Macroeconomics*, Vol. 1a (J. Taylor and M. Woodford, eds.), 65–148, North Holland. [1117]
- Dedola, L. and S. Neri (2007), “What does a technology shock do? A VAR analysis with model-based sign restrictions.” *Journal of Monetary Economics*, 54, 512–549. [1088]
- Faust, J. (1998), “The robustness of identified VAR conclusions about money.” *Carnegie-Rochester Conference Series on Public Policy*, 49, 207–244. [1088]
- Faust, J., J. Rogers, E. Swanson, and J. Wright (2003), “Identifying the effects of monetary policy shocks on exchange rates using high frequency data.” *Journal of the European Economic Association*, 1, 1031–1057. [1088]
- Faust, J., E. Swanson, and J. Wright (2004), “Identifying VARs based on high frequency futures data.” *Journal of Monetary Economics*, 51, 1107–1131. [1088]
- Gafarov, B., M. Meier, and J. L. Montiel (2016), “Projection inference for set-identified SVARs.” Manuscript, Columbia University. [1090]

- Gafarov, B., M. Meier, and J. L. Montiel Olea (2018), “Delta-method inference for a class of set-identified SVARs.” *Journal of Econometrics*, 203, 316–327. [1090]
- Gambetti, L. and A. Musso (2017), “Loan supply shocks and business cycle.” *Journal of Applied Econometrics*, 32, 764–782. [1088]
- Giacomini, R. and T. Kitagawa (2015), “Robust inference about partially-identified SVARs.” Manuscript, UCL. [1090]
- Giraitis, L. and P. C. Phillips (2004), “Uniform limit theory for stationary autoregression.” Cowles Foundation Discussion Paper 1475. [1104]
- Granziera, E., H. R. Moon, and F. Schorfheide (2018), “Supplement to ‘Inference for VARs identified with sign restrictions.’” *Quantitative Economics Supplemental Material*, 86, <https://doi.org/10.3982/QE978>. [1087, 1091]
- Hristov, N., O. Hülsewig, and T. T. Wollmershäuser (2012), “Loan supply shocks during the financial crisis: Evidence for the euro area.” *Journal of International Money and Finance*, 31, 569–592. [1088]
- James, A. T. (1954), “Normal multivariate analysis and the orthogonal group.” *Annals of Mathematical Statistics*, 25, 40–75. [1099]
- Kaido, H., F. Molinari, and J. Stoye (2016), “Confidence intervals for projections of partially identified parameters.” Manuscript, Cornell University. [1089]
- Kilian, L. (1998), “Small-sample confidence intervals for impulse response functions.” *Review of Economics and Statistics*, 80, 218–230. [1106]
- Kilian, L. and H. Lütkepohl (2017), *Structural Vector Autoregressive Analysis*. Cambridge University Press. [1088, 1101, 1106]
- Kilian, L. and D. Murphy (2012), “Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models.” *Journal of the European Economic Association*, 10, 1166–1188. [1088]
- Lütkepohl, H. (1990), “Asymptotic distributions of impulse response functions and forecast error variance decompositions of vector autoregressive models.” *Review of Economic and Statistics*, 72 (1), 116–125. [1101]
- Manski, C. F. (2003), *Partial Identification of Probability Distributions*. Springer Verlag, New York. [1088]
- McCloskey, A. (2017), “Bonferroni-based size-correction for nonstandard testing problems.” *Journal of Econometrics*, 200 (1), 17–35. [1089, 1113]
- Mikusheva, A. (2007), “Uniform inference in autoregressive models.” *Econometrica*, 75, 1411–1452. [1104]
- Mittnik, S. and P. A. Zdrozny (1993), “Asymptotic distributions of impulse responses, step responses, and variance decompositions of estimated linear dynamic models.” *Econometrica*, 61 (4), 857–870. [1101]

- Moon, H. R., F. Schorfheide, E. Granziera, and M. Lee (2011), “Inference for VARs identified with sign restrictions.” NBER Working Paper, 17140. [1088, 1089]
- Moon, H. R. and F. Schorfheide (2012), “Bayesian and frequentist inference in partially identified models.” *Econometrica*, 80 (2), 755–782. [1088, 1116]
- Moon, H. R., F. Schorfheide, and E. Granziera (2013), “Inference for VARs identified with sign restrictions.” Manuscript, University of Pennsylvania. [1106]
- Mountford, A. and H. Uhlig (2009), “What are the effects of fiscal policy shocks?” *Journal of Applied Econometrics*, 24, 960–992. [1088]
- Pappa, E. (2009), “The effects of fiscal shocks on employment and the real wage.” *International Economic Review*, 50, 217–244. [1088]
- Peersman, G. (2005), “What caused the early millennium slowdown? Evidence based on vector autoregressions.” *Journal of Applied Econometrics*, 20, 185–207. [1106]
- Peersman, G. and R. Straub (2009), “Technology shocks and robust sign restrictions in a euro area SVAR.” *International Economic Review*, 50, 727–750. [1088]
- Romano, J. P., A. M. Shaikh, and M. Wolf (2014), “A practical two-step method for testing moment inequalities.” *Econometrica*, 82, 1979–2002. [1090]
- Rosen, A. M. (2008), “Confidence sets for partially identified parameters that satisfy a finite number of moment inequalities.” *Journal of Econometrics*, 146, 107–117. [1090, 1099]
- Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010), “Structural vector autoregressions: Theory of identification and algorithms for inference.” *Review of Economic Studies*, 77, 665–696. [1088, 1090]
- Runkle, D. E. (1987), “Vector autoregressions and reality.” *Journal of Business & Economic Statistics*, 5 (4), 437–442. [1101]
- Sims, C. A. (1980), “Macroeconomics and reality.” *Econometrica*, 48, 1–48. [1087]
- 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. [1088, 1104, 1115]
- Yildiz, N. (2012), “Consistency of plug-in estimators of upper contour and level sets.” *Econometric Theory*, 28, 309–327. [1104]

Co-editor Andres Santos handled this manuscript.

Manuscript received 15 September, 2017; final version accepted 2 February, 2018; available online 16 February, 2018.