

DELTA-METHOD INFERENCE FOR A CLASS OF SET-IDENTIFIED SVARS

BULAT GAFAROV¹, MATTHIAS MEIER² AND JOSÉ LUIS MONTIEL OLEA³

We study vector autoregressions that impose equality and/or inequality restrictions to set-identify the dynamic responses to a single structural shock of interest (e.g., a monetary policy shock). We make three contributions. First, we present an algorithm to compute the largest and smallest value that an impulse-response coefficient can attain over its identified set. Second, we provide conditions under which these largest and smallest values are directionally differentiable functions of the model's reduced-form parameters (i.e., the autoregressive coefficients and the covariance matrix of residuals). Third, we propose a computationally convenient delta-method approach to conduct inference about the coefficients of the structural impulse-response function. We illustrate our results by set-identifying the responses of different U.S. macroeconomic variables to an unconventional monetary policy shock. We use the theory and methods developed in this paper to assess the effects of the announcement of the second part of the Quantitative Easing program in August 2010. (JEL-Classification: C1, C32, E47).

KEYWORDS: Set-Identification, SVAR, Delta Method, Directional Differentiability, Unconventional Monetary Policy.

1. INTRODUCTION

An increasingly popular practice in empirical macroeconomics is to set-identify the parameters of a structural vector autoregression [SVAR] by means of exclusion and/or sign restrictions. Most studies working with this type of models have relied on Bayesian methods to construct posterior credible sets for the structural parameters of interest (for example, Inoue and Kilian (2013), Arias, Rubio-Ramirez, and Waggoner (2014), and Baumeister and Hamilton (2015)).

A practical concern with Bayesian analysis in set-identified SVARs is that posterior inference continues to be influenced by prior beliefs even if the sample size is infinite (Poirier (1998), Gustafson (2009), Moon and Schorfheide (2012)). This observation has motivated the study of alternative inference approaches that dispense with the specification of a prior distribution over structural parameters that are only set-identified.

¹ Penn State University, Department of Economics and National Research University Higher School of Economics. E-mail: bzg134@psu.edu.

² Bonn University, Department of Economics. E-mail: matthias.meier@uni-bonn.de.

³ Columbia University, Department of Economics. E-mail: jm4474@columbia.edu. This draft: July 10, 2017. First draft: November 19th, 2014.

⁴ We would like to thank seminar participants at the econometrics workshop at Bonn, Brown, Cornell, ITAM, Michigan, Ohio State University, Pennsylvania State University, Rutgers, UCSD, Vanderbilt, Wisconsin-Madison, the 2014 NSF-Time Series Conference, and the 2016 summer meetings of the Econometric Society (Europe and North America) for helpful comments on an earlier draft of this paper. We would also like to thank Oliver Linton, Andres Santos, and two anonymous referees for extremely helpful comments and suggestions. Bulat Gafarov gratefully acknowledges support from the Basic Research Program of the National Research University Higher School of Economics. The usual disclaimer applies.

There are now two interesting and innovative suggestions to summarize uncertainty in set-identified SVARs. On the one hand, [Moon, Schorfheide, and Granziera \(2013\)](#) [MSG13] have proposed a *frequentist* confidence interval for structural impulse-response coefficients based on a moment-inequality-minimum-distance framework. On the other hand, [Giacomini and Kitagawa \(2015\)](#) [GK15] have proposed a *robust Bayes* credible interval that achieves a given credibility level regardless of the prior specified over the model’s set-identified structural parameters.

We contribute to the analysis of set-identified SVARs by proposing a novel delta-method interval for the coefficients of the impulse-response function [IRF]. We show that our delta-method interval is *point-wise consistent in level* and, under certain regularity conditions, it also has *asymptotic robust Bayesian credibility* of at least the nominal level. Thus, our inference approach can be interpreted both from a frequentist and a robust Bayes perspective. We also argue that the computational cost of our procedure compares very favorably with MSG13, GK15, and also with the standard Bayesian algorithm described in [Uhlig \(2005\)](#).

Broadly speaking, our approach is based on a closed-form characterization of the endpoints of the identified set and its directional derivative. Our delta-method interval—which may be viewed as a generalization of the pioneering work of [Lütkepohl \(1990\)](#) on delta-method inference for point-identified VARs—takes the form of a plug-in estimator for the identified set plus/minus standard errors.

The main limitation of our approach is that the delta-method interval is only defined for SVAR models that impose equality and inequality restrictions on a single structural shock (e.g., a monetary policy shock). Admittedly, this is problematic, as some popular applications of set-identified SVARs feature restrictions on multiple structural innovations.¹ In spite of this observation, single-shock set-identified models have been applied in several empirical studies: for example, to study the effects of monetary policy on output [[Uhlig \(2005\)](#)], the impact of monetary policy on the housing market [[Vargas-Silva \(2008\)](#)], the effects of labor market shocks on worker flows [[Fujita \(2011\)](#)], the effects of exchange rates on aggregate prices [[An and Wang \(2012\)](#)], and the effect of optimism shocks on business cycles fluctuations [[Beaudry, Nam, and Wang \(2014\)](#)]. Thus, we think there is room for our results to have an impact on empirical work.

To illustrate the usefulness of our main results, we estimate a monetary structural vector autoregression using monthly U.S. data from July 1979 to December 2007 (a sample that deliberately ends one semester before the financial crisis begins). The goal of our exercise is to use pre-crisis data to learn about the responses of macroeconomic variables to shocks that have effects similar to the ‘unconventional’ monetary policy interventions implemented

¹SVAR applications for the oil market set-identify both demand and supply shocks using sign restrictions and elasticity bounds [[Kilian and Murphy \(2012\)](#)]. The same is true for recent labor market applications [[Baumeister and Hamilton \(2015\)](#)]. Also [Mountford and Uhlig \(2009\)](#)—one of the most cited applications of set-identified SVARs—use sign restrictions to identify a government revenue shock as well as a government spending shock, while controlling for a generic business cycle shock and a monetary policy shock.

after the crisis.

We set-identify an *unconventional* monetary policy [UMP] shock as an innovation that decreases the two-year government bond rate upon impact, but has no effect over the nominal federal funds rate.² We consider two additional sign restrictions on the contemporaneous responses of inflation and output. Namely, we assume that—upon impact—neither inflation nor output can respond negatively to a UMP shock. Since the model is only set-identified, our analysis effectively captures the effects of any historical economic shock that affected the economy in the same way as an UMP shock.

We apply our delta-method approach to construct a confidence interval for the dynamic responses of industrial production, inflation, the two-year government bond rate, and the nominal federal funds rate. We use our delta-method intervals to assess the effects of the announcement of the second part of the so-called Quantitative Easing program (QE2) in August 2010. Pre-crisis data turns out to be extremely useful to learn about the post-crisis response of macroeconomic aggregates to unconventional monetary policy.

The remainder of the paper is organized as follows. Section 2 presents an overview of the main methodological results in this paper. Section 3 introduces our empirical application, which is used as a running example throughout the paper. Section 4.1 presents our algorithm to evaluate the endpoints of the identified set. Section 4.2 establishes the differentiability properties of the endpoints. Section 4.3 presents our delta-method approach and establishes its asymptotic frequentist validity as well as its asymptotic robust Bayesian credibility. Section 5 presents the delta-method intervals for the dynamic responses to the QE2 program. Section 6 concludes. All of our proofs are collected in Appendix A. Additional figures and implementation details of different procedures are collected in Appendix B.

GENERIC NOTATION: If A is a matrix, A_{ij} denotes the ij -th element of A , $\text{vec}(A)$ denotes the vectorization of A , and $\text{vech}(A)$ denotes half-vectorization (applicable only if A is symmetric). The Kronecker product between matrices A and B is denoted by $A \otimes B$. The vector $e_i^m \in \mathbb{R}^m$ denotes the i -th column of the identity matrix—denoted \mathbb{I}_m —of dimension m . If B is a matrix of dimension $n \times n$, $B_i \equiv Be_i^n$ denotes its i -th column. If the dimension of e_i^n is obvious, we ignore the superscript n .

²The paper focuses on the two-year rate as this variable changed considerably after the announcement of the second round of the Quantitative Easing program. See [Krishnamurthy and Vissing-Jorgensen \(2011\)](#)

2. MODEL, SET-IDENTIFYING RESTRICTIONS, AND OVERVIEW OF MAIN THEORETICAL RESULTS

This section presents the baseline SVAR model, discusses the class of set-identifying restrictions that we consider, and provides an overview of our main methodological results.

2.1. SVAR model and impulse-response coefficients

We study the n -dimensional structural vector autoregression (SVAR) with p lags; i.i.d. structural innovations distributed according to F ; and unknown $n \times n$ structural matrix B :

$$(2.1) \quad Y_t = A_1 Y_{t-1} + \dots + A_p Y_{t-p} + B \varepsilon_t, \quad \mathbb{E}_F[\varepsilon_t] = 0_{n \times 1}, \quad \mathbb{E}_F[\varepsilon_t \varepsilon_t'] \equiv \mathbb{I}_n.$$

The object of interest is the k -th period ahead structural impulse response function of variable i to a particular shock j (e.g., a monetary policy shock):

$$(2.2) \quad \lambda_{k,i,j}(A, B) \equiv e_i' C_k(A) B_j,$$

where $B_j \equiv B e_j$ and e_i and e_j denote the i -th and j -th column of \mathbb{I}_n .³ We refer to the parameter in (2.2) as the (k, i, j) -coefficient of the structural impulse-response function.

An auxiliary object in the estimation of (2.2) is the vector of *reduced-form VAR parameters*:

$$(2.3) \quad \mu \equiv (\text{vec}(A)', \text{vec}(\Sigma)')' \in \mathcal{M} \subseteq \mathbb{R}^d, \quad \text{where } A \equiv (A_1, A_2, \dots, A_p), \quad \Sigma \equiv B B'.$$

The reduced-form parameter space is denoted as \mathcal{M} . The parameter A denotes the autoregressive coefficients of the VAR model, while Σ denotes the covariance matrix of residuals. These parameters can be estimated directly from the data by multivariate Least-Squares (LS). Our main high-level assumption will be the approximate normality of the distribution of the LS estimator of μ . This condition will be satisfied even in the presence of unit roots and possible cointegration of unknown form; see [Sims, Stock, and Watson \(1990\)](#), [Toda and Yamamoto \(1995\)](#), [Dolado and Lütkepohl \(1996\)](#), [Inoue and Kilian \(2002\)](#), and Proposition 7.1 in [Lütkepohl \(2007\)](#). Our main assumption is less demanding than the asymptotic normality of the reduced-form impulse-responses in MSG13 (see [Kilian \(1998\)](#), [Benkowitz, Neumann, and Lütkepohl \(2000\)](#)).⁴

³The transformation $C_k(A)$ that appears in equation (2.2) is defined recursively by the formula $C_0 \equiv \mathbb{I}_n$:

$$C_k(A) \equiv \sum_{m=1}^k C_{k-m}(A) A_m, \quad k \in \mathbb{N},$$

$A_m = 0$ if $m > p$; see [Lütkepohl \(1990\)](#), p. 116.

⁴We would like to thank an anonymous referee for suggesting this clarification.

2.2. *Set-Identifying Restrictions*

A common practice in empirical macroeconomics is to use equality and inequality restrictions to *set-identify* the structural IRFs in (2.2). An example of an equality restriction in a monetary VAR is that prices do not react contemporaneously to monetary policy shocks. An example of an inequality restriction is that a contractionary monetary policy shock cannot increase prices.

Let $\mathcal{R}(\mu) \subseteq \mathbb{R}^n$ be the set of values of B_j that satisfy the inequality and equality restrictions. In our paper, the set $\mathcal{R}(\mu)$ takes the form:

$$(2.4) \quad \mathcal{R}(\mu) \equiv \left\{ B_j \in \mathbb{R}^n \mid Z(\mu)'B_j = \mathbf{0}_{m_z \times 1} \text{ and } S(\mu)'B_j \geq \mathbf{0}_{m_s \times 1} \right\},$$

where $Z(\mu)$ is a matrix of dimension $n \times m_z$ and $S(\mu)$ is a matrix of dimension $n \times m_s$. The matrix $Z(\mu)$ collects the equality restrictions specified by the researcher (we assume there are m_z of them). The matrix $S(\mu)$ collects the inequality restrictions (we assume there are m_s of them).

The simple formulation in (2.4) allows the researcher to incorporate the following identifying restrictions:

- a) Sign restrictions on the responses of variable i at horizon k to an impulse on the j -th shock:

$$e_i' C_k(A) B_j \geq \text{ or } = 0,$$

as in Uhlig (2005).

- b) Long-run restrictions on the response of variable i to an impulse on the j -th shock:

$$e_i' (\mathbb{I}_n - A_1 - \dots - A_p)^{-1} B_j \geq \text{ or } = 0,$$

as in Blanchard and Quah (1989).

- c) Short-run restrictions on the coefficients of the j -th structural equation. For example, the contemporaneous coefficient of the i -th variable in the j -th structural equation:

$$e_i' (B')^{-1} e_j = e_i' \Sigma^{-1} B_j \geq \text{ or } = 0,$$

as in Rubio-Ramirez, Caldara, and Arias (2015).

- d) Elasticity bounds as in Kilian and Murphy (2012); for example, for some $b \in \mathbb{R}$:

$$e_i' B_j / e_{i'}' B_j \geq b \iff (e_i - b e_{i'})' B_j \geq 0,$$

provided $e_{i'}' B_j > 0$.

SIGN-NORMALIZATION: In order to make sure that the impulse response of interest is with respect to a fixed-sign shock one should always impose a sign-normalization. Our framework allows at least two different ways of imposing such a normalization: i) restricting the sign of the direct effect of the j -th variable on the j -th equation, or ii) restricting the sign of an arbitrary IRF coefficient. The first type of sign normalization is covered in c) as the short-run restriction $e_j' B^{-1} e_j \geq 0$, while the second is covered in a) as a typical sign restriction on the IRFs.

2.3. Overview of the main results

The main results in this paper concern the ‘endpoints’ of the identified set for a given structural impulse-response coefficient. These endpoints (which we sometimes refer to as the *maximum and minimum* response) are defined as follows:

DEFINITION 1: Given a vector of reduced-form parameters μ we define the endpoints of the identified set for $\lambda_{k,i,j}$ as the functions:

$$(2.5) \quad \bar{v}_{k,i,j}(\mu) \equiv \sup_{B \in \mathbb{R}^{n \times n}} e_i' C_k(A) B e_j, \text{ s.t. } BB' = \Sigma \text{ and } B e_j \in \mathcal{R}(\mu),$$

and

$$(2.6) \quad \underline{v}_{k,i,j}(\mu) \equiv \inf_{B \in \mathbb{R}^{n \times n}} e_i' C_k(A) B e_j, \text{ s.t. } BB' = \Sigma \text{ and } B e_j \in \mathcal{R}(\mu).$$

The function $\bar{v}_{k,i,j}(\mu)$ corresponds to the largest value of the structural parameter over its identified set. The lower bound is defined analogously.

Our delta-method approach is supported by the three results described in the abstract, which can be summarized as follows:

- **THEOREM 1** (*Algorithm to evaluate the maximum and minimum response*): We present an algorithm that allows a researcher to evaluate the endpoints of the identified set given a vector of reduced-form parameters. The algorithm evaluates all different collections of ‘active’ constraints and selects those that generate the largest (or smallest) value function—after checking that the inequality constraints not included in the set of active constraints are satisfied.⁵

Our algorithm does not require sampling from the space of structural matrices B . Instead, we show that $\bar{v}_{k,i,j}(\mu)$ and $\underline{v}_{k,i,j}(\mu)$ are the *value functions* of a mathematical program whose

⁵Given a point x , we refer to any collection of binding restrictions defining $\mathcal{R}(\mu)$ as *active* constraints at x . The term ‘active constraints’ or ‘active set of constraints’ is the common terminology used in numerical optimization; see p. 308 in Nocedal and Wright (2006).

Karush-Kuhn-Tucker points can be described analytically—up to a set of active inequality constraints. More concretely, Lemma 1 shows that the maximum response for $\lambda_{k,i,j}$ is equal to either plus or minus the function:

$$v_{k,i,j}(\mu; r) \equiv \left(e_i' C_k(A) \Sigma^{1/2} M_{\Sigma^{1/2}r} \Sigma^{1/2} C_k(A)' e_i \right)^{1/2},$$

where

$$M_{\Sigma^{1/2}r} \equiv \mathbb{I}_n - \Sigma^{1/2} r (r' \Sigma r)^{-1} r' \Sigma^{1/2},$$

and r is a matrix collecting the gradient vectors of the constraints in $\mathcal{R}(\mu)$ that are active at a maximum. Evaluating the function above for different values of r and checking the feasibility of the corresponding solution yields the maximum response. The minimum response is obtained analogously.

- **THEOREM 2 (Directional Differentiability of the endpoints):** We show that the functions $\bar{v}_{k,i,j}(\cdot)$ and $\underline{v}_{k,i,j}(\cdot)$ are *directionally* differentiable. More precisely, let $X^*(\mu)$ denote set of maximizers of program (2.5). Consider a sequence of ‘perturbations’ of μ each of them in a ‘direction’ $h_N \in \mathbb{R}^d$. We show that for any sequence $h_N \in \mathbb{R}^d$ such that $h_N \rightarrow h \in \mathbb{R}^d$, and any sequence $t_N \rightarrow \infty$:

$$t_N \left(\bar{v}_{k,i,j}(\mu + h_N/t_N) - \bar{v}_{k,i,j}(\mu) \right) \rightarrow \max_{x \in X^*(\mu)} \left[\dot{v}_{k,i,j}(\mu; r(\mu; x))' h \right],$$

where $r(\mu; x)$ collects the gradient of the constraints that are active at a point x . The proof of the result above builds on Lemma 2 which establishes the differentiability of the function $v_{k,i,j}$ for a fixed set of active constraints. We relate the expression of the directional derivative with the generalized versions of the envelope theorems in the work of [Fiacco and Ishizuka \(1990\)](#) and [Bonnans and Shapiro \(2000\)](#). We argue that directional differentiability of the value functions (as opposed to full differentiability) arises due to the possibility that different structural models lead to the maximum (or minimum) response.

- **THEOREM 3 (Large-sample properties):** We establish the point-wise consistency in level and the asymptotic robust Bayes credibility of our delta-method interval. Our suggested interval takes the form

$$CS_T(1 - \alpha; \lambda_{k,i,j}) \equiv \left[\underline{v}_{k,i,j}(\hat{\mu}_T) - z_{1-\alpha/2} \hat{\sigma}_{(k,i,j),T} / \sqrt{T}, \bar{v}_{k,i,j}(\hat{\mu}_T) + z_{1-\alpha/2} \hat{\sigma}_{(k,i,j),T} / \sqrt{T} \right],$$

where $\hat{\mu}_T$ is the typical LS estimator for the VAR reduced-form parameters, $z_{1-\alpha/2}$ is the $(1 - \alpha/2)$ quantile of a standard normal, and $\hat{\sigma}_{(k,i,j),T}$ is our formula for the standard errors based on the directional derivatives.

3. RUNNING EXAMPLE: UNCONVENTIONAL MONETARY POLICY SHOCKS

This section introduces our empirical application, which will be used as a running example to illustrate our assumptions and results.

We consider a simple 4-variable model that includes the Consumer Price Index (CPI_t), the Industrial Production Index (IP_t), the 2-year Treasury Bond rate ($2yTB_t$), and the Federal Funds rate (FF_t).⁶ We take a logarithmic transformation of CPI_t , IP_t and then work with first differences for all variables. Thus, our vector of macro variables is:

$$Y_t \equiv \left(\ln CPI_t - \ln CPI_{t-1}, \quad \ln IP_t - \ln IP_{t-1}, \quad 2yTB_t - 2yTB_{t-1}, \quad FF_t - FF_{t-1} \right)'$$

We set the number of lags equal to $p = 12$ following [Gertler and Karadi \(2015\)](#). The time span of the monthly series is July 1979 to August 2008 ($T = 342$). To keep our exposition as simple as possible, we ignore potential co-integration issues between short-term and long-term interest rates. Without loss of generality, we assume that column of B corresponding to an UMP shock is the first column; $B_1 \equiv Be_1$. Our equality/inequality restrictions are summarized in the following Table I. These sign restrictions can be justified by the DSGE model calibrated in the work of [Bhattarai, Eggertsson, and Gafarov \(2014\)](#).

TABLE I
SET-IDENTIFICATION OF AN UNCONVENTIONAL MONETARY POLICY SHOCK: RESTRICTIONS

Series	Acronym	UMP	Notation
Consumer Price Index	CPI	+	$e'_1 B_1 \geq 0$
Industrial Production	IP	+	$e'_2 B_1 \geq 0$
2-year Treasury Bond rate	2yTB	-	$e'_3 B_1 \leq 0$
Fed Funds Rate	FF	0	$e'_4 B_1 = 0$

DESCRIPTION: Restrictions on contemporaneous responses to a UMP shock. '0' stands for a zero restriction, '-' stands for a negative sign restriction and '+' for positive sign restriction.

[Baumeister and Benati \(2013\)](#) study a related identification scheme. They consider a Bayesian SVAR to study an analogous 'spread' monetary policy shock that leaves the short-term nominal rate unchanged, but affects the spread between the ten-year Treasury-bond yield and the policy rate.

OUTLINE FOR THE REST OF OUR PAPER: We have already presented an overview of our main results and described our running example. In the remaining part of the paper, we formalize Theorems 1, 2, 3 and use them to conduct inference about the responses to an *unconventional monetary policy shock*.

⁶All these variables are sourced from the data set of [Gertler and Karadi \(2015\)](#). We thank Peter Karadi for making their data set available to us.

4. THEOREMS

4.1. *Theorem 1*

In this section we consider the problem of finding the maximum response to an impulse in the j -th structural shock subject to m_z equality ('zero') restrictions and m_s inequality ('sign') restrictions. The focus on the maximum and the minimum is an intermediate step to conduct inference about the coefficients of the impulse-response function.

4.1.1. *Assumptions*

We make two assumptions on the sign and zero restrictions allowed in the model:

ASSUMPTION 1 The choice set in program (2.5) is not empty at μ .

This assumption simply requires that the identifying restrictions do not contradict each other.

Now, let $e_1^{m_s}, e_2^{m_s}, \dots, e_{m_s}^{m_s}$ denote the m_s different columns of the identity matrix \mathbb{I}_{m_s} . Let $e(k)$ denote an $m_s \times k$ matrix formed by collecting any of the $k \leq m_s$ columns of \mathbb{I}_{m_s} .

DEFINITION 2: We say that $Z(\mu)$ and $S(\mu)$ are linearly independent at μ if for any $k \in \mathbb{Z}$, $0 \leq k \leq m_s$ and any $e(k)$ the matrix

$$R(\mu; e(k)) \equiv [Z(\mu), S(\mu)e(k)] \in \mathbb{R}^{n \times (m_z + k)},$$

has full rank.

ASSUMPTION 2 $Z(\mu)$ and $S(\mu)$ are linearly independent at μ .

This assumption has two important implications. The first implication is that at most $n - 1$ constraints can be active at a solution of program (2.5) (in particular, it implies $m_z \leq n - 1$). The second implication is that it will allow us to characterize the maximum and minimum response in terms of *Karush-Kuhn-Tucker* conditions.

We verify (and discuss) this assumption for the UMP example in section 4.1.3.

4.1.2. *Algorithm*

We now show that the value function $\bar{v}_{k,i,j}(\mu)$ in (2.5) can be obtained by applying a simple algorithm.

Let r be the matrix that collects the columns of $Z(\mu)$ and the columns of $S(\mu)$ that are active at a solution of the program in (2.5). Our first observation is that the value function $\bar{v}_{k,i,j}(\mu)$ equals plus or minus

$$v_{k,i,j}(\mu; r) = \left(e_i' C_k(A) \Sigma^{1/2} M_{\Sigma^{1/2} r} \Sigma^{1/2} C_k(A)' e_i \right)^{1/2},$$

and the corresponding maximizer equals either:

$$x_+^*(\mu; r) \equiv \Sigma^{1/2} \left(M_{\Sigma^{1/2} r} \right) \Sigma^{1/2} C_k(A)' e_i / v_{k,i,j}(\mu; r),$$

or

$$x_-^*(\mu; r) \equiv -\Sigma^{1/2} \left(M_{\Sigma^{1/2} r} \right) \Sigma^{1/2} C_k(A)' e_i / v_{k,i,j}(\mu; r),$$

where $M_{\Sigma^{1/2} r} \equiv \mathbb{I}_n - \Sigma^{1/2} r (r' \Sigma r)^{-1} r' \Sigma^{1/2}$.

This result is shown formally in Lemma 1 in Appendix A.1 (where we also provide intuition). The lemma implies that if we knew the program's binding constraints, the value function could be computed directly—up to its sign—as $v_{k,i,j}(\mu, r)$. Moreover, the sign of value function is positive if $x_+^*(\mu; r)$ satisfies the inequality restrictions that are not included in r , and negative otherwise.

Let R denote the set of all possible matrices r that could arise from collecting the m_z columns of the matrix $Z(m)$ and $n - m_z - 1$ columns of the matrix $S(m)$. Take any c larger than

$$\bar{c} \equiv \max_{i,k} \left(e_i' C_k(A) \Sigma C_k(A)' e_i \right)^{1/2}.$$

The parameter c will be used to ‘penalize’ candidate solutions that do not satisfy the inequality restrictions in $S(\mu)$.⁷ The penalization works as follows:

Consider first the case in which $v_{k,i,j}(\mu; r) \neq 0$. Since $x_+^*(\mu; r)$ and $x_-^*(\mu; r)$ above are well defined, we can verify if these candidate solutions satisfy the sign restrictions that were not included in r (that is, we verify the *primal feasibility* of the solutions). If the primal feasibility condition is satisfied we store the candidate values; else we penalize them to guarantee that they are never a solution. More concisely, we define the auxiliary functions:

⁷ The constant \bar{c} is the maximum value of the following programs:

$$(4.1) \quad \bar{c} \equiv \max_{i,k} \sup_{B \in \mathbb{R}^{n \times n}} e_i' C_k(A) B e_j, \text{ s.t. } B B' = \Sigma.$$

$$\begin{aligned} f_{max}^+(\mu; r) &\equiv v_{k,i,j}(\mu; r) - 2(1 - \mathbf{1}_{m_s}(x_+(\mu; r)))c, \\ f_{max}^-(\mu; r) &\equiv -v_{k,i,j}(\mu; r) - 2(1 - \mathbf{1}_{m_s}(x_-(\mu; r)))c, \end{aligned}$$

where $\mathbf{1}_{m_s}(x) \equiv \mathbf{1}\{S(\mu)'x \geq \mathbf{0}_{m_s \times 1}\}$ is 1 if and only if x satisfies all the inequality restrictions in $S(\mu)$. The functions f_{max}^+, f_{max}^- allow us to keep track of the candidate values (and their feasibility) for each combination of active restrictions.

Consider now the penalization in the case in which $v_{k,i,j}(\mu; r) = 0$. This case is slightly different as the one considered in the previous paragraph, as the candidate solutions (x_+ and x_-) are not defined in this case. If there is a point $x^* \neq 0$ satisfying the equality restrictions in r and also the inequality restrictions that are not included in r , we set:

$$f_{max}^+(\mu; r) = f_{max}^-(\mu; r) = 0.$$

If no such point $x^* \neq 0$ exists, $v_{k,i,j}(\mu, r) = 0$ cannot be a solution and we set:

$$f_{max}^+(\mu; r) = f_{max}^-(\mu; r) = -2c.$$

The following theorem shows that we can compute the value function of the mathematical program (2.5) by selecting the maximum value of $\max\{f_{max}^+(\mu; r), f_{max}^-(\mu; r)\}$ over $r \in R$. That is, we can solve for $\bar{v}_{k,i,j}(\mu)$ by considering the different combinations of active restrictions and select the maximum value $\pm v_{k,i,j}(\mu, r)$ over them.

THEOREM 1 *Suppose that Assumptions 1 and 2 hold, then:*

$$\bar{v}_{k,i,j}(\mu) = \max_{r \in R} \left(\max\{f_{max}^+(\mu; r), f_{max}^-(\mu; r)\} \right).$$

The minimum value is obtained analogously.

PROOF: The intuition behind the proof is as follows. Note that any combination of active sign restrictions r for which $x_+(\mu; r)$ (or $x_-(\mu; r)$) is well-defined and feasible must be, by definition, no larger than $\bar{v}_{k,i,j}(\mu)$. Thus, we only have to show that

$$\max_{r \in R} \left(\max\{f_{max}^+(\mu; r), f_{max}^-(\mu; r)\} \right) \geq \bar{v}_{k,i,j}(\mu).$$

Since Lemma 1 showed that the value of the program (2.5) should be of the form $f_{max}^+(\mu; r)$ or $f_{max}^-(\mu; r)$ for some $r \in R$, the result must follow. The proof is formalized in Appendix A.2.

4.1.3. *Using the algorithm in the UMP example*

We verify Assumptions 1 and 2 at the estimated LS values of μ , denoted $\widehat{\mu}_T$. The simplest way to verify Assumption 1 is to consider the different candidate solutions for the different combinations of active constraints and check whether one of these solutions is feasible. For Assumption 2, note that regardless of the number of k columns selected from S the resulting matrix $R(\mu, e(k))$ will always have full column rank. Thus, Assumption 2 is also verified.⁸

We now use our algorithm to evaluate the identified set and report $\overline{v}_{k,i,j}(\widehat{\mu}_T)$ and $\underline{v}_{k,i,j}(\widehat{\mu}_T)$ for the cumulative IRFs.⁹ The scale in Figure 1 corresponds to a one standard deviation structural UMP shock.

We consider first the equality/inequality restrictions in Table I. Evaluating the endpoints of the identified for the 4 variables in the VAR, over 36 horizons, takes around 0.1 seconds. We then include an additional inequality restriction on the response of output to an expansionary UMP shock. Namely, we assume that even one period after the shock, the cumulative effect on IP cannot be negative ($e_2'(C_0 + C_1(A))B_1 \geq 0$). Figure 1 shows that the upper bounds of the identified sets under the two identification schemes almost overlap. The figure suggests that the noncontemporaneous constraint has thus little additional identification power.

There are at least two other ways of evaluating the maximum and minimum response (although only our algorithm is guaranteed to provide a global solution in a finite number of steps). One approach is to simply use a numerical solver (such as Matlab's `fmincon`) to get the value of the non-linear, non-convex program in (2.5). The result in Theorem 1 allows us to avoid the specification of the standard tuning parameters for numerical optimization routines (such as initial conditions, algorithms for the solver, tolerance levels for the solutions, and number of iterations).

Another approach is to rely on a version of the Bayesian algorithm in Uhlig (2005). Given reduced-form parameters μ and D draws of a unit vector $q \in \mathbb{R}^n$, one could report the maximum and minimum value over $\{\lambda_{k,i,j}(\mu, q^d)\}_{d=1}^D$. Note that such algorithm is essentially a random grid search approach to solve the program (2.5). The grid search approach underestimates the identified set. In our application the bias is negligible for $D = 10,000$ draws (the algorithm, however, takes around 300 seconds to run).


⁸Verifying Assumption 2 with more general restrictions requires additional work. For example, suppose that the researcher is interested in including the restriction:

$$e_2' C_1(A) B_1 \geq 0.$$

This restriction says that the UMP shock cannot decrease the growth rate in Industrial Production even one-period after the shock. Since $C_1(A) = A_1$, the vector $e_2' C_1(A)$ is equal to the second row of A_1 , which we can denote as $(A_{1,(2,1)}, A_{1,(2,2)}, A_{1,(2,3)}, A_{1,(2,4)})$. Assumption 2 will be satisfied as long as μ is such that $A_{1,(2,j)} \neq 0$ for all $j = 1, \dots, 4$, which means that each of the entries in the first lag of Y_{t-1} has predictive power on Y_t after controlling for the rest of the lags.

⁹The formula for the maximum (minimum) k -th period ahead *cumulative* IRF replaces $C_k(\widehat{A}_T)$ by $C_0(\widehat{A}_T) + C_1(\widehat{A}_T) + \dots + C_k(\widehat{A}_T)$.

Figure 1: Identified Set for the Cumulative Impulse Response Functions to a one standard deviation UMP shock (given $\hat{\mu}_T$) for two different identification schemes



(SOLID, BLUE LINE) Endpoints of the identified set for the cumulative responses given $\hat{\mu}_T$ and the equality/inequality restrictions in Table I. (BLUE, CROSSES) Endpoints of the identified set with the additional restriction that the cumulative response of IP to a UMP shock one month after impact is non-negative, $e_2'(C_0 + C_1(A))B_1 \geq 0$. Note that the upper bounds of the identified sets under the two identification schemes almost overlap.

4.2. Theorem 2

In this section we show that the endpoints of the identified set— $\underline{v}_{k,i,j}(\cdot)$ and $\bar{v}_{k,i,j}(\cdot)$ —are directionally differentiable functions of the reduced-form parameter μ . This result is the basis of our delta-method approach to conduct inference in set-identified SVARs.

4.2.1. Assumptions

In order to establish our differentiability result we need an additional regularity condition. Our key assumption is as follows:

ASSUMPTION 3 The matrices $Z(\cdot)$ and $S(\cdot)$ are differentiable at μ .

We are not aware of equality/inequality restrictions in the SVAR literature that do not satisfy this property. In particular, all the examples given in Subsection 2.2 of this paper satisfy Assumption 3 for every value of $\mu \in \mathcal{M}$.

4.2.2. Directional Differentiability

We will continue working with the auxiliary function $v_{k,i,j}(\mu; r(\mu))$, where we now explicitly acknowledge the possible dependence of r on μ . Lemma 2 in Appendix A.3 shows that if Assumptions 1-3 hold and $v_{k,i,j}(\mu; r(\mu)) \neq 0$, then the function $v_{k,i,j}(\mu; r(\mu))$ is differentiable with respect to $(\text{vec}(A)', \text{vec}(\Sigma)')$ with the derivative $\dot{v}_{k,i,j}(\mu; r(\mu))$ given by:

$$\begin{bmatrix} \frac{\partial v_{k,i,j}(\mu; r(\mu))}{\partial \text{vec}(A)} \\ \frac{\partial v_{k,i,j}(\mu; r(\mu))}{\partial \text{vec}(\Sigma)} \end{bmatrix} = \begin{bmatrix} \frac{\partial \text{vec}(C_k(A))}{\partial \text{vec}(A)} (x^*(\mu; r(\mu)) \otimes e_i) - \sum_{k=1}^l w_k^* \frac{\partial \text{vec}(r_k(\mu))}{\partial \text{vec}(A)} x^*(\mu; r(\mu)) \\ \lambda^* \Sigma^{-1} x^*(\mu; r(\mu)) \otimes \Sigma^{-1} x^*(\mu; r(\mu)) - \sum_{k=1}^l w_k^* \frac{\partial \text{vec}(r_k(\mu))}{\partial \text{vec}(\Sigma)} x^*(\mu; r(\mu)) \end{bmatrix},$$

where $r_k(\mu)$ denotes the k -th column of $r(\mu)$,

$$\begin{aligned} x^*(\mu; r(\mu)) &= \Sigma^{1/2} \left(M_{\Sigma^{1/2} r(\mu)} \right) \Sigma^{1/2} C_k(A)' e_i / v_{k,i,j}(\mu; r(\mu)), \\ \lambda^* &\equiv \frac{1}{2} v_{k,i,j}(\mu; r(\mu)), \quad w^* \equiv [r(\mu)' \Sigma r(\mu)]^{-1} r(\mu)' \Sigma C_k(A) e_i, \end{aligned}$$

and w_k^* is the k -th component of the vector w^* .¹⁰

¹⁰The *envelope theorem* sheds light on the derivative formula provided in Lemma 2. Note first that, by definition,

$$v_{k,i,j}(\mu; r(\mu)) = \max_{x \in \mathbb{R}^n} e_i' C_k(A) x \quad \text{s.t.} \quad x' \Sigma^{-1} x = 1 \quad \text{and} \quad r'(\mu) x = \mathbf{0}_{l \times 1}.$$

The auxiliary *Lagrangian function* of this problem is given by:

$$\mathcal{L}(x; \mu, r(\mu)) = (x' \otimes e_i') \text{vec}(C_k(A)) - \lambda \left((x' \otimes x') \text{vec}(\Sigma^{-1}) - 1 \right) - w' (r(\mu)' x),$$

where λ is the Lagrange multiplier corresponding to the quadratic equality restriction and $w \in \mathbb{R}^l$ is the

We now state the definition of directional differentiability and present our second Theorem.

DEFINITION 3: We say that the real-valued function v with domain $\mathcal{M} \subseteq \mathbb{R}^d$ is *directionally differentiable* at μ if for any $h \in \mathbb{R}^d$, any sequence $t_N \rightarrow \infty$, and any sequence $h_N \in \mathbb{R}^d$ such that $h_N \rightarrow h$ ($\mu + t_N h_N \in \mathcal{M}$), there exists a continuous function $\dot{v}_\mu : \mathbb{R}^d \rightarrow \mathbb{R}$ such that:

$$t_N \left(v(\mu + h_N/t_N) - v_{k,i,j}(\mu) \right) \rightarrow \dot{v}_\mu(h).$$

We refer to the function \dot{v}_μ as the directional derivative of $v(\cdot)$ at μ .¹¹

Let $X^*(\mu)$ denote the argmax of program (2.5). For $x \in X^*(\mu)$ let $r(\mu; x)$ denote the matrix that collects *all* elements of $Z(\mu)$ and $S(\mu)$ that are active at x .

THEOREM 2 *Suppose that Assumptions 1-3 hold. Suppose in addition $\bar{v}_{k,i,j}(\mu) > 0$. Then $\bar{v}_{k,i,j}$ is a directionally differentiable function of the reduced-form parameter μ with the directional derivative given by:*

$$(4.2) \quad \max_{x \in X^*(\mu)} \left[\dot{v}_{k,i,j}(\mu; r(\mu; x))' h \right].$$

Whenever $X^(\mu) = \{x^*\}$ is a singleton, the value function $\bar{v}_{k,i,j}(\mu)$ is fully differentiable with the derivative $\dot{v}_{k,i,j}(\mu; r(\mu; x^*))$.*¹²

PROOF: See Appendix A.4.

Theorem 4.2, p. 223 in [Fiacco and Ishizuka \(1990\)](#) and Theorem 4.24, p. 280 in the book of [Bonnans and Shapiro \(2000\)](#) present a generalized version of the envelope theorem. They show that—under suitable constraint qualifications—the directional derivative (in direction h and evaluated at parameter μ) of the largest and smallest value in a mathematical program with equality and inequality constraints is given by:

$$\sup_{x \in X^*(\mu)} \left[\nabla_\mu \mathcal{L}(x; \mu) h \right],$$

vector of Lagrange multipliers corresponding to the l equality restrictions. The envelope theorem suggests that $\dot{v}_{k,i,j}(\mu; r(\mu))$ is given by the formula in Lemma 2. This intuition is confirmed in the proof of Lemma 2 provided $v_{k,i,j}(\mu; r(\mu)) \neq 0$.

¹¹See p.172 in [Shapiro \(1991\)](#).

¹²If $\bar{v}_{k,i,j}(\mu) < 0$ the directional derivative simply becomes:

$$\max_{x \in X_*(\mu)} \left[-\dot{v}_{k,i,j}(\mu; r(\mu; x))' h \right],$$

and

$$\inf_{x \in X_*(\mu)} \left[\nabla_{\mu} \mathcal{L}(x; \mu) h \right],$$

provided there is a unique set of Lagrange Multipliers supporting the optimal solutions. Theorem 2 uses the results in Lemma 1 and Lemma 2 to verify this formula.

DELTA-METHOD VS. BOOTSTRAP: We also note that directionally differentiable functions have been a topic of recent research. [Fang and Santos \(2015\)](#) show that the standard bootstrap is not consistent when applied to parameters of the form $v(\mu)$, where v is a directionally differentiable function. [Kitagawa, Payne, and Montiel Olea \(2016\)](#) show that Bayesian credible sets based on the quantiles of the posterior distribution of $v(\mu)$ will be asymptotically equivalent to the frequentist bootstrap (which is not consistent in this case). These results imply that typical frequentist and Bayesian inference for directionally differentiable functions is not guaranteed to be consistent.

The next section shows that the special form of the directional derivative that arises in the class of SVARs models studied in this paper allows the researcher to conduct (computationally convenient) delta-method inference, with a slight adjustment on the standard errors. We note that the recent paper of [Hong and Li \(2017\)](#) provides an alternative frequentist point-wise consistent inference procedure for directionally differentiable functions of general form. Such an approach, however, has two drawbacks compared to our delta method. First, implementing the numerical delta-method in [Hong and Li \(2017\)](#) requires a user specified tuning parameter. Second, their procedure requires the evaluation of the value function for a large number of re-sampled values of μ (whereas our delta-method only requires the evaluation of the value functions at $\hat{\mu}$).

4.3. *Theorem 3*

This section proposes a delta-method interval of the form

$$\text{CS}_T(1 - \alpha) \equiv \left[\underline{v}_{k,i,j}(\hat{\mu}_T) - z_{1-\alpha/2} \hat{\sigma}_{(k,i,j),T}/\sqrt{T}, \bar{v}_{k,i,j}(\hat{\mu}_T) + z_{1-\alpha/2} \hat{\sigma}_{(k,i,j),T}/\sqrt{T} \right],$$

where

$$\hat{\mu}_T \equiv (\text{vec}(\hat{A}_T)', \text{vec}(\hat{\Sigma}_T)'),$$

is the LS estimator for μ defined as:

$$\hat{A}_T \equiv \left(\frac{1}{T} \sum_{t=1}^T Y_t X_t' \right) \left(\frac{1}{T} \sum_{t=1}^T X_t X_t' \right)^{-1}, \quad \hat{\Sigma}_T \equiv \frac{1}{T - np - 1} \sum_{t=1}^T \hat{\eta}_t \hat{\eta}_t',$$

with

$$X_t \equiv (Y_{t-1}', \dots, Y_{t-p}')', \quad \hat{\eta}_t \equiv Y_t - \hat{A}_T X_t.$$

We work under the assumption that $\sqrt{T}(\hat{\mu}_T - \mu)$ is asymptotically normal with some covariance matrix Ω .¹³

We use the results in Theorem 2 and the asymptotic normality of $\hat{\mu}_T$ to suggest the following formula for $\hat{\sigma}_{(k,i,j),T}$:

$$(4.3) \quad \hat{\sigma}_{(k,i,j),T} \equiv \max_{r \in R(\hat{\mu}_T)} \left(\dot{v}_{k,i,j}(\hat{\mu}_T; r)' \hat{\Omega}_T \dot{v}_{k,i,j}(\hat{\mu}_T; r) \right)^{\frac{1}{2}},$$

where $R(\hat{\mu}_T)$ is the set of *all* possible active constraints in program (2.5) evaluated at $\hat{\mu}_T$. Note that our procedure does not attempt to estimate neither the argmax nor the argmin of program (2.5).

FREQUENTIST COVERAGE: Let P denote the data generating process and let $\mathcal{I}_{k,i,j}^R(\mu(P))$ denote the identified set for the structural parameter $\lambda_{k,i,j}$ given the equality/inequality restrictions in $\mathcal{R}(\mu)$. This section shows that under our proposed specification of $\hat{\sigma}_{(k,i,j),T}$,

$$\liminf_{T \rightarrow \infty} \inf_{\lambda \in \mathcal{I}_{k,i,j}^R(\mu(P))} P(\lambda \in \text{CS}_T(1 - \alpha)) \geq 1 - \alpha.$$

Consequently, the delta-method interval presented in this paper is *point-wise consistent in level*.

¹³A common formula for $\hat{\Omega}$ based on the assumption of uncorrelated, possibly heteroskedastic structural innovations is given by:

$$\hat{\Omega}_T \equiv \left(\frac{1}{T} \sum_{t=1}^T \text{vec} \left([\hat{\eta}_t X_t', \hat{\eta}_t \hat{\eta}_t' - \hat{\Sigma}_T] \right) \text{vec} \left([\hat{\eta}_t X_t', \hat{\eta}_t \hat{\eta}_t' - \hat{\Sigma}_T] \right)' \right).$$

Our delta-method approach is also valid under the presence of time-series dependence in η_t (we only need a heteroskedasticity and autocorrelation robust estimator of Ω).

ROBUST BAYESIAN CREDIBILITY: We also show that under some regularity conditions our delta-method interval has, asymptotically, robust Bayesian credibility of at least the nominal level. To formalize this statement Let P^* denote some prior for the structural parameters (A_1, \dots, A_p, B) and let $\lambda_{k,i,j}(A, B) \in \mathbb{R}$ denote the structural coefficient of interest. For a given square root of $\Sigma \equiv BB'$ define the orthogonal matrix $Q \equiv \Sigma^{-1/2}B$. It is well known that a prior P^* can be written as $(P_\mu^*, P_{Q|\mu}^*)$, where P_μ^* is a prior on the reduced-form parameters, and $P_{Q|\mu}^*$ is a prior on the orthogonal matrix, conditional on μ . Following this notation, let $\mathcal{P}(P_\mu^*)$ denote the class of prior distributions such that $\mu^* \sim P_\mu^*$.

Define the *Robust Bayes Credibility* of our delta-method region as:

$$(4.4) \quad RBC(Y_1, \dots, Y_T) \equiv \inf_{P^* \in \mathcal{P}(P_\mu^*)} P^* \left(\lambda(A, B) \in \text{CS}_T(1 - \alpha) \mid Y_1, \dots, Y_T \right).$$

We show that if the prior for the reduced-form parameters μ satisfies the *Bernstein-von Mises Theorem* and the bounds of the identified set are *differentiable* then for any $\epsilon > 0$:

$$\lim_{T \rightarrow \infty} P(RBC(Y_1, \dots, Y_T) < 1 - \alpha - \epsilon) = 0$$

Thus, our delta-method interval has *asymptotic* robust Bayesian credibility of at least $1 - \alpha$.

We now describe the main large-sample assumptions used to establish the frequentist coverage and the robust Bayesian credibility of our delta-method interval.

4.3.1. *Assumptions*

The SVAR parameters (A_1, \dots, A_p, B, F) define a probability distribution, denoted P , over the data observed by the econometrician. Our main assumptions concerning P are as follows. First, we assume that the LS estimator $\hat{\mu}_T$ is asymptotically normal with a covariance matrix that can be estimated consistently.

ASSUMPTION 4 (Asymptotic Normality of $\hat{\mu}_T$) The data generating process P is such that for $\mu(P) \in \mathbb{R}^d$:

$$\sqrt{T}(\hat{\mu}_T - \mu(P)) \xrightarrow{d} \zeta(P) \sim \mathcal{N}_d(\mathbf{0}, \Omega(P)),$$

and

$$\hat{\Omega}_T \xrightarrow{P} \Omega(P).$$

Second, we will assume that the prior P_μ^* used to compute robust Bayesian credibility and the data generating process P satisfy the Bernstein von-Mises Theorem in Ghosal, Ghosh, and Samanta (1995). More precisely, we assume that:

ASSUMPTION 5 (Bernstein-von Mises Theorem)

$$\sup_{B \in \mathcal{B}(\mathbb{R}^d)} \left| P_\mu^* \left(\sqrt{T}(\mu^* - \hat{\mu}_T) \in B \mid Y_1, \dots, Y_T \right) - \mathbb{P}(\zeta(P) \in B) \right| \xrightarrow{P} 0,$$

where $\zeta(P) \sim \mathcal{N}_d(\mathbf{0}, \Omega(P))$, and $\mathcal{B}(\mathbb{R}^d)$ is the set of all Borel measurable sets in \mathbb{R}^d .

Assumption 5 is satisfied for Normal-Inverse Wishart prior (see Uhlig (2005)) in a VAR model with Gaussian i.i.d. errors (see Gafarov, Meier, and Montiel Olea (2016)). More generally, if the VAR reduced-form errors are i.i.d. Gaussian, Theorem 1 and 2 in Ghosal et al. (1995) imply that Assumption 5 will be satisfied whenever P_μ^* has a continuous density at μ with polynomial majorants.

4.3.2. Large-sample coverage and robust Bayesian credibility

Dümbgen (1993), Shapiro (1991), and Fang and Santos (2015) have shown if v is a directionally differentiable function with directional derivative $\dot{v}_\mu(h)$ (in direction h evaluated at μ) then:

$$\sqrt{T}(v(\hat{\mu}_T) - v(\mu)) \xrightarrow{d} \dot{v}_\mu(\zeta),$$

whenever Assumption 4 holds. Theorem 2 in the previous section established that the directional derivative of $\bar{v}_{k,i,j}$ —in direction h evaluated at μ —is given by:

$$\max_{x \in X^*(\mu)} \left[\dot{v}_{k,i,j}(\mu; r(\mu; x))' h \right].$$

where $X^*(\mu)$ is the argmax of program (2.5) at μ . Thus, Theorem 2 and Assumption 4 imply that:

$$\sqrt{T}(\bar{v}_{k,i,j}(\hat{\mu}_T) - \bar{v}_{k,i,j}(\mu)) \xrightarrow{d} \max_{x \in X^*(\mu)} \left[\dot{v}_{k,i,j}(\mu; r(\mu; x))' \zeta \right],$$

where

$$\dot{v}_{k,i,j}(\mu; r(\mu; x))' \zeta \sim \mathcal{N}_1 \left(0, \dot{v}_{k,i,j}(\mu; r(\mu; x))' \Omega \dot{v}_{k,i,j}(\mu; r(\mu; x)) \right).$$

Our suggestion—which exploits the specific form of the directional derivative in the SVAR context—is to consider:

$$\hat{\sigma}_{(k,i,j),T} \equiv \max_{r \in R(\hat{\mu}_T)} \left(\dot{v}_{k,i,j}(\hat{\mu}_T; r)' \hat{\Omega}_T \dot{v}_{k,i,j}(\hat{\mu}_T; r) \right)^{\frac{1}{2}},$$

where $R(\hat{\mu}_T)$ is the set of *all* the different collections of active constraints evaluated at $\hat{\mu}_T$.

The idea is that $\hat{\sigma}_{(k,i,j),T}$ converges in probability to:

$$\max_{r \in R(\mu)} \left(\dot{v}_{k,i,j}(\mu; r)' \Omega \dot{v}_{k,i,j}(\mu; r) \right)^{\frac{1}{2}}$$

which is larger than or equal to:

$$\max_{x \in X^*(\mu)} \left(\dot{v}_{k,i,j}(\mu; r(\mu, x))' \Omega \dot{v}_{k,i,j}(\mu; r(\mu, x)) \right)^{\frac{1}{2}}.$$

Thus, our formula for the standard error implies there is no need to estimate neither the argmax nor the argmin of the program defining $\bar{v}(\mu)$. The suggested confidence interval is shown to be point-wise consistent in level.¹⁴ We also show that our delta-method interval has, asymptotically, robust Bayesian credibility of at least the nominal level (provided some regularity conditions are satisfied). These two properties are formalized in the following theorem.

THEOREM 3 *Let $\hat{\sigma}_{(k,i,j),T}$ be defined as in (4.3). Suppose that the asymptotic variance of the candidate value functions in $X^*(\mu)$ and $X_*(\mu)$ are strictly positive; that is:*

$$\min_{x \in X_*(\mu(P)) \cup X^*(\mu(P))} \|\Omega^{1/2}(P) \dot{v}_{k,i,j}(\mu(P); r(\mu(P); x))\| > 0.$$

a) *If Assumptions 1-4 are satisfied at $\mu = \mu(P)$, then:*

$$\liminf_{T \rightarrow \infty} \inf_{\lambda \in \mathcal{I}_{k,i,j}^R(\mu(P))} P(\lambda \in CS_T(1 - \alpha)) \geq 1 - \alpha.$$

b) *If in addition Assumption 5 holds and $X^*(\mu(P))$ and $X_*(\mu(P))$ are both singletons, then for any $\epsilon > 0$:*

$$\lim_{T \rightarrow \infty} P \left(\inf_{P^* \in \mathcal{P}(P_\mu^*)} P^* \left(\lambda(A, B) \in CS_T(1 - \alpha) \mid Y_1, \dots, Y_T \right) < 1 - \alpha - \epsilon \right) = 0.$$

PROOF: See Appendix A.5.

Note that we have assumed that the identified set is non-empty at μ , and we have also showed that under Assumptions 1-4 the probability of observing an empty identified set at $\hat{\mu}_T$ converges to zero as the sample size grows to infinity. It is of course still possible to observe an empty identified set at a given realization of $\hat{\mu}_T$. In this case, our algorithm will report a maximum response equal to $-c$ and a minimum response equal to c .¹⁵

¹⁴The question of how to build a *uniformly consistent in level*, delta-method confidence set for a set-identified parameter is beyond the scope of this paper. For the readers interested in uniform inference for set-identified parameters in SVARs our suggestion is to apply the projection approach developed in [Gafarov et al. \(2016\)](#). In comparison, the delta-method approach described in this paper is faster to implement.

¹⁵In our Matlab implementation, this will generate a warning message asking the user to drop restrictions.

4.3.3. Monte-Carlo Evidence

FREQUENTIST COVERAGE: We conduct a simple Monte-Carlo exercise to study the coverage probability of our delta-method interval. We set $(1 - \alpha) = .68$ implying that $z_{1-\alpha/2} = .9945$. Instead of generating new draws of (Y_1, \dots, Y_T) , we generate 10,000 draws of $\hat{\mu}_T$ directly from its asymptotic normal distribution $\mathcal{N}_d(\mu, \Omega/T)$ (where we fix the values of μ and Ω at its estimated values in the UMP example). We decided to proceed in this way in order to ‘enforce’ the asymptotic normality assumption for $\hat{\mu}_T$ (which is the key requirement in part a) of Theorem 3). We set $T = 342$ which corresponds to the number of periods in our empirical application.

For each ‘draw’ of $\hat{\mu}_T$ (denoted μ^*) we compute the confidence interval:

$$\left[\underline{v}_{k,i,j}(\mu^*) - .9945 \sigma_{(k,i,j),T}^* / \sqrt{T}, \bar{v}_{k,i,j}(\mu^*) + .9945 \sigma_{(k,i,j),T}^* / \sqrt{T} \right],$$

where we treat Ω as known to compute the formula for the standard error $\hat{\sigma}_{(k,i,j),T}$. We do this to assume away any problem concerning the estimation of Ω (as Theorem 3 assumes that we have a consistent estimator for the asymptotic variance of $\hat{\mu}_T$).

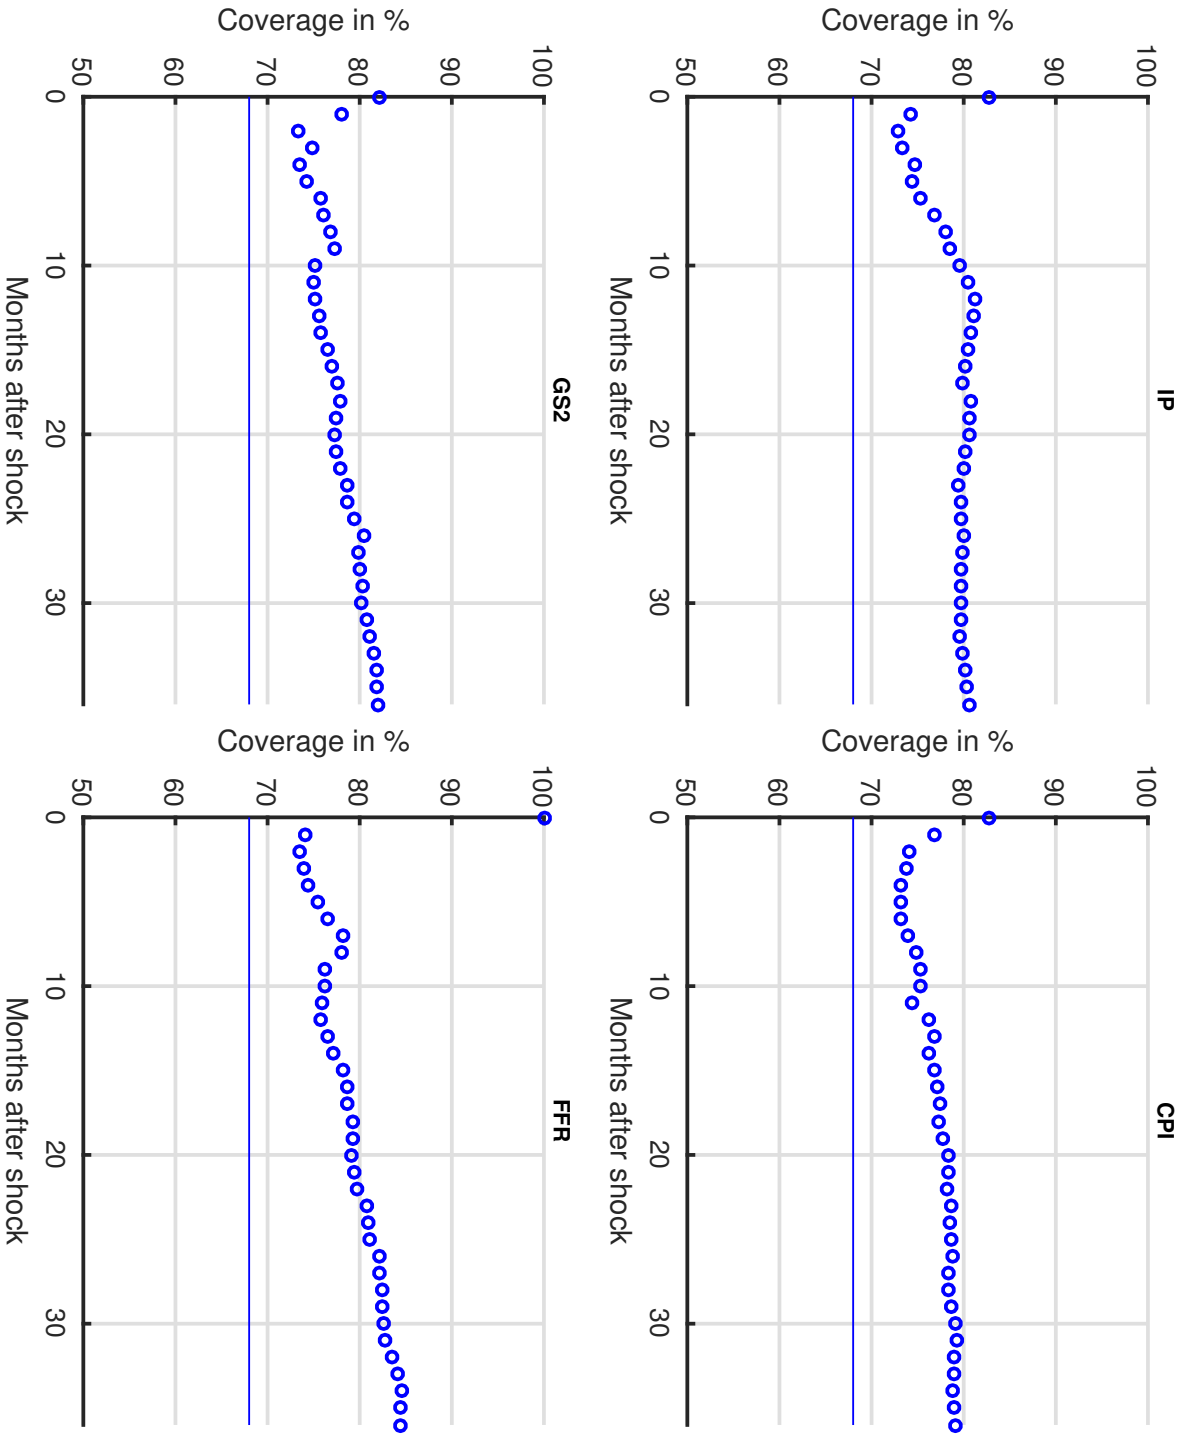
Finally, we check whether $[\underline{v}_{k,i,j}(\mu), \bar{v}_{k,i,j}(\mu)]$ is contained in the confidence interval corresponding to each draw μ^* from $\mathcal{N}_d(\hat{\mu}_T, \hat{\Omega}_T)$. The estimated probability provides a lower bound on the coverage of the identified parameter. The results are reported in Figure 2. We note that the simulated coverage probability lies between 68% and 84% (except for the contemporaneous IRF for FFR which is equal to zero by assumption). The simulated coverage probability is higher than the nominal size of 68%. This is consistent with our theorem, as we are using a standard error that protects against potential violations of full differentiability (even when the function is differentiable at μ).¹⁶

ROBUST BAYESIAN CREDIBILITY IN THE UMP APPLICATION: We also compute the robust Bayesian credibility of our delta-method interval based on an uninformative Normal-Inverse Wishart prior on μ (following Uhlig (2005)). Namely, we generate 10,000 draws of μ^* from the posterior distribution and report the share of draws for which $[\underline{v}_{k,i,j}(\mu^*), \bar{v}_{k,i,j}(\mu^*)]$ is contained in

$$\left[\underline{v}_{k,i,j}(\hat{\mu}_T) - .9945 \hat{\sigma}_{(k,i,j),T} / \sqrt{T}, \bar{v}_{k,i,j}(\hat{\mu}_T) + .9945 \hat{\sigma}_{(k,i,j),T} / \sqrt{T} \right].$$


The results are provided in Figure 3. The simulated credibility is larger or close to the nominal level of 68%, which is consistent with the part b of Theorem 3. We also report the robust Bayesian credibility based on the asymptotic normal approximation in Figure 5 on Appendix B.2.

¹⁶One can use the ideas of Freyberger and Horowitz (2015) to propose an alternative estimator for the standard error which could deliver yet tighter CS. We leave this extension for further research.

Figure 2: Monte-Carlo coverage probability based on the model $\mu^* \sim \mathcal{N}(\hat{\mu}_T, \hat{\Omega}_T/T)$, $T = 342$.

(CIRCLES) Monte-Carlo estimate of the probability $P\left(\left[\hat{u}_{k,i,j}(\hat{\mu}_T), \bar{u}_{k,i,j}(\hat{\mu}_T)\right] \subset \left[u_{k,i,j}(\mu^*) - .9945 \sigma_{(k,i,j),T}^* / \sqrt{T}, \bar{u}_{k,i,j}(\mu^*) + .9945 \sigma_{(k,i,j),T}^* / \sqrt{T}\right]\right)$ for the model $\mu^* \sim \mathcal{N}(\hat{\mu}_T, \hat{\Omega}_T)$, with $T = 342$. The values $\hat{\mu}_T$ and $\hat{\Omega}_T$ correspond, respectively, to the estimators of the reduced-form parameter and its asymptotic covariance matrix in the UMP application. (SOLID LINE) Nominal confidence level for the delta-method confidence interval (68%).

Figure 3: Robust Bayesian credibility of the delta-method interval based on the posterior distribution corresponding to an uninformative Normal-Inverse Wishart prior on μ^* (as in Uhlig (2005)), $T = 342$.



(CIRCLES) Monte-Carlo estimate of the probability $P_{\mu}^* \left(\left[\underline{v}_{k,i,j}(\mu^*), \bar{v}_{k,i,j}(\mu^*) \right] \subset \left[\underline{v}_{k,i,j}(\hat{\mu}_T) - .9945\hat{\sigma}_{(k,i,j),T}/\sqrt{T}, \bar{v}_{k,i,j}(\hat{\mu}_T) + .9945\hat{\sigma}_{(k,i,j),T}/\sqrt{T} \right] \mid Y_1, \dots, Y_T \right)$ based on the posterior distribution associated to an uninformative Normal-Inverse Wishart prior on μ^* (as in Uhlig (2005)) with $T = 342$. The values $\hat{\mu}_T$ and $\hat{\Omega}_T$ correspond, respectively, to the estimators of the reduced-form parameter and its asymptotic covariance matrix in the UMP application. (SOLID LINE) Nominal level of the delta-method interval (68%).

5. UNCONVENTIONAL MONETARY POLICY SHOCKS

In August 2010 the Federal Open Market Committee announced: “*The Committee will keep constant the Federal Reserve’s holdings of securities at their current level by reinvesting principal payments from agency debt and agency mortgage-backed securities in longer-term Treasury securities.*” This announcement was an important prelude for the second part of the Quantitative Easing program (QE2) (see p. 244 in [Krishnamurthy and Vissing-Jorgensen \(2011\)](#) for a detailed discussion). In addition, this announcement generated a drop in the intraday yield for two- and ten- year treasury bond. In fact, from the end of July 2010 to the end of August 2010 the 2 year Treasury bond rate fell by 10 basis points.


Figure 4 uses our delta-method approach to construct confidence bands for the evolution of the levels of the four variables in the monetary SVAR. We fix all the variables at their level on July 2010 and we trace their evolution (over a 12-month window) according to the confidence set for their cumulative responses. The motivation for this exercise is as follows. Suppose that—back in August 2010—an econometrician is asked to provide confidence bands for the evolution of IP, CPI, 2YTB, and FF after the August 2010 announcement of the Federal Open Market Committee (FOMC). The econometrician observes the realization of the macroeconomic variables from July 1979 until August 2010, but decides to deliberately ignore the two years of data after the crisis (to avoid introducing structural changes, stochastic volatility, or any other feature that will complicate the estimation of the VAR).

The econometrician uses the data until December 2007—one semester before the financial crisis—to conduct delta-method inference on the cumulative responses to a one standard deviation unconventional monetary policy shock. The econometrician then uses these cumulative responses to get a rough idea of the evolution of the variables (in levels) following the announcement of the Federal Reserve in August 2010. The econometrician assumes there is a linear trend for CPI/IP, and ignores sampling uncertainty coming from the trend estimation in reporting the bands.

An ex-post evaluation of this exercise (over a window of 12 months) is reported in Figure 4.¹⁷ We note that the observed dynamics for CPI, IP, GS2, and FFR from August 2010 to July 2011 fall within the bounds motivated by our delta-method interval. We also note that our delta-method interval misses the observed value at most three out of 12 months, which means that our 68% confidence set covers each of these variables at least 75% of the time. We also report the 68% Bayesian credible sets.

¹⁷The reason to focus in a 12-month window is to cover the period between the QE2 announcement and the announcement of the so-called “Operation Twist” in September 2011. See <http://www.federalreserve.gov/newsevents/press/monetary/20110921a.htm>.

Figure 4: Delta-Method Interval for CPI, IP, 2yTB, FF after the August 2010 announcement



(SHADED AREA) Evolution of the Levels CPI, IP, 2yTB, and FF based on our 68% delta method confidence bands for the coefficients of Cumulative Impulse-Response Functions. (SOLID LINE) Observed Levels of CPI, IP, 2yTB, and FF from December 2009 to July 2011. Both the CPI index and the IP index were normalized to have a starting value of 100. (DASHED LINE) Evolution of the Levels CPI, IP, 2yTB, and FF based on the 68% credible set constructed using the priors in Uhlig (2005).

COMPUTATIONAL COST: We close this section with some comments regarding the computational cost of our delta-method procedure. Most of the work to compute the endpoints of the identified set and its derivatives is analytical. Consequently, practitioners can expect the computational burden of our procedure to be low. We note that the implementation of our delta-method interval in the running example takes only around .15 seconds (using a standard Laptop @2.4GHz IntelCore i7).

COMPARISON WITH THE PROJECTION APPROACH: Figure 6 in Appendix B.2 presents a comparison between the delta-method approach and the *projection* approach recently proposed by Gafarov et al. (2016) [GMM16]. The projection approach has two theoretical properties that we were not able to verify for the delta-method. First, projection is consistent in level *uniformly* over a reasonable class of data generating processes. Second, projection yields valid *simultaneous* inference; that is, it covers the whole impulse-response function (across different horizons and different variables) and not only its scalar coefficients.¹⁸ We note that in our application the projection confidence interval (which is wider than the delta-method bands) contains the realized value of IP, CPI, 2YTB, and FF for every horizon under consideration.

COMPARISON WITH GK ROBUST APPROACH: Figure 7 in the Appendix reports the robust-Bayesian credible set in Giacomini and Kitagawa (2015). The implementation of the robust-Bayes credible set (based on 10,000 posterior draws and using our algorithm to evaluate the endpoints) took around 9,106 seconds.¹⁹

COMPARISON WITH MSG: Figure 9 in the Appendix reports the 68% Bonferroni confidence set of Moon et al. (2013).²⁰ Appendix B.1.1 describes the algorithm and related computational issues. The computational cost is approximately 4,000 seconds on a single core machine for 10,000 grid points.

It is hard to provide a general theoretical comparison of the length of the Bonferroni CS and the delta method. The efficiency ranking of the two procedures is likely depend on the particular DGP. One can see that, in our illustrative example, the 68% delta method CS is tighter than the corresponding Bonferroni CS with the same nominal level for almost all combinations of the horizons and time series. One possible explanation behind the larger length of Moon et al. (2013) is that their procedure is *uniformly* consistent in level over the class of GDPs for which the reduced form impulse response functions converge to a normal

¹⁸While our paper focuses on point-wise inference, it is straightforward to provide joint inference by applying Bonferroni correction to the significance level. Figure 8 compares confidence sets that cover not only a single impulse response but the impulse response functions of all variables and all horizons of interest. We compare our Delta-method results when using a Bonferroni-correction with Inoue and Kilian (2013)'s joint Bayes credible set for impulse response functions using the priors for the reduced-form parameters in Uhlig (2005). See Appendix B.1.2.

¹⁹Out of which 1,266 seconds were used just to compute the identified set for each posterior draw, and the remaining time to translate the posterior bounds into the GK robust bounds

²⁰Moon et al. (2013) also propose a projection-based CS which is a special case of the Bonferroni CS. There is no clear theoretical ranking of the various CS proposed in that paper so we chosen the least computationally intensive variation.

distribution. We note that our delta-method does not have this property.

6. CONCLUSION

This paper focused on set-identified structural VAR models that impose equality and inequality restrictions to set-identify only one structural shock. For this class of models, the endpoints of the identified set have special properties that allow an intuitive and computationally simple approach to conduct frequentist and (asymptotic) robust Bayes inference. Specifically, the paper made three contributions:

(i) We presented an algorithm to compute—for each horizon, each variable, a fixed vector of reduced-form parameters, and a given collection of equality and/or inequality restrictions—the largest and smallest value of the coefficients of the structural IRF (see Theorem 1). Our algorithm did not require random sampling from the space of orthogonal matrices or unit vectors. Instead, we treated the bounds of the identified set as the *maximum and minimum value* of a mathematical program whose solutions we were able to characterize analytically. Our algorithm can be used outside our delta-method framework (for example, in computing the maximum and minimum response for the [Giacomini and Kitagawa \(2015\)](#) robust Bayes approach).

(ii) We provided sufficient conditions under which the largest and smallest value of the structural parameters are directionally differentiable functions of the reduced-form parameters (see Theorem 2). This result also seems to be of interest in its own right and could be used to explore the frequentist properties of the robust-Bayesian procedure in [Giacomini and Kitagawa \(2015\)](#).

(iii) Finally, we proposed a computationally convenient delta-method approach to conduct inference for the set-identified coefficients of the structural IRF. We presented sufficient conditions to guarantee the point-wise consistency in level and asymptotic robust Bayes credibility of our suggested inference approach. We note that the delta-method in this paper exploited the structure of the directional derivative.

REFERENCES

- AN, L. AND J. WANG (2012): “Exchange rate pass-through: Evidence based on vector autoregression with sign restrictions,” *Open Economies Review*, 23, 359–380.
- ARIAS, J., J. F. RUBIO-RAMIREZ, AND D. F. WAGGONER (2014): “Inference Based on SVAR Identified with Sign and Zero Restrictions: Theory and Applications,” *Working paper, Duke University*.
- BAUMEISTER, C. AND L. BENATI (2013): “Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound,” *International Journal of Central Banking*, 9, 165–212.
- BAUMEISTER, C. AND J. HAMILTON (2015): “Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information,” *Econometrica*, 5, 1963–1999.

- BEAUDRY, P., D. NAM, AND J. WANG (2014): “Do mood swings drive business cycles and is it rational?” *NBER Working Paper (w17651)*.
- BENKWITZ, A., M. H. NEUMANN, AND H. LÜTEKPOHL (2000): “Problems related to confidence intervals for impulse responses of autoregressive processes,” *Econometric Reviews*, 19, 69–103.
- BHATTARAI, S., G. B. EGGERTSSON, AND B. GAFAROV (2014): “Time Consistency and the Duration of Government Debt: A Signalling Theory of Quantitative Easing,” *Working paper, Penn State University*.
- BLANCHARD, O. J. AND D. QUAH (1989): “The dynamic effects of aggregate demand and supply disturbances,” *The American Economic Review*, 79, 655–673.
- BONNANS, J. F. AND A. SHAPIRO (2000): *Perturbation analysis of optimization problems*, Springer.
- DOLADO, J. J. AND H. LÜTEKPOHL (1996): “Making Wald tests work for cointegrated VAR systems,” *Econometric Reviews*, 15, 369–386.
- DÜMBGEN, L. (1993): “On nondifferentiable functions and the bootstrap,” *Probability Theory and Related Fields*, 95, 125–140.
- FANG, Z. AND A. SANTOS (2015): “Inference on Directionally Differentiable Functions,” *Working paper, University of California at San Diego*.
- FIACCO, A. V. AND Y. ISHIZUKA (1990): “Sensitivity and stability analysis for nonlinear programming,” *Annals of Operations Research*, 27, 215–235.
- FREYBERGER, J. AND J. L. HOROWITZ (2015): “Identification and shape restrictions in nonparametric instrumental variables estimation,” *Journal of Econometrics*, 189, 41–53.
- FUJITA, S. (2011): “Dynamics of worker flows and vacancies: evidence from the sign restriction approach,” *Journal of Applied Econometrics*, 26, 89–121.
- GAFAROV, B., M. MEIER, AND J. L. MONTIEL OLEA (2016): “Projection Inference for Set-Identified SVARs,” *Working paper, New York University*.
- GERTLER, M. AND P. KARADI (2015): “Monetary Policy Surprises, Credit Costs and Economic Activity,” *American Economic Journal: Macroeconomics*, 7, 44–76.
- GHOSAL, S., J. K. GHOSH, AND T. SAMANTA (1995): “On convergence of posterior distributions,” *The Annals of Statistics*, 2145–2152.
- GIACOMINI, R. AND T. KITAGAWA (2015): “Robust Inference about partially identified SVARs,” *Working Paper, University College London*.
- GUSTAFSON, P. (2009): “What are the limits of posterior distributions arising from nonidentified models, and why should we care?” *Journal of the American Statistical Association*, 104, 1682–1695.
- HONG, H. AND J. LI (2017): “The numerical delta-method,” *Working Paper, Stanford University*.
- INOUE, A. AND L. KILIAN (2002): “Bootstrapping autoregressive processes with possible unit roots,” *Econometrica*, 70, 377–391.
- (2013): “Inference on impulse response functions in structural VAR models,” *Journal of Econometrics*, 177, 1–13.

- KILIAN, L. (1998): “Small-sample confidence intervals for impulse response functions,” *Review of economics and statistics*, 80, 218–230.
- KILIAN, L. AND D. P. 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.
- KITAGAWA, T., J. PAYNE, AND J. L. MONTIEL OLEA (2016): “Calibrated Inference for nondifferentiable functions,” *Working paper, New York University*.
- KRISHNAMURTHY, A. AND A. VISSING-JORGENSEN (2011): “The effects of quantitative easing on interest rates: channels and implications for policy,” *Brookings Papers on Economic Activity*.
- LÜTKEPOHL, H. (1990): “Asymptotic distributions of impulse response functions and forecast error variance decompositions of vector autoregressive models,” *The Review of Economics and Statistics*, 116–125.
- (2007): *New Introduction to Multiple Time Series Analysis*, Springer.
- MOON, H. R. AND F. SCHORFHEIDE (2012): “Bayesian and frequentist inference in partially identified models,” *Econometrica*, 80, 755–782.
- MOON, H. R., F. SCHORFHEIDE, AND E. GRANZIERA (2013): “Inference for VARs Identified with Sign Restrictions,” *University of Southern California*.
- MOUNTFORD, A. AND H. UHLIG (2009): “What are the Effects of Fiscal Policy Shocks?” *Journal of Applied Econometrics*, 24, 960–992.
- NOCEDAL, J. AND S. WRIGHT (2006): *Numerical optimization*, Springer Science & Business Media, second edition ed.
- OK, E. A. (2007): *Real analysis with economic applications*, vol. 10, Princeton University Press.
- POIRIER, D. J. (1998): “Revising beliefs in nonidentified models,” *Econometric Theory*, 14, 483–509.
- RUBIO-RAMIREZ, J., D. CALDARA, AND J. ARIAS (2015): “The Systematic Component of Monetary Policy in SVARs: An Agnostic Identification Procedure,” *Working Paper, Board of Governors of the Federal Reserve*.
- SHAPIRO, A. (1991): “Asymptotic analysis of stochastic programs,” *Annals of Operations Research*, 30, 169–186.
- SIMS, C. A., J. H. STOCK, AND M. W. WATSON (1990): “Inference in linear time series models with some unit roots,” *Econometrica: Journal of the Econometric Society*, 113–144.
- TODA, H. Y. AND T. YAMAMOTO (1995): “Statistical inference in vector autoregressions with possibly integrated processes,” *Journal of econometrics*, 66, 225–250.
- 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.
- VARGAS-SILVA, C. (2008): “Monetary policy and the US housing market: A VAR analysis imposing sign restrictions,” *Journal of Macroeconomics*, 30, 977–990.