Joint Bayesian Inference about Impulse Responses in VAR Models

Atsushi Inoue and Lutz Kilian

Joint Bayesian Inference about Impulse Responses in VAR Models[*]

Atsushi Inoue[†] and Lutz Kilian[‡]

First draft: January 21, 2020
This version: July 5, 2020

Abstract

Structural VAR models are routinely estimated by Bayesian methods. Several recent studies have voiced concerns about the common use of posterior median (or mean) response functions in applied VAR analysis. In this paper, we show that these response functions can be misleading because in empirically relevant settings there need not exist a posterior draw for the impulse response function that matches the posterior median or mean response function, even as the number of posterior draws approaches infinity. As a result, the use of these summary statistics may distort the shape of the impulse response function which is of foremost interest in applied work. The same concern applies to error bands based on the upper and lower quantiles of the marginal posterior distributions of the impulse responses. In addition, these error bands fail to capture the full uncertainty about the estimates of the structural impulse responses. In response to these concerns, we propose new estimators of impulse response functions under quadratic loss, under absolute loss and under Dirac delta loss that are consistent with Bayesian statistical decision theory, that are optimal in the relevant sense, that respect the dynamics of the impulse response functions and that are easy to implement. We also propose joint credible sets for these estimators derived under the same loss function. Our analysis covers a much wider range of structural VAR models than previous proposals in the literature including models that combine short-run and long-run exclusion restrictions and models that combine zero restrictions, sign restrictions and narrative restrictions.

[*†]Atsushi Inoue, Vanderbilt University, Department of Economics, Nashville, TN 37235-1819. E-mail: at-sushi.inoue@vanderbilt.edu.

[‡]Lutz Kilian, Federal Reserve Bank of Dallas, 2200 N. Pearl St., Dallas, TX 75201, USA, and CEPR. E-mail:lkilian2019@gmail.com (corresponding author).

1 Introduction

It is standard practice in empirical macroeconomics to report measures of impulse response functions based on structural vector autoregressive (VAR) models (see Kilian and Lütkepohl 2017). Structural VAR models are routinely estimated by Bayesian methods. Applied VAR users typically treat the posterior median response function, defined as the vector of medians obtained from the marginal posterior distribution of each individual impulse response, as the estimate of the impulse response function. Inference about median response functions is based on vectors of the $\alpha/2$ and $1 - \alpha/2$ probability quantiles of the marginal posterior distributions of the impulse responses.

The existing literature has identified three main concerns with this approach. One concern voiced by Fry and Pagan (2011), Kilian and Murphy (2012) and Inoue and Kilian (2013), among others, is that the marginal medians underlying the median response function may come from different structural models. Thus, median response functions tend to conflate the dynamic responses implied by different structural models, calling into question the economic interpretation of the median response function. The same concern applies to the impulse response error bands reported in the literature.

Another concern is that these two-dimensional error bands are consistent with any shape of the response functions that fits within the error band. This point was first raised by Sims and Zha (1999, p. 1113) who observe that "it has been conventional in the applied literature to construct a $1 - \alpha$ probability interval separately at each point of a response horizon, then to plot the response itself with the upper and lower limits of the probability intervals as three lines. The resulting band ... does not directly give much information about the forms of deviation from the ... estimate of the response function that are most likely."

A third concern is that the lower and upper interval endpoints of conventional error bands do not represent likely variation in the impulse response functions. As noted by Sims and Zha (1999),

bands constructed by "connecting the dots" tend to be misleading because it is extremely unlikely in economic applications that the uncertainty about the impulse response estimates is independent across horizons. Moreover, as emphasized in Inoue and Kilian (2016) and Kilian and Lütkepohl (2017), the dependence across impulse response estimators matters not only for conducting joint inference across a range of horizons for a given impulse response function, but it matters equally for conducting joint inference about multiple impulse response functions. Since most researchers are interested in assessing the implications of economic theory for a range of different impulse response functions simultaneously, as exemplified by Blanchard (1989) and Sims and Zha (1999), it is essential to take account of the dependence of all structural impulse responses of interest, not just of the responses in a given impulse response function.

Although posterior median response functions are ubiquitous in applied work, some researchers prefer to evaluate the structural impulse responses under quadratic loss and report the vector of posterior means of the impulse responses. The latter approach, however, still conflates the dynamics of different structural models. Moreover, it is not obvious how to conduct inference about the mean response function.

In this paper, we provide a unified approach to evaluating vectors of structural impulse responses from a Bayesian point of view. We consider three alternative loss functions, namely absolute loss, quadratic loss and Dirac delta loss, corresponding to estimators based on the median, mean, and mode. Our analysis covers both the construction of optimal impulse response estimators and joint inference under the same loss function used for constructing the impulse response estimate.

First, we propose a new class of estimators of structural impulse response functions designed to be valid under absolute or quadratic loss, respectively, for a broad range of identifying assumptions. These estimators are derived from Bayesian statistical decision theory, they respect the dynamics of the impulse response functions, they are optimal in the relevant sense, and they are easy to

implement. We show that, in contrast, the posterior median response functions and mean response functions used in many applied studies do not respect the dynamics of the impulse responses or their dependence across response functions. In particular, in the highly empirically relevant case when the number of estimated structural impulse responses ($nirf$) exceeds the number of estimated structural model parameters (np), these estimators need not coincide with any posterior draw for the structural impulse response function, even when the number of posterior draws approaches infinity. As a result, median and mean response functions may distort the shape of impulse response functions (and their comovement), which is of foremost interest in applied work. This makes them unappealing for applied work. The intuitive reason is that, when $nirf > np$, these estimators violate the constraint that the minimized value of the expected loss must be a member of the set of structural impulse response functions that can be generated from the structural VAR model. When $nirf \leq np$, in contrast, the median and mean response functions will converge to our estimator under absolute and under quadratic loss, respectively, as the number of posterior draws approaches ∞. For any finite number of posterior draws, there may be differences.

Second, we show how to construct joint credible sets of the structural response functions under absolute and under quadratic loss. These credible sets differ conceptually and practically from conventional impulse response error bands based on vectors of quantiles of the marginal impulse response distributions. They not only account for the full estimation uncertainty about the impulse responses, but convey the directions in which the estimates are likely to depart from the baseline estimate. They also are guaranteed to generate solutions contained in the set of impulse response functions that can be generated by an autoregressive model. In contrast, conventional pointwise error bands are unsuitable for joint inference.

Third, we discuss how to construct estimators of the impulse responses and joint credible sets under a Dirac delta loss function. This approach was originally proposed by Inoue and Kilian (2013,

3

2019). That study showed how to rank posterior draws for structural VAR models based on the joint posterior density of the structural impulse responses when $nirf = np$. Closed-form solutions for this density can be derived analytically, facilitating the implementation of this approach. Inoue and Kilian (2013) did not consider the case of $nirf < np$ or of $nirf > np$. In this paper, we generalize this earlier approach to allow for arbitrary combinations of $nirf$ and np. We also extend the original analysis to allow for nonrecursive models with short-run exclusion restrictions, long-run exclusion restrictions, or combinations of sign and exclusion restrictions, as discussed in Arias, Rubio-Ramirez and Waggoner (2018). Especially, the latter situation is increasingly important in applied work for modeling the transmission of global shocks to domestic economies and for modeling economies with frictions (e.g., Mumtaz and Surico 2009; Mountford and Uhlig 2009; Aastveit, Bjørnland and Thorsrud 2015; Beaudry, Nam and Wang 2018; Kilian and Zhou 2020a,b). In contrast, the original analysis in Inoue and Kilian (2013) only allowed for recursively identified models and models identified by sign restrictions. Our analysis also allows for narrative restrictions, as discussed in Antolin-Diaz and Rubio-Ramirez (2018). Such inequality restrictions are increasingly used in sign-identified VAR models (e.g., Kilian and Zhou 2020a,b). We achieve this added generality by exploiting a new approach to deriving a closed-form solution for the joint posterior density function of the structural impulse responses that builds on results in Arias, Rubio-Ramirez and Waggoner (2018).

The remainder of the paper is organized as follows. Section 2 highlights our two main contributions. In section 3, we discuss joint inference under absolute and quadratic loss In section 4, we discuss joint inference under the Dirac delta loss function. Section 5 contains two empirical illustrations. The concluding remarks are in section 6.

2 Overview

The Bayesian estimate of a parameter is defined as the parameter value that minimizes in expectation the loss function of the user. In the interest of clarity, we begin by acknowledging that the choice of the loss function for impulse respone analysis is subjective and that estimators derived under one loss function need not be optimal under another. Our point in this paper is not to argue for one loss function over another. In fact, we present results for three commonly used alternative loss functions (absolute, quadratic and Dirac delta loss). We leave it to the applied user to choose among these loss functions.

Instead, our first main point is that, regardless of the loss function, it is important to restrict estimates of impulse response functions to lie in the relevant parameter space spanned by the underlying VAR model, when evaluating the expected loss. This restriction is not imposed in the construction of conventional median or mean response functions, defined as vectors of pointwise posterior medians or means. We propose alternative Bayesian estimators of the impulse response function that satisfy this restriction. Restricting the space of impulse responses to response functions that can be generated by the underlying structural VAR model is crucial if we are interested in the comovement across impulse responses.

The key difference between our approach and conventional median and mean response functions can be summarized as follows. The impulse response function, $\theta = g(\lambda)$, is a vector-valued function that maps from the space $\Lambda \subset \Re^{np}$ of the underlying structural VAR parameters, λ, to the space $\Theta \subset \Re^{nirf}$ of the structural responses. Note that this function is surjective on the co-domain $\Theta = g(\Lambda)$, but not on its complement Θ^c.[1] For concreteness, suppose that the loss function is quadratic and denote the data by Y. Then the conventional optimal Bayes decision, $\tilde{\theta} = E(\theta|Y)$, known as

---

[1] A function g from a set Λ to a set Θ is surjective, if for every element θ in the codomain Θ of g, there is at least one element λ in the domain Λ of g such that $g(\lambda) = \theta$. It is not required that λ be unique; the function g may map one or more elements of Λ to the same element of Θ.

the mean response function, minimizes $E\left[\|\delta - \theta\|^2 | Y\right]$ among all $\delta \in \Re^{nirf}$, where the expectation is taken with respect to the posterior distribution of θ and $\|\cdot\|$ denotes the Euclidian norm. In contrast, the optimal Bayes decision proposed in this paper, denoted $\overline{\theta}$, minimizes $E\left[\|\delta - \theta\|^2 | Y\right]$ among all $\delta \in \Re^{nirf}$, for which there exists a λ in Λ such that $g(\lambda) = \delta$. Although the loss function in our analysis is the same as in the conventional approach, the constraint set is different. By construction, our optimal solution $\overline{\theta}$ always belongs to Θ, so there always exists a $\overline{\lambda} \in \Lambda$ such that $\overline{\theta} = g(\overline{\lambda})$. This is not the case for the conventional estimator. If $\tilde{\theta}$, belongs to Θ, there exists $\widetilde{\lambda} \in \Lambda$ such that $\tilde{\theta} = g(\widetilde{\lambda}) = E(\theta|Y)$ because g is surjective on Θ. In contrast, if $\tilde{\theta}$ belongs to Θ^c, $\overline{\theta}$ will differ from $\tilde{\theta} = E(\theta|Y)$. We demonstrate in section 3 that, when $nirf > np$, Θ is of measure zero in $\Re^{nirf}$ and hence, with probability one, there is no $\widetilde{\lambda} \in \Lambda$ such that $\tilde{\theta} = g(\widetilde{\lambda})$. Thus, the two optimal Bayes estimators differ in general. Analogous arguments apply to the median response function under an additively separable absolute loss function.

The reason for preferring $\overline{\theta}$ over $\tilde{\theta}$ is that only $\overline{\theta}$ respects the dynamics of θ implied by the underlying structural model, as shown in section 3. We document that median and mean response functions can be quite misleading because they may distort the shape of impulse response functions and their comovement. Because most applied users are interested in the dynamics of the impulse response function, we make the case that restricting the space of impulse response functions is essential for applied work. This is not to deny that conventional median or mean response functions, remain useful summary statistics for studying individual impulse responses in isolation, but there are few examples of economically interesting questions that can be answered based on the marginal distribution of impulse responses.

Our second main point is that conventional Bayesian error bands, defined as vectors of upper and lower quantiles of the marginal impulse response posterior distributions misrepresent the esti- mation uncertainty about estimates of impulse response functions. This point is neither new nor

controversial. It is well documented in the Bayesian literature, and it applies to both point identified and set-identified models (e.g., Sims and Zha 1999; Inoue and Kilian 2013; Montiel Olea and Plagborg-Møller 2019). Analogous concerns have been raised in the growing frequentist literature on simultaneous confidence regions for impulse responses (e.g., Lütkepohl, Staszewska-Bystrova and Winker 2015a,b, 2018; Inoue and Kilian 2016; Bruder and Wolf 2018; Montiel Olea and Plagborg-Møller 2019). The only reason that the use of pointwise Bayesian error bands has persisted in applied work has been the lack of an easy to implement alternative that can be applied to a wide range of structural VAR models. Our contribution to this literature is to propose an alternative approach to the construction of joint credible sets that can be adapted to any of the three loss functions of interest. These credible sets provide a constructive alternative to the use of pointwise quantile error bands in Bayesian inference.[2]

3 Joint inference under absolute and under quadratic loss

It is useful to first discuss joint impulse response inference under absolute and under quadratic loss, before turning to the Dirac delta loss function. The techniques discussed in this section may be applied to impulse responses in scalar autoregressions as well as in structural vector autoregressions. Let θ denote the vector of unknown impulse responses obtained by appropriately stacking all n^2 impulse response functions for horizon $h = 0, 1, 2, ..., H$, in the n-dimensional autoregressive model into one vector of dimension $n^2(H + 1)$. $\theta \in \Theta$, where Θ is the set of all structural impulse response functions that satisfy the identifying restrictions imposed on the structural VAR model. We abstract from the details of the construction of θ, which may differ from one structural VAR model to another, because our analysis does not depend on these details (see Kilian and Lütkepohl

---

[2]In closely related work, Montiel Olea and Plagborg-Møller (2019) propose a simultaneous credible band for impulse response functions based on the product of component-wise equal-tailed credible intervals, where the tail probability has been calibrated to yield simultaneous credibility $1 - \alpha$.

2017). The space Θ may be approximated by drawing from the posterior of the model parameters and simulating the distribution of the structural impulse responses that are consistent with the identifying assumptions. We first derive Bayesian estimators of this impulse response vector under absolute and under quadratic loss. We then show how to construct joint credible sets for this estimator under the same loss function.

3.1   Constructing the impulse response estimator

Let $\overline{\theta}$ denote a summary statistic of the posterior distribution of θ. $L(\theta, \overline{\theta})$ is a loss function that maps $\Theta \times \Theta$ to $\Re$ such that

$$\overline{\theta} \;=\; \mathrm{argmin}_{\overline{\theta} \in \Theta} E[L(\theta, \overline{\theta})], \tag{1}$$

where it is assumed that there exists a unique minimum and the expectation is with respect to the joint posterior distribution of θ over Θ. Our analysis in this section focuses on the loss functions

$$L(\theta, \overline{\theta}) \;=\; \begin{cases} \sum_{j=1}^{k} (\theta_j - \overline{\theta}_j)^2 & \text{(quadratic loss)} \\[2mm] \sum_{j=1}^{k} |\theta_j - \overline{\theta}_j| & \text{(absolute loss)} \end{cases}, \tag{2}$$

where θ_j and $\overline{\theta}_j$ are the jth elements of θ and $\overline{\theta}$, respectively. Given the additive separability of these loss functions, it suffices to evaluate the expectation under each of the marginal posterior distributions. As is well known, the quadratic loss function for a scalar random variable is minimized in expectation by the mean. Moreover, under additive separability, the absolute loss function for a scalar random variable is minimized in expectation by the median.

In practice there is no closed-form solution to (1), so we need to approximate the solution by

simulation. The estimators are given by

$$\widehat{\overline{\theta}}_M = \text{argmin}_{\overline{\theta} \in \widehat{\Theta}_M} \frac{1}{M} \sum_{i=1}^{M} L(\theta^{(i)}, \overline{\theta}), \tag{3}$$

where $\theta^{(i)}$ is the ith posterior draw and $\widehat{\Theta}_M$ consists of M posterior draws, possibly after reweighting the draws, as discussed in Arias et al. (2018) and Antolin-Diaz and Rubio-Ramirez (2018). Specifically, under quadratic and under absolute loss, respectively,

$$\widehat{\overline{\theta}}_M^{mean} = \text{argmin}_{\overline{\theta} \in \widehat{\Theta}_M} \frac{1}{M} \sum_{i=1}^{M} \sum_{j=1}^{k} (\theta_j^{(i)} - \overline{\theta}_j)^2, \tag{4}$$

$$\widehat{\overline{\theta}}_M^{median} = \text{argmin}_{\overline{\theta} \in \widehat{\Theta}_M} \frac{1}{M} \sum_{i=1}^{M} \sum_{j=1}^{k} |\theta_j^{(i)} - \overline{\theta}_j|, \tag{5}$$

where $\theta_j^{(i)}$ denotes the jth element of the ith draw of θ, denoted $\theta^{(i)}$, from its posterior distribution. Unlike conventional estimators, these mean and median estimators respect the dynamics of the impulse response function θ and restrict attention to feasible model solutions.

3.2 Comparison with conventional posterior median (or mean) impulse response functions

The key difference from the conventional median or mean response functions in the literature is that under absolute loss, the solution $\widehat{\overline{\theta}}_M$ is always contained in the set of structural impulse responses that can be obtained by drawing from the posterior of θ. This is true for any combination of p and H. In contrast, in the empirically relevant case when the number of estimated structural responses ($nirf$) exceeds the number of estimated structural model parameters (np), the posterior median response function will not necessarily be contained in Θ, even as the number of posterior draws approaches infinity. The reason is that in this case the posterior distribution of the structural impulse response functions is defined over a lower-dimensional manifold. Our analysis recognizes

9

that the space of θ may be a lower dimensional smooth manifold, whereas the conventional approach postulates that it is the Euclidian space. This fact does not change the posterior distribution for the structural model parameters, but it affects optimal impulse response inference based on the posterior distribution of the structural parameters under any loss function, including absolute loss and quadratic loss.

Even when $nirf \leq np$, the vector of marginal posterior medians will coincide with $\widehat{\overline{\theta}}_M$, only as the number of posterior draws approaches ∞. When $nirf < np$, the vector of pointwise medians obtained from the marginal posterior distributions does not solve expression (1) even in the limit. Analogous results hold for the mean response function, defined as the vector of pointwise means from the marginal posterior impulse response distributions. This point is illustrated by the following stylized example.

Consider estimating the scalar-valued Gaussian AR(1) model

$$y_t \quad = \quad \rho y_{t-1} + u_t, \tag{6}$$

$$u_t \quad \overset{iid}{\sim} \quad N(0,1), \tag{7}$$

from the observations $y_0, y_1, ..., y_T$ with a diffuse conjugate Gaussian prior with prior mean ρ_{mean}. Then the posterior distribution of ρ is

$$\rho \quad \sim \quad N\left(\rho_{mean}, \frac{1}{n_T}\right). \tag{8}$$

where $n_T = \sum_{t=1}^{T} y_{t-1}^2$.

3.2.1 The Conventional Approach

For expository purposes, we consider the reduced-form impulse response function and, with a slight abuse of notation, denote this impulse response vector by θ. Focusing on the reduced-form responses allows us to isolate the essence of the problem with the conventional approach, which are unrelated to the uncertainty about the error variance. Since the response to a unit shock at horizon 0 is known to be 1, we focus on higher horizons. Suppose that we are interested in joint inference at horizons 1 and 2 only. Then $\theta = [\rho, \ \rho^2]$. In the context of this example, the condition that $nirf > np$ in a structural VAR model translates to the number of estimated reduced-form responses, which is 2 in this example, exceeding the lag order of 1.

We consider first the posterior mean response function and then the posterior median response function. The posterior mean of $[\rho, \ \rho^2]$ is

$$\left[\rho_{mean}, \ \rho_{mean}^2 + \frac{1}{n_T} \right] \tag{9}$$

Note that there is no possible value of ρ in the AR(1) model that produces the posterior mean impulse response function because $1/n_T > 0$.

Next, suppose that the posterior median of ρ, denoted as ρ_{median}, equals ρ. Because

$$
\begin{aligned}
P(\rho^2 \leq x) &= P(-\sqrt{x} \leq \rho \leq \sqrt{x}) \\
&= \Phi\left(\sqrt{n_T}(\sqrt{x} - \rho_{median})\right) - \Phi\left(\left(\sqrt{n_T}(-\sqrt{x} - \rho_{median})\right)\right),
\end{aligned}
\tag{10}
$$

it is always true that

$$P(\rho^2 \leq \rho_{median}^2) < \frac{1}{2}. \tag{11}$$

In other words, ρ_{median}^2 is not the posterior median of ρ^2. Thus, there is no value of ρ in the

AR(1) model such that the posterior median response function matches the actual impulse response function. In contrast, the solution $\widehat{\overline{\theta}}_M$ satisfies this constraint by construction.

Note that these results do not depend on the choice of the prior or the specification of the Bayesian model. The reason that the conventional mean and median response functions in this example fail to recover $\widehat{\overline{\theta}}_M$ is that they do not respect the dynamics of the impulse response function. Our analysis calls into question the conventional claim that the posterior median response function is necessarily optimal under the loss function $\sum_{i=1}^{k} |\theta_i - \overline{\theta}_i|$. The reason that we reach a different conclusion is not that we changed the loss function. Rather the difference is about the range over which to integrate when computing the expected loss. Whereas we restrict ourselves to the set of impulse response functions contained in Θ, the conventional approach considers in addition infeasible solutions not contained in Θ.[3]

3.2.2 Mean and median response functions distort the dynamics of impulse response functions

It may seem that it is a matter of taste whether to restrict the minimization problem for θ to elements of Θ or not. We would argue that it is not. The reason why we do not find the conventional formulation of the minimization problem underlying the Bayesian estimator of θ compelling is that the construction of the statistics reported in econometrics should reflect the preferences of applied users. For example, it is widely recognized there is no point in econometricians exploring loss functions that are at odds with the objectives of applied users. Likewise, there is no point in defining the space of possible solutions in impulse response analysis in a way that conflicts with the objectives of applied users.

In many studies it is the shape of the response function that users of structural VAR models

---

[3]It may be tempting to argue that the conventional solution for θ is preferred because in this AR(1) example it generates the same expected loss at $h = 1$ and lower expected loss at $h = 2$. This argument would be misleading because the conventional solution only reduces the expected loss by violating the constraint that solutions for θ must be feasible.

are most interested in. For example, macroeconomists are often interested in whether the response function for real output to demand shocks is hump shaped or not (e.g., Cochrane 1994). In fact, many macroeconomists have abandoned frictionless neoclassical models and adopted models with nominal or real rigidities based on VAR evidence of sluggish or delayed responses of inflation and output (e.g., Woodford 2003). Similarly, users of structural VAR models in international economics tend to be interested in whether there is delayed overshooting in the response of the exchange rate to monetary policy shocks (e.g., Eichenbaum and Evans 1995).

Moreover, many empirical questions can only be answered by considering several impulse response functions in conjunction. For example, the question of whether a global oil price shock creates stagflation in the domestic economy by necessity involves studying the responses of inflation as well as real output. Likewise, researchers interested in the effects of a U.S. monetary policy shock care about the responses of both real output and inflation because the loss function of the Federal Reserve depends on both real output and inflation, making it necessary to consider these response functions jointly. Another good example is Blanchard (1989). This study uses a macroeconomic VAR model to evaluate the implication of standard Keynesian models that (1) positive demand innovations increase output and decrease unemployment persistently, and (2) that a favorable supply shock triggers an increase in unemployment without a decrease in output. Similarly, Sims and Zha (1999) stress that it is often thought that plausible patterns of the response of the economy to a disturbance in monetary policy should have interest rates rising, the money stock falling, output falling and prices falling.

Thus, it is fair to conclude that most applied researchers are interested in the shape of impulse response functions and in the comovement across impulse response functions. This fact is important because the rationale for median and mean response functions is based on the implicit premise that applied users are not interested in the shape of a given impulse response function (or in the

comovement across impulse response functions). It is this premise that motivates their approach of minimizing the expected loss of the response function θ over the Euclidian space rather than over Θ.

Our concern with this approach is that it is not possible to answer questions about the shape of a given impulse response function based on statistics that do not respect the dynamics of the impulse response function. Likewise, it is not possible to evaluate the comovement across the n^2 response functions contained in θ based on statistics that do not respect that comovement. As a result, computing median or mean response functions may distort the shape of the response function, which is of foremost interest in applied work, and more generally the comovement across response functions in multivariate models. This point is best illustrated by example.

Consider the same AR(1) process that served as our example earlier. The posterior distribution of the slope parameter ρ is $N\left(\rho_{mean}, \frac{1}{n_T}\right)$. The impulse response horizons of interest are $h = 1, 2, 3, 4$. Then the posterior mean of $\theta = \left[\rho, \rho^2, \rho^3, \rho^4\right]$ is given by

$$\left[\rho_{mean}, \rho_{mean}^2 + 1/n_T, \rho_{mean}^3 + 3\rho_{mean}/n_T^2, \rho_{mean}^4 + 6\rho_{mean}^2/n_T + 3/n_T^2\right],$$

where we make use of the fact that $\rho = \rho_{mean} + z/\sqrt{n_T}$ with $z \sim N(0,1)$. It is easy to construct examples in which the shape of this mean response function differs from the correct shape of the response function implied by the AR(1) model. For example, let $\rho_{mean} = 0.7$ and $n_T = 5$. Then the mean impulse response function is $[0.70, 0.69, 0.76, 0.95]$. Even though the posterior mean of the root of the AR(1) process is stationary, which implies that the response function at the mean value must be smoothly decaying, the mean response function is increasing with the horizon. Similar problems arise with the median response function. Figure 1 shows an example, in which the median response function follows an irregular pattern that is clearly inconsistent with the properties of the AR(1) model. A user of the median response function would incorrectly conclude that the response

14

function is not smoothly decaying over time. The optimal estimator under absolute loss derived in this paper, in contrast, exhibits the expected smooth decline. These examples show that mean and median response functions cannot be trusted to capture the shape of the impulse response functions across horizons (or the comovement of responses across variables in Θ).

Since there is overwhelming evidence that most applied users do care about the shape of impulse response functions and the comovement across impulse response functions, the posterior mean response function and the posterior median response function clearly are not valid summary statistics of the impulse response dynamics, whereas the alternative statistics derived in this section are. We conclude that median response functions and mean response functions cannot be recommended for applied work.

3.2.3   The role of weights in the loss function

It should be noted that it is possible to further weight individual elements of θ in the loss function, depending on the variable of interest and the horizon of the response function. For example, a user may want to assign more weight to impulse responses at horizons of interest to policymakers. We do not consider this possibility here, since this option has not been considered in applied work. We do note, however, that the conventional median response function solves the loss minimization problem under additively separable absolute loss regardless of the weights applied to the elements of θ, whereas our approach will generate different solutions under absolute loss as well as quadratic loss, as the weights are varied.

3.3   Extensions to partially identified models

Although so far we have focused on fully identified structural VAR models, our analysis can be easily adapted to structural VAR models, in which only a subset of the structural shocks is identified.

The only change is the dimensionality of the θ vector when evaluating the expected loss. For example, when only one structural shock is identified, as is often the case in structural VAR models of monetary policy shocks, θ is of length $n(H+1)$. Otherwise, the implementation of our methods is unchanged.

Like in the fully identified model, the median response function may distort the dynamics of the impulse response function when there are more structural impulse responses to be estimated than structural model parameters. This situation is common in applied work based on partially identified structural VAR models. For example, Eichenbaum and Evans (1995) specify a recursive semi-structural model of monetary policy shocks with $n = 5$ variables, $p = 6$ autoregressive lags, and $H = 35$. They are recovering $nirf = nH + 1 = 176$ structural impulse responses from $np = 1 + pn^2 = 151$ structural parameter estimates, so the number of structural impulse responses in their model exceeds the number of structural model parameters. Similarly, Uhlig (2005) specifies a sign-identified structural VAR model of monetary policy shocks with $n = 6$, $p = 6$ and $H = 60$. In that study, there are $nirf = n(H+1) = 366$ structural impulse responses, which exceeds the number of structural model parameters, $np = n^2(p+1) = 252$. It should be noted that the methods proposed in this section may be applied regardless of whether $nirf > np$ or $nirf \leq np$. There is no reason for an applied user to be concerned about whether $nirf > np$. This question only matters for the valdity of conventional mean and median response functions.

3.4 Joint credible sets

In response to the concerns we reviewed in the introduction, there has been growing interest in recent years in quantifying the joint uncertainty of vectors of VAR impulse response estimates. There is no consensus in the Bayesian literature on how to construct such credible sets. For example, Berger (1985, p. 145) does not view credible sets as having a clear decision-theoretic

16

role at all, and therefore is leery of 'optimality' approaches to the selection of a credible set. He views credible sets as an easily reportable crude summary of the posterior distribution. In this section, we show how to construct joint credible sets for $\widehat{\overline{\theta}}_M$ under quadratic and under absolute loss.[4] Although our approach is optimal under these loss function, it may be also be viewed as a pragmatic approach to conveying the posterior uncertainty about θ in the spirit of Berger's remark. It should be noted that joint credible sets need not be constructed under the same loss function as the original estimate, but doing so ensures that $\widehat{\overline{\theta}}_M$ is contained within the credible set.

We define the $(1-\alpha)100\%$ joint credible set based on a loss function $L(\theta, \overline{\theta})$ by

$$\Theta_{1-\alpha,L} = \{\theta \in \Theta : L(\theta, \overline{\theta}) \leq c_{1-\alpha,L}\}, \tag{12}$$

where $c_{1-\alpha,L}$ is the smallest number such that the posterior probability of $\Theta_{1-\alpha,L}$ is $1-\alpha$ and L refers to the loss function. In practice, this credible set is estimated as

$$\widehat{\Theta}_{1-\alpha,M,quadratic} = \left\{\theta \in \widehat{\Theta}_M : \sum_{i=1}^{k}(\theta_i - \widehat{\overline{\theta}}_{M,i}^{mean})^2 \leq c_{1-\alpha,quadratic}\right\}, \tag{13}$$

$$\widehat{\Theta}_{1-\alpha,M,absolute} = \left\{\theta \in \widehat{\Theta}_M : \sum_{i=1}^{k}|\theta_i - \widehat{\overline{\theta}}_{M,i}^{median}| \leq c_{1-\alpha,absolute}\right\}, \tag{14}$$

where $\widehat{\overline{\theta}}_M^{mean}$ and $\widehat{\overline{\theta}}_M^{median}$ denote the respective solutions to (3) for the loss functions in (2), and where $c_{1-\alpha,quadratic}$ and $c_{1-\alpha,absolute}$ are the smallest values such that $\widehat{\Theta}_{1-\alpha,M,quadratic}$ and $\widehat{\Theta}_{1-\alpha,M,absolute}$, respectively, have posterior probability $1-\alpha$. Unlike the conventional credible sets used in the literature, our joint credible sets are designed to contain only members of the set of structural response functions. In practice, these joint credible set may be computed by sorting in ascending

---

[4]For further discussion see Berger (1985, subsection 4.3.3) and Bernardo and Smith (1994, subsection 5.1.5), for example.

order the M posterior draws of the structural response function by the value of $\sum_{i=1}^{M} L(\theta^{(i)}, \overline{\theta})/M$ and retaining the first $(1-\alpha)100\%$ draws, starting with the draw with the lowest value. Although we focus on impulse response analysis, it should be noted that the same techniques may also be adapted to characterize the uncertainty about the future time path of an economic variable in a forecasting application.

3.5   Comparison with conventional error bands

The conventional approach in the literature has been to report posterior median (or mean) response functions along with pointwise error bands constructed from the upper and lower quantiles of the marginal posterior distributions of the structural impulse responses. While this approach makes sense when conducting inference on individual impulse responses, it is readily apparent that this approach is misguided if we are interested in inference on θ.

Constructing response functions and error bands by stringing together quantiles of the marginal posterior distributions of individual impulse responses, is akin to arguing that in studying a system of regression equations it is sufficient to base inference on t-test statistics for each coefficient in a given equation, ignoring that the t-statistics are dependent within and across equations. It is immediately obvious that this approach fails to capture the true uncertainty about the vector θ. The only defense of this approach may be that there has been no constructive alternative to the use of such pointwise error bands to date. As we have shown in this section, there are easy-to-implement alternative methods that avoid this drawback and that are valid under absolute and under quadratic loss, respectively, rendering the use of pointwise error bands obsolete.

As noted in the introduction, there are three distinct concerns with the use of conventional pointwise error bands. First, vectors of upper and lower quantiles derived from the marginal impulse response posterior distributions, like vectors of medians, need not be contained in the set

Θ of feasible solutions and may distort the shape of the response functions. Second, as pointed out by Sims and Zha (1999), two-dimensional error bands are consistent with any shape of the response functions that fits within the error band. They do not tell us anything about likely departures from the baseline estimate of the impulse response function. Third, the uncertainty about the impulse response function θ clearly depends on the covariance across individual impulse response estimates, which cannot be captured without considering the joint distribution of the impulse responses.

Our approach to inference differs sharply from the conventional approach of reporting vectors of quantiles of marginal impulse response posterior distributions. By restricting all members of the credible set for θ to lie in the set Θ, we not only ensure that each member of that set is feasible, but we also make sure that $\overline{\theta}$ corresponds to a draw from the joint distribution of the impulse responses that satisfies the covariance structure of the impulse responses. In addition, even when posterior median (or mean) response functions are numerically close to $\widehat{\overline{\theta}}_M$, these joint credible sets will not resemble conventional pointwise error bands. Rather, they will look like shot-gun trajectories employed in ballistics, with each structural model represented by a set of n^2 such trajectories up to horizon H. This feature addresses the concern raised by Sims and Zha (1999) that we need to be able to visualize likely departures from the baseline estimate of θ.[5] Similar techniques have been employed in Inoue and Kilian (2013, 2016) and a number of applied studies to characterize the joint estimation uncertainty in structural VAR models.

Of course, in practice, M is typically so large that we cannot distinguish individual response functions or, for that matter, responses from different structural models. The shotgun trajectory plot does, however, contain the information required to make these assessments. If we are interested in evidence that a positive oil price shocks causes stagflation in the U.S. economy, for example, we may display the response functions that satisfy the definition of a stagflationary response in

---

[5]These trajectories are also conceptually similar to the forecast paths reported in weather forecasting, when outlining the predicted path of a hurricane.

a different color than response functions that do not, as illustrated in Inoue and Kilian (2016). Similarly, if we are interested in whether the most likely models satisfy the requirement that a monetary policy shock raises the interest rate, lowers the price level and lowers real output, as discussed by Sims and Zha (1999), we can plot the responses for this subset of structural models in a different color. We can also make probability statements about the conditional comovement of several response functions contained in the joint credible set. Finally, if we are interested in the shape of a given impulse response function, we can make probability statements about whether the response is hump-shaped or whether there is delayed overshooting, for example. Further discussion of techniques for extracting such information from shot-gun trajectory plots can be found in Inoue and Kilian (2016).[6]

It may be tempting to construct two-dimensional impulse response error bands from these joint credible sets. For example, one might wish to bound the joint credible set by constructing the upper and lower envelope around the response estimates in the credible region. Such error bands may indeed help summarize the range of uncertainty about the impulse responses, but it is important to understand that they would not be a good substitute for the plot of the trajectories. First, these error bands, much like the conventional pointwise error bands, need not be contained in the set Θ of feasible solutions. Second, when $nirf > np$, one cannot necessarily rule out that the credible region may be disjoint, even as the number of posterior draws approaches ∞. Third, by focusing on error bands rather than the shot-gun trajectory plots, we suppress the information that applied users are most interested in. Consider the example of a structural VAR model of monetary policy shocks. For starters, we lose the information required to judge whether the response of output to

---

[6]One may object that if we care about features such as hump-shaped responses, we could instead use a loss function that reflects this objective. This would allow the user to directly evaluate the posterior support for a hump shape. In practice, however, one cannot target all possible loss functions a potential user may have in mind. It therefore makes sense for researchers to settle on default conventions for reporting impulse response estimates that allow users with different loss functions to extract as much information as possible from the impulse response plot. This is indeed the approach typically taken both in the Bayesian and in the frequentist structural VAR literature (e.g., Montiel Olea, Luis and Plagborg-Møller 2019).

a monetary policy shock is hump-shaped. Next, as stressed by Sims and Zha (1999), when looking at the response of output to a monetary policy shock, we need to know whether a stronger output decline within the credible set remains consistent with the implications of a monetary policy shock for the other model variables. This question cannot be answered based on error bands. It can only be answered based on the information contained in the shot-gun trajectory plots.

4 Joint inference under Dirac delta loss

Although the traditional approach in applied studies has been to focus on absolute or quadratic loss functions, respectively, starting with Inoue and Kilian (2013, 2019) there has been growing interest in the application of the Dirac delta loss function to the joint distribution of the structural impulse responses. Given an additively separable loss function, as discussed in section 3.1, it would be straightforward in principle to evaluate expression (1) under Dirac delta loss. There are two concerns. One concern is that additive separability is implied by the use of a quadratic loss function, but not absolute loss or Dirac delta loss. Hence, it is not clear why we would want to restrict the loss function to be additively separable under Dirac delta loss. The other concern is that, unlike under quadratic or absolute loss, we need to know the marginal posterior distributions of each impulse response coefficient in order to be able to determine the modes for each element of θ. Deriving these marginal posterior distributions would require us not only to derive the joint posterior distribution, but to use Monte Carlo integration methods to derive the marginal distribution of each element of θ. The computational cost of this last step is prohibitive in practice (see Inoue and Kilian 2013).

This is why in this section we take the less restrictive and computationally simpler approach of evaluating the mode of the joint posterior distribution of θ. We assume for now that the structural VAR model is fully identified. Let θ denote the $n^2(H+1) \times 1$ vector of structural impulse response functions that satisfy the identifying restrictions, where $\theta \in \Theta$. As before, Θ is a lower-dimensional

smooth manifold when the number of structural impulse responses to be estimated exceeds the number of structural model parameters used to construct these responses. We are interested in the estimator $\overline{\theta}$ of θ under the Dirac delta loss function, $L(\theta, \overline{\theta}) = -\delta(\overline{\theta} - \theta)$, where $\delta(\cdot)$ denotes the Dirac delta function. Minimizing the expected loss is equivalent to maximizing the posterior density of θ, which occurs when $\overline{\theta}$ is the posterior mode of the joint posterior density.

4.1 The most likely impulse response estimate and joint HPD credible sets

The use of the mode has a long tradition in Bayesian inference, as does the use of highest posterior density (HPD) regions for impulse response inference (e.g., Koop 1996; Zha 1999; Waggoner and Zha 2012; Plagborg-Møller 2019). Similarly, Sims and Zha (1999) suggest that impulse response inference should focus on deviations from the baseline estimate that are "most likely". Focusing on the modal draw from the joint posterior distribution of θ and on joint HPD sets, as proposed by Inoue and Kilian (2013, 2019) thus is quite natural.[7]

The procedure proposed in Inoue and Kilian (2013, 2019) for finding the modal posterior draw is simple. We first rank the M posterior draws for θ that satisfy the identifying restrictions of the structural VAR model based on the value of the joint posterior density function $f(\theta|y_1, ..., y_T)$. Then the modal (or most likely) posterior draw for θ, denoted $\widehat{\overline{\theta}}_M$, will be the draw with the highest joint posterior value. By construction, the mode of the joint posterior distribution will be within Θ. In practice, the responses of the modal model are obtained by sorting the M posterior draws for θ in ascending order based on the value of the joint posterior.

The corresponding $1 - \alpha$ probability joint credible sets may be constructed by retaining the $(1 - \alpha)100\%$ draws for θ with the highest joint posterior density values among the M draws that satisfy the identifying restrictions. The most likely impulse response estimate thus may be viewed

---

[7]This approach has been applied in a number of recent studies including Herwartz and Plödt (2016), Herrera and Rangaraju (2020), Zhou (2020), and Cross, Nguyen and Tran (2019).

as the limit of a joint $1 - \alpha$ probability highest-posterior density (HPD) credible set. Like in section 3, the plot of this joint credible set resembles a shot-gun trajectory plot and may be evaluated in the same way.

The main challenge, in practice, is how to compute the values of the joint posterior density of the structural impulse responses responses for each posterior draw. In our earlier work, we analytically derived a closed-form solution for the function $f(\cdot)$ that makes this approach quite computationally efficient. In this section, we propose an alternative derivation of this density that is substantially more general than the expressions derived in Inoue and Kilian (2013, 2019) under the assumption that $nirf = np$.

4.2 The joint posterior density of the structural impulse responses

Following Arias et al. (2018), consider a pth-order structural VAR model

$$
\begin{aligned}
y_t' A_0 &= \sum_{j=1}^{p} y_{t-j}' A_j + c + \varepsilon_t' \\
&= x_t' A_+ + \varepsilon_t', \quad \text{for } t = 1, 2, ..., T,
\end{aligned} \tag{15}
$$

where y_t is an $n \times 1$ vector of observed variables, A_j is an $n \times n$ matrix of parameters for $j = 0, 1, ..., p$, c is a $1 \times n$ vector of intercepts, ε_t is an $n \times 1$ vector of structural shocks, $x_t' = [y_{t-1}' \ y_{t-2}' \ \cdots \ y_{t-p}' \ 1]$ and $A_+' = [A_1' \ \cdots \ A_p' \ c']$. A_0 is assumed to be nonsingular.The further analysis differs depending on the nature of the identifying restrictions. We first discuss the case of structural VAR models that are set-identified by sign restrictions (possibly in conjunction with exclusion restrictions and narrative inquality restrictions). We then consider the case of structural VAR models exactly identified by any combination of short-run and long-run exclusion restrictions.

4.2.1 Models identified by sign restrictions

The key technical difference from Inoue and Kilian (2013, 2019) is that in the current paper we start with the posterior of A_0 and A_+, as defined in Arias et al. (2018), whereas in our earlier work we applied the change-of-variable method to the joint posterior of the reduced-form slope parameters, the error covariance matrix and the rotation matrix. Given a uniform-normal-inverse-Wishart family of prior densities with parameters $\bar{\nu}$, $\bar{\Phi}$, $\bar{\Psi}$ and $\bar{\Omega}$, the posterior density of A_0 and A_+ is given by

$$NGN_{(\tilde{\nu},\tilde{\Phi},\tilde{\Psi},\tilde{\Omega})} \propto |\det(A_0)|^{\tilde{\nu}-n}e^{-\frac{1}{2}\text{vec}(A_0)'(I_n\otimes\Phi)\text{vec}(A_0)} \times e^{-\frac{1}{2}\text{vec}(A_+-\tilde{\Psi}A_0)'(I_n\otimes\Omega)^{-1}\text{vec}(A_+-\tilde{\Psi}A_0)},$$
(16)

where $\tilde{\nu} = T + \bar{\nu}$, $\tilde{\Omega} = (X'X + \bar{\Omega})^{-1}$, $\tilde{\Psi} = \tilde{\Omega}(X'Y + \bar{\Omega}^{-1}\bar{\Psi})$, $\tilde{\Phi} = Y'Y + \bar{\Phi} + \bar{\Psi}'\bar{\Omega}^{-1}\bar{\Psi} - \tilde{\Psi}'\tilde{\Omega}^{-1}\tilde{\Psi}$, $Y = [y_1 \cdots y_T]$ and $X = [x_1 \cdots x_T]'$ (see equation 2.8 in Arias et al. 2018).[8]

Let Θ_h denote the structural impulse response matrix at horizon h implied by (15). Let $J_{H,p}$ denote the $(H+1)n^2 \times n^2(p+1)$ Jacobian matrix of $\theta = \text{vec}([\Theta_0 \ \Theta_1 \ \cdots \ \Theta_H])$ with respect to $\text{vec}([A_0 \ A_1 \ \cdots \ A_p])$. The following proposition, which is a corollary of the results in Arias et al. (2018), states the joint posterior density of θ (up to scale).

Proposition 1:

$$f(\theta|y_1,...,y_T) = NGN_{(\tilde{\nu},\tilde{\Phi},\tilde{\Psi},\tilde{\Omega})}|J'_{H,p}J_{H,p}|^{\frac{1}{2}}.$$
(17)

If the number of estimated impulse responses ($nirf$) is no larger than the number of estimated structural parameters used in computing these structural responses, $f(\theta|y_1,...,y_T)$ simplifies to

---

[8]The approach of specifying a marginally uniform prior on the rotation matrix has been standard in the literature on sign-identified structural VAR models (see, e.g., Rubio-Ramirez et al. 2010; Arias et al. 2018; Antolin-Diaz and Rubio-Ramirez 2018). As is well known, this prior may be inadvertently informative about the structural impulse responses, although the practical importance of this problem remains unclear. For an alternative approach with its own advantages and disadvantages see Plagborg-Møller (2019).

$NGN_{(\tilde{\nu},\tilde{\Phi},\tilde{\Psi},\tilde{\Omega})}|J_{H,p}|.$

Proof: This result follows from Arias et al.'s (2018) propositions 1 and 2 and their equation 2.8. To make our procedure operational, what remains to be derived is the Jacobian $J_{H,p}$, which can be written as the product of three Jacobian matrices.

The first Jacobian matrix is that of $\theta = \text{vec}([\Theta_0 \ \cdots \ \Theta_H])$ with respect to $\text{vec}([A_0'^{-1} \ \Phi_1 \ \cdots \ \Phi_p])$, where Φ_h denotes the reduced-form impulse response matrix at horizon h. Because $\Theta_h = \Phi_h A_0'^{-1}$, this Jacobian is

$$J_1 = \begin{bmatrix} I_{n^2} & 0_{n^2 \times n^2} & \cdots & 0_{n^2 \times n^2} \\ (I_n \otimes \Phi_1) & A_0^{-1} \otimes I_n & \cdots & 0_{n^2 \times n^2} \\ \vdots & \vdots & \ddots & \vdots \\ (I_n \otimes \Phi_H) & 0_{n^2 \times n^2} & \cdots & A_0^{-1} \otimes I_n \end{bmatrix}, \tag{18}$$

where K_n is the $n \times n$ commutation matrix.

The second Jacobian matrix is that of $\text{vec}([A_0'^{-1} \ \Phi_1 \ \cdots \ \Phi_H])$ with respect to $\text{vec}([A_0 \ B_1 \ \cdots \ B_p])$, where $B_i = A_i A_0^{-1}$, $i = 1, ..., p$. It is given by the $n^2(H+1) \times n^2(H+1)$ matrix

$$J_2 = \begin{bmatrix} -(A_0^{-1} \otimes A_0'^{-1})K_n & 0_{n^2 \times n^2} & 0_{n^2 \times n^2} & \cdots & 0_{n^2 \times n^2} \\ 0_{n^2 \times n^2} & I_{n^2} & 0_{n^2 \times n^2} & \cdots & 0_{n^2 \times n^2} \\ 0_{n^2 \times n^2} & X & I_{n^2} & \cdots & 0_{n^2 \times n^2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0_{n^2 \times n^2} & X & X & \cdots & I_{n^2} \end{bmatrix}, \quad \text{if } nirf \le np, \tag{19}$$

and by the $n^2(H+1) \times n^2(p+1)$ matrix

$$J_2 = \begin{bmatrix} -(A_0^{-1} \otimes A_0'^{-1})K_n & 0_{n^2 \times n^2} & 0_{n^2 \times n^2} & \cdots & 0_{n^2 \times n^2} \\ 0_{n^2 \times n^2} & I_{n^2} & 0_{n^2 \times n^2} & \cdots & 0_{n^2 \times n^2} \\ 0_{n^2 \times n^2} & X & I_{n^2} & \cdots & 0_{n^2 \times n^2} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0_{n^2 \times n^2} & X & X & \cdots & I_{n^2} \\ 0_{n^2 \times n^2} & X & X & \cdots & X \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0_{n^2 \times n^2} & X & X & \cdots & X \end{bmatrix} \;, \quad \text{if } nirf > np, \qquad (20)$$

where X denotes $n^2 \times n^2$ matrices.

The third Jacobian matrix is that of $\mathrm{vec}([A_0 \; B_1 \; \cdots \; B_p])$ with respect to $\mathrm{vec}([A_0 \; A_1 \; \cdots \; A_p])$ and is given by

$$J_3 = \begin{bmatrix} I_{n^2} & 0_{n^2 \times n^2} & \cdots & 0_{n^2 \times n^2} \\ -(A_1 A_0^{-1} \otimes A_0'^{-1})K_n & (I_n \otimes A_0'^{-1})K_n & \cdots & 0_{n^2 \times n^2} \\ \vdots & \vdots & \ddots & \vdots \\ -(A_p A_0^{-1} \otimes A_0'^{-1})K_n & 0_{n^2 \times n^2} & \cdots & (I_n \otimes A_0'^{-1})K_n \end{bmatrix}. \qquad (21)$$

Thus, the Jacobian $J_{p,H}$ is the product of (18), (19)/(20) and (21).

Because the Jacobian matrix is of dimension $(H+1)n^2 \times (H+1)n^2$, when $nirf > np$, it becomes very large in typical macroeconomic applications. Although its determinant simplifies to $|A_0|^{-(H+1)n}$ when $nirf \leq np$, it does not when $nirf > np$. This poses a computational challenge. We work around this problem by using the LU decomposition and computing the sum of the log

of the diagonal elements for the log of the Jacobian determinant.[9]

4.2.2 Models exactly identified by short-run and/or long-run exclusion restrictions

Consider a normal-inverse-Wishart (NIW) family of prior for the reduced-form parameters. Suppose that there are $\frac{n(n-1)}{2}$ short-run and/or long-run exclusion restrictions of the form

$$C \begin{bmatrix} \text{vec}(A_0^{-1\prime}) \\ \\ \text{vec}(B(1)^{-1}A_0^{-1\prime}) \end{bmatrix} = 0_{\frac{n(n-1)}{2} \times 1}, \tag{22}$$

where C is a $\frac{n(n-1)}{2} \times 2n^2$ matrix of constants and $B(1) = I_n - B_1 - \cdots - B_p$ is assumed to be invertible. Typicaly, each of the rows of C consists of a 1 and $2n^2 - 1$ 0's. We assume that there is a unique A_0 that satisfies $A_0^{-1\prime}A_0^{-1} = \Sigma$ given (22). In this case the joint posterior density of θ can be derived directly from the posterior density of the reduced-form parameters, building on the results in Arias et al. (2018). The following proposition, which is not a special case of Proposition 1 because it does not allow for sign restrictions, states the joint posterior density of θ (up to scale).

Proposition 2:

$$f(\theta|y_1, ..., y_T) = NIW_{(\tilde{\nu}, \tilde{\Phi}, \tilde{\Psi}, \tilde{\Omega})}|J'_{H,p}J_{H,p}|^{\frac{1}{2}}, \tag{23}$$

where $J_{H,p} = J_1 J_2 J_3$, $J_3 = D^{-1}N$, D and N are the $n^2(\min(p, H) + 1) \times n^2(\min(p, H) + 1)$ and

---

[9] Our code utilizes the function logdet.m, provided on the Mathwork file exchange. © 2009, Dahua Lin.

$n^2(\min(p, H) + 1) \times \frac{n(n+1)}{+} n^2 \min(p, H)$ matrices such that

$$
D = \begin{bmatrix}
-D_n^+(A_0^{-1\prime} \otimes \Sigma + (\Sigma \otimes A_0^{-1\prime})K_n) & 0 \\
-C \begin{bmatrix} (A_0^{-1} \otimes A_0^{-1\prime})K_n \\ (A_0^{-1} \otimes B(1)^{-1}A_0^{-1\prime})K_n \end{bmatrix} & 0 \\
0 & I
\end{bmatrix}, \tag{24}
$$

$$
N = \begin{bmatrix}
I & 0 \\
0 & -C \begin{bmatrix} 0 \\ A_0^{-1}B(1)^{-1\prime} \otimes B(1)^{-1} \end{bmatrix} \\
0 & I
\end{bmatrix}, \tag{25}
$$

$D_n^+ = (D_n'D_n)^{-1}D_n'$ and D_n is the $n^2 \times \frac{n(n+1)}{2}$ duplication matrix such that $\text{vec}(\Sigma) = D_n\text{vech}(\Sigma)$.

Proof: Because the posterior distribution of B and Σ is $NIW_{(\tilde{\nu},\tilde{\Phi},\tilde{\Psi},\tilde{\Omega})}$, we only need to find the Jacobian matrix $J_{H,p}$ for transforming $[\text{vech}(\Sigma)' \ \text{vec}[B_1 \ \cdots \ B_p]']$ to $\theta = \text{vec}([\Theta_0 \ \Theta_1 \ \cdots \ \Theta_p])$, taking into account the zero restrictions in (22).

This Jacobian matrix is the product of three Jacobian matrices. The first two of these matrices are the same as J_1 and J_2 in Proposition 1, whereas J_3 is replaced by the Jacobian matrix of $\text{vec}([A_0 \ B_1 \ \cdots \ B_p])$ with respect to $[vech(\Sigma)' \ \text{vec}([B_1 \ \ldots \ B_p])']$. It follows from (22) that

$$
C \begin{bmatrix}
-(A_0^{-1} \otimes A_0^{-1\prime})K_n & 0 \\
-(A_0^{-1} \otimes B(1)^{-1}A_0^{-1\prime})K_n & A_0^{-1} \otimes B(1)^{-1\prime} \otimes B(1)^{-1}
\end{bmatrix} = 0. \tag{26}
$$

Thus the differential of $\text{vec}([A_0\ B_1\ \cdots\ B_p])$ can be written as

$$
-\begin{bmatrix} D_n^+(A_0^{-1\prime}\otimes\Sigma + (\Sigma\otimes A_0^{-1\prime})K_n) & 0 \\ -C\begin{bmatrix} (A_0^{-1}\otimes A_0^{-1\prime}))K_n \\ (A_0^{-1}\otimes B(1)^{-1}A_0^{-1\prime})K_n \end{bmatrix} & 0 \\ 0 & I \end{bmatrix}\begin{bmatrix} \text{vec}(A_0) \\ \text{vec}(B_1) \\ \vdots \\ \text{vec}(B_p) \end{bmatrix}
$$

$$
= \begin{bmatrix} I & 0 \\ -C\begin{bmatrix} 0 \\ A_0^{-1}B(1)^{-1\prime}\otimes B(1)^{-1} \end{bmatrix} & 0 \\ 0 & I \end{bmatrix}\begin{bmatrix} \text{vech}(d\Sigma) \\ \text{vec}[dB_1\ \cdots\ dB_p] \end{bmatrix}. \tag{27}
$$

Thus, we have $J_3 = D^{-1}N$ and it follows from Theorem 2 of Arias et al. (2018) that the volume element is given by

$$
|J_{H,p}'J_{H,p}|^{\frac{1}{2}} \tag{28}
$$

where $J_{H,p} = J_1 J_2 J_3$.

This proposition extends the analysis in Inoue and Kilian (2013) to nonrecursive exactly identified models based on any combination of short-run and long-run exclusion restrictions, while relaxing the assumption that $nirf = np$. It also covers, as a special case, models identified based on long-run exclusion restrictions only and models that are recursively identified by zero restrictions on the structural impact multiplier matrix. For $nirf = np$, our results for the latter model match the closed-form solution derived in Inoue and Kilian (2013).

29

4.3 Discussion

Our analysis of the Dirac delta loss function contributes to the literature in three dimensions. First, whereas in our earlier work on fully identified recursive and sign-identified structural VAR models we focused on the case of $nirf = np$, in the current paper, Propositions 1 and 2 allow for $nirf \leq np$ as well as for $nirf > np$. It should be noted that the first $nirf < np$ responses of the modal model obtained from the joint posterior distribution derived under the assumption that there are np responses, will not in general coincide with the responses of the modal model derived under the assumption that $nirf < np$ because the minimizer of the expected loss may change with the dimensionality of θ.

Second, Propositions 1 and 2 cover not only the recursively identified and sign-identified structural VAR models studied by Inoue and Kilian (2013), but also cover structural VAR models for which the joint posterior density of the structural impulse responses has not been derived to date. One example is nonrecursive models based on short-run and/or long-run restrictions. Another example is models based on combinations of sign and zero restrictions in the structural impact multiplier matrix. Especially the latter extension is practically important. Applications of structural VAR models to the transmission of global shocks to the domestic economy tend to be block recursive (e.g., Cushman and Zha 1997; Zha 1999). When global and/or domestic shocks are identified by sign restrictions, this results in a structural VAR model that combines sign and zero restrictions on A_0^{-1} (e.g., Mumtaz and Surico 2009; Kilian and Zhou 2020a) Similar situations arise, more generally, in sign-identified models when the feedback from one variable to another is delayed by informational or institutional frictions (e.g., Mountford and Uhlig 2009; Aastveit, Bjørnland and Thorsrud 2015; Nam and Wang 2018; Kilian and Zhou 2020b). When evaluating a model that combines zero and sign restrictions on the structural impact multiplier matrix, the posterior draws must be reweighted based on the importance sampler described in Arias et al. This

reweighting does not affect the value of the posterior density of a given model. It only affects how often this density value arises when resampling the original set of posterior draws.

Third, Proposition 1 generalizes the results in Inoue and Kilian (2013, 2019) by allowing for the imposition of additional narrative inequality restrictions in sign-identified models (with or without additional zero restrictions), as discussed in Antolin-Diaz and Rubio-Ramirez (2018). Since narrative restrictions restrict the likelihood, they require further reweighting of the posterior draws based on the importance sampler described in Antolin-Diaz and Rubio-Ramirez (2018). This reweighting again leaves unaffected the value of the posterior density for a given draw.[10]

4.4 Extensions to partially identified models

As in section 3, it is conceptually straightforward to extend the analysis under Dirac delta loss to partially identified structural VAR models. This requires the user to integrate out by numerical methods the elements of θ in the joint posterior distribution that are not identified (see Inoue and Kilian 2013). For example, in the model of Uhlig (2005) one would integrate out all responses to shocks other than the monetary policy shock. It should be noted, however, that this Monte Carlo integration approach is much more computationally costly than evaluating the relevant subset of θ under quadratic or absolute loss.

5 Empirical illustrations

In this section, we present two empirical VAR examples based on diffuse Gaussian-inverse Wishart priors for the reduced-form parameters. All models include an intercept. One example is based on Proposition 2 and covers the case of $nirf > np$ whereas the other example is based on Proposition

---

[10]It should be noted that one does not necessarily have to estimate the structural VAR model in the notation used by Arias et al. (2018) and Antolin-Diaz and Rubio-Ramirez (2018) to make use of our results. When using existing code, VAR model estimates can typically be expressed in terms of A_0 and A_+ matrices for the purpose of evaluating the posterior density value of a given posterior draw without having to re-simulate the posterior draws using their notation.

1 and covers the case of $nirf < np$.

5.1 Example for $nirf > np$

Our first empirical example is a stylized model of the U.S. macroeconomy based on short-run and long-run exclusion restrictions, as discussed in Rubio-Ramirez, Waggoner and Zha (2010). This specific example was chosen because models combining short-run and long-run restrictions are not covered by the theoretical results in Inoue and Kilian (2013, 2019). We specify a quarterly VAR(4) model for real GNP growth ($\Delta \log Y$), the federal funds rate (R) and GNP deflator inflation ($\Delta \log P$). Variation in these data is explained in terms of an aggregate demand shock, an aggregate supply shock, and a monetary policy shock. The identifying restrictions are that aggregate demand shocks have no long-run effect on real output and that monetary policy shocks have neither an impact effect on real output nor a long-run effect on real output. The model is exactly identified. The estimation period is 1954.IV-2007.IV. We solve for the structural impact multiplier matrix, as discussed in Rubio-Ramirez, Waggoner and Zha (2010). The sign of the monetary policy shock is normalized to imply an increase in the federal funds rate on impact. The sign of the aggregate supply shock is normalized to ensure that the aggregate supply shock raises real output on impact. Given that the structural model is only loosely restricted, we follow the convention of reporting 68% credible sets. The maximum impulse response horizon is 12 quarters. The number of estimated structural impulse responses, $nirf = n^2(H + 1) - 1 = 116$. exceeds the number of estimated structural parameters, $np = n^2(p + 1) - 3 = 42$, where 3 refers to the number of short-run and long-run exclusion restrictions.

Although we are considering all 9 structural impulse functions when minimizing the expected loss, in the interest of space, only a subset of the results are shown. Figure 2 focuses on the response of the model variables to an unanticipated monetary tightening. Such a shock would be expected

to raise the interest rate, lower real GNP and lower GNP deflator inflation. Figure 2 allows us to assess the impact of the choice of alternative loss functions. The 68% joint credible sets look broadly similar across loss functions, but there is more uncertainty under Dirac delta loss than under absolute or quadratic loss in this example, as measured by the range of responses. Likewise, the preferred impulse response estimates are broadly similar. They indicate a transitory increase in the interest rate and a persistent decline in real GNP. Inflation is not very responsive. Only the estimator derived under Dirac delta loss is consistent with a persistent decline in the price level in response to a monetary policy tightening.

Figure 3 illustrates that we would have underestimated the extent of the estimation uncertainty if we had relied on conventional pointwise 68% error bands. In some cases, the range spanned by the responses in the joint credible set is two or three times as wide. In other words, the use of pointwise error bands may cause the user to underestimate the estimation uncertainty at short as well as long horizons. The median and mean response functions in this example are not too far from the corresponding estimates under absolute and quadratic loss in Figure 1, but the mean and median inflation responses in Figure 3 differs from the estimate derived under Dirac delta loss in Figure 2.

5.2 Example for $nirf < np$

The second empirical example is a monthly VAR(24) model of the global market for crude oil. The model variables include the percent change in global crude oil production, an appropriate measure of the global business cycle, the log real price of oil, the change in global crude oil inventories, the ex ante real market rate of interest, and the trade-weighted U.S. real exchange rate. The model is block recursive. The structural shocks in the oil market block are identified based on static and dynamic sign restrictions, as is standard in the literature, complemented by elasticity bounds. The

remaining structural shocks are identified based on empirically motivated exclusion restrictions and economically motivated dynamic sign restrictions.

Because the model is estimated subject to both sign and exclusion restrictions in the structural impact multiplier matrix, the posterior draws have to be reweighted based on an importance sampler, as described in Arias et al. (2018). The model is estimated subject to additional narrative inequality restrictions, which require a second layer of importance sampling, as discussed in Antolin-Diaz and Rubio-Ramirez.(2018). The structural impulse responses are set-identified. The estimation period is 1973.2-2018.6. Details of the construction of the data and of the identifying assumptions can be found in Kilian and Zhou (2020a). The maximum impulse response horizon is 18 months. Given the strength of the identifying restrictions, we report 95% credible sets. The number of estimated structural parameters, $np = n^2(p + 1) - 9 = 891$, exceeds the number of estimated structural impulse responses, $nirf = n^2(H + 1) - 9 = 639$, where 9 refers to the number of zero restrictions in the impact period.

Like in the previous example, we evaluate the model based on all 36 impulse response functions jointly, but, for illustrative purposes, we report only a subset of the results. Figure 4 focuses on the response of selected model variables to an exogenous decline in the real value of the dollar. Such a shock is expected to raise global real activity and the real price of oil and to possibly lower oil production and oil inventories. This is largely what we find in Figure 4, except that the response of the real price of oil is clearly positive only at short horizons. At longer horizons the estimation uncertainty about the oil price response grows. In contrast, the responses of global real activity and of oil inventories are precisely estimated under all three loss functions. Likewise, the joint credible sets in Figure 4 are largely identical, regardless of the choice of loss function.

Unlike in the previous example, the 95% pointwise error bands in Figure 5 cover roughly the range of the response functions in Figure 4. The median response function for the real price of

oil in Figure 5, however, suggests a more persistently positive response of the real price of oil to an exogenous U.S. real exchange rate depreciation than the estimate based on the absolute loss function in Figure 4. The latter estimate largely agrees with the estimate derived under the Dirac delta loss function. The differences between the mean response function and the impulse response estimate derived under quadratic loss are less pronounced.

6   Concluding remarks

Several studies have voiced concerns about the use of posterior median response functions in applied VAR analysis and about the use of error bands based on quantiles of the marginal posterior distribution of the structural impulse responses (e.g., Sims and Zha 1999; Fry and Pagan 2011; Kilian and Murphy 2012; Inoue and Kilian 2013). In this paper, we make precise what these concerns are and we propose alternative methods of Bayesian impulse response inference that avoid these drawbacks.

We first show that there may not exist an impulse response vector in the set of all possible posterior draws for the impulse responses that matches the posterior median response function, even as the number of posterior draws approaches ∞, calling into question the use of posterior median response functions. Analogous concerns apply to posterior mean response functions. As a result, the use of mean and median response functions may distort the dynamics of the impulse response function. Since the shape of response functions (as well as the comovement across impulse response functions) are of central interest in applied work, we conclude that posterior median response functions and posterior mean response functions should not be used in applied work. In response to this problem, we propose new estimators of the impulse response function under quadratic and under absolute loss that are consistent with Bayesian statistical decision theory, that are optimal in the relevant sense, that respect the dynamics of the impulse response function, and

that are easy to implement.

We then discuss in detail why the construction of impulse response error bands based on the quantiles of the marginal posterior distributions is also inappropriate. There are three main concerns. First, vectors of upper and lower quantiles derived from the marginal impulse response posterior distributions, like vectors of medians, need not be contained in the set of impulse response functions that can be generated by the model and hence may distort the shape of the response functions. Second, two-dimensional error bands are consistent with any shape of the response functions that fits within the error band. They do not tell us anything about likely departures from the baseline estimate of the impulse response function. Third, the uncertainty about the impulse response function clearly depends on the covariance across individual impulse response estimates, which cannot be captured without considering the joint distribution of the impulse responses. In response to these concerns, we provide a constructive alternative to the use of conventional pointwise error bands. We propose the construction of joint credible sets under absolute and under quadratic loss that capture the full uncertainty about the impulse response estimates, that provide information about likely departures from the path of the estimated impulse response function and that are guaranteed to be contained in the set of feasible impulse responses functions.

Even though our focus has been on structural impulse response analysis in VAR models, we note that the same methods may be easily adapted to estimate the path of a forecast and to characterize the joint uncertainty about this path. Similar problems also arise in the evaluation of impulse responses based on Bayesian estimates of DSGE models. For example, Herbst and Schorfheide (2015) report posterior mean response functions and pointwise quantile error bands computed based on draws from the posterior distribution of the structural DSGE model parameters. Such estimates are subject to analogous concerns. For example, the posterior mean response function generates estimates inconsistent with the DSGE model structure, if the number of responses exceeds

the number of structural DSGE model parameters. More importantly, the pointwise error bands by construction misrepresent the uncertainty about the response functions. These concerns can be addressed using the tools developed in our paper.

Although the traditional approach in applied studies has been to focus on absolute or quadratic loss functions, respectively, there has been growing interest in recent years in the application of the Dirac delta loss function to the estimation of structural impulse responses. In this paper, we provide a substantial generalization of the approach of Inoue and Kilian (2013, 2019), which involved evaluating the joint density of the set of structural impulses under a Dirac delta loss function when the number of estimated structural responses equals the number of estimated structural model parameters. Inoue and Kilian identified the modal (or most likely) posterior draw of the structural impulse responses by ranking the posterior draws by the value of the implied joint posterior density of these responses. They then reported this set of responses as the most likely impulse response estimate, along with joint $1 - \alpha$ probability HPD credible sets based on the $(1 - \alpha)100\%$ most likely draws in the set of solutions that satisfy the identifying restrictions. Like the new estimators we derived under absolute and under quadratic loss, this approach addresses all the concerns that have been raised about posterior median (and mean) response functions and the corresponding pointwise error bands.

We improve on this work by presenting alternative analytic derivations of the joint posterior density of the set of structural impulse responses that accommodate arbitrary horizons and lag orders and that allow for nonrecursive identification schemes based on short- and/or long-run exclusion restrictions, for combinations of sign and exclusion restrictions, and for narrative restrictions, none of which are covered by the results in Inoue and Kilian (2013, 2019). Our Bayesian estimators and joint credible sets based on the quadratic or the absolute loss function may also be applied to the alternative Bayesian approaches discussed in Baumeister and Hamilton (2018) and Plagborg-Møller

(2019). We leave for future research the question of how to derive analogous results for the joint posterior density of structural responses under their assumptions.

There is little to choose between the three loss functions based on computational cost or ease of implementation. Two empirical illustrations suggest that the three Bayesian impulse response estimators proposed in this paper in some case may generate similar results, but at other times can differ. The same is true for the difference between conventional median or mean response functions and the improved impulse response estimators proposed in this paper. Since there is no a priori reason for conventional impulse response estimators to be close to the optimal impulse response function estimators derived in this paper under absolute (and under quadratic) loss, the latter should routinely replace median (and mean) response functions in applied work. We also showed by example that conventional inference based on pointwise error bands may greatly understate the estimation uncertainty and cannot be recommended for applied work. They should be routinely replaced by the joint credible sets proposed in this paper.

Acknowledgments

References

1. Aastveit, K.A, Bjørnland, H.C., Thorsrud, L.A., 2015. What drives oil prices? Emerging vs. developed economies. Journal of Applied Econometrics 30, 1013-1028. https://doi.org/10.1002/jae.2406
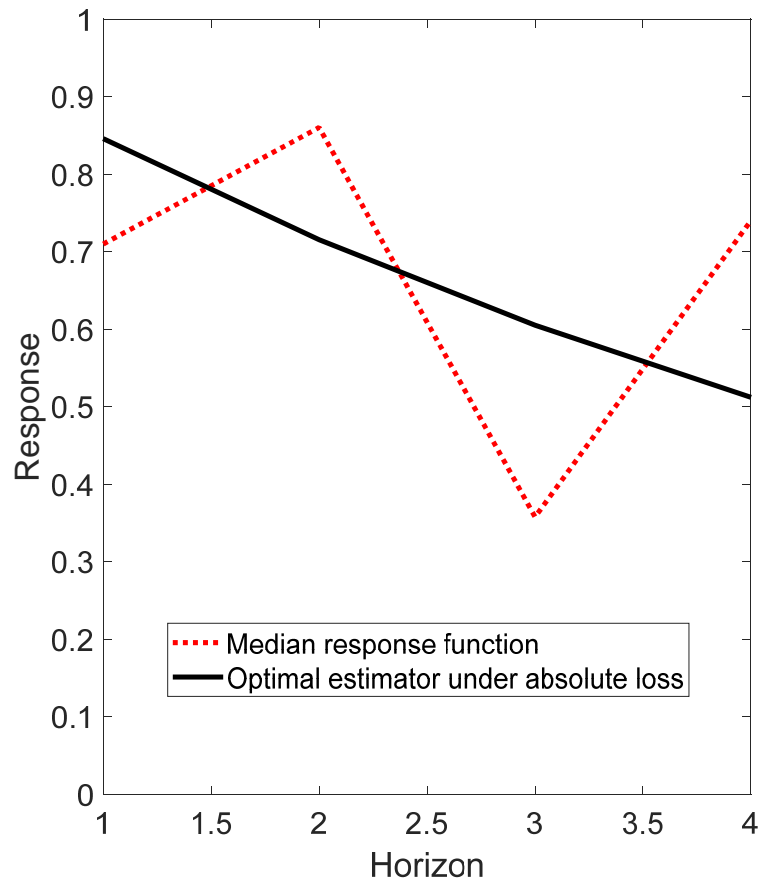
2. Antolin-Diaz, J., Rubio-Ramirez, J.F., 2018. Narrative sign restrictions for SVARs. American Economic Review 108, 2802-2839. https://doi.org/10.1257/aer.20161852

3. Arias, J., Rubio-Ramirez, J.F., Waggoner, D.F., 2018. Inference based on SVARs identified with sign and zero restrictions: Theory and applications. Econometrica 86, 685-720. https://doi.org/10.3982/ECTA14468

4. Baumeister, C., Hamilton, J.D., 2018. Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations," Journal of Monetary Economics 100, 48-65. https://doi.org/10.1016/j.jmoneco.2018.0

5. Berger, J.O., 1985. Statistical decision theory and Bayesian analysis. 2nd ed. Springer, New York. https://doi.org/10.1007/978-1-4757-4286-2

6. Bernardo, J.M., Smith, A.F.M., 1994. Bayesian Theory. John Wiley and Sons, Chichester, UK. https://doi.org/10.1002/9780470316870

7. Blanchard, O.J., 1989. A traditional interpretation of macroeconomic fluctuations. American Economic Review 79, 1146–1164.

8. Bruder, S., Wolf, M., 2018. Balanced bootstrap joint confidence bands for structural impulse response functions. Journal of Time Series Analysis 39, 641-664. https://doi.org/10.1111/jtsa.12289

9. Cochrane, J.H., 1994. Shocks. Carnegie-Rochester Conference Series on Public Policy, 41, 295-364. https://doi.org/10.1016/0167-2231(94)00024-7

10. Cross, J.L., Nguyen, B.H., Tran, T.D, 2019. The role of precautionary and speculative demand in the global market for crude oil. Manuscript, BI Norwegian Business School.

11. Cushman, D.O., Zha, T., 1997. Identifying monetary policy in a small open economy under

flexible exchange rates. Journal of Monetary Economics 39, 433-448. https://doi.org/10.1016/S0304-3932(97)00029-9

12. Eichenbaum, M., Evans, C.L. 1995. Some empirical evidence on the effects of monetary policy shocks on exchange rates. Quarterly Journal of Economics 110, 975-1009. https://doi.org/10.2307/2946646

13. Fry, R., Pagan, A.R., 2011. Sign restrictions in structural vector autoregressions: A critical review. Journal of Economic Literature 49, 938-960. https://doi.org/10.1257/jel.49.4.938

14. Herbst, E., Schorfheide, F., 2015. Bayesian estimation of DSGE models. Princeton University Press, Princeton. https://doi.org/10.23943/princeton/9780691161082.001.0001

15. Herrera, A.M., Rangaraju, S.K., 2020. The effects of oil supply shocks on U.S. economic activity: What have we learned? Journal of Applied Econometrics, 35, 141-159. https://doi.org/10.1002/jae.2735

16. Herwartz, H., Plödt, M., 2016. The macroeconomic effects of oil price shocks: Evidence from a statistical identification approach. Journal of International Money and Finance 61, 30-44. https://doi.org/10.1016/j.jimonfin.2015.11.001

17. Inoue, A., Kilian, L., 2013. Inference on impulse response functions in structural VAR models. Journal of Econometrics, 177, 1-13. https://doi.org/10.1016/j.jeconom.2013.02.009

18. Inoue, A., Kilian, L., 2016. Joint confidence sets for structural impulse responses. Journal of Econometrics 192, 421-432. https://doi.org/10.1016/j.jeconom.2016.02.008

19. Inoue, A., Kilian, L. 2019. Corrigendum to 'Inference on impulse response functions in structural VAR models' [J. Econometrics 177 (2013), 1-13]. Journal of Econometrics. 209, 139-143. https://doi.org/10.1016/j.jeconom.2017.08.020

20. Kilian, L., Lütkepohl, H., 2017. Structural vector autoregressive analysis. Cambridge University Press, New York. https://doi.org/10.1017/9781108164818

21. Kilian, L., Murphy, D.P., 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. https://doi.org/10.1111/j.1542-4774.2012.01080.x

22. Kilian, L., Zhou, X., 2020a. Oil prices, exchange rates and interest rates. Manuscript, Federal Reserve Bank of Dallas.

23. Kilian, L., Zhou, X., 2020b. Does drawing down the U.S. strategic petroleum reserve help stabilize oil prices? Journal of Applied Econometrics, forthcoming.

24. Koop, G., 1996. Parameter uncertainty and impulse response analysis. Journal of Econometrics 72, 135-149. https://doi.org/10.1016/0304-4076(94)01717-4

25. Lütkepohl, H., Staszewska-Bystrova, A., Winker, P., 2015a. Confidence bands for impulse responses: Bonferroni vs. Wald. Oxford Bulletin of Economics and Statistics 77, 800-821. https://doi.org/10.1111/obes.12114

26. Lütkepohl, H., Staszewska-Bystrova, A., Winker, P., 2015b. Comparison of methods for constructing joint confidence bands for impulse response functions. International Journal of Forecasting 31, 782-798. https://doi.org/10.1016/j.ijforecast.2013.08.003

27. Lütkepohl, H., Staszewska-Bystrova, A., Winker, P., 2018. Calculating joint confidence bands for impulse response functions using highest density regions. Empirical Economics 55, 1389-1411. https://doi.org/10.1007/s00181-017-1325-3

28. Montiel Olea, J.L., Plagborg-Møller, M., 2019. Simultaneous confidence bands: Theory, implementation, and an application to SVARs. Journal of Applied Econometrics 34, 1-17. https://doi.org/10.1002/jae.2656

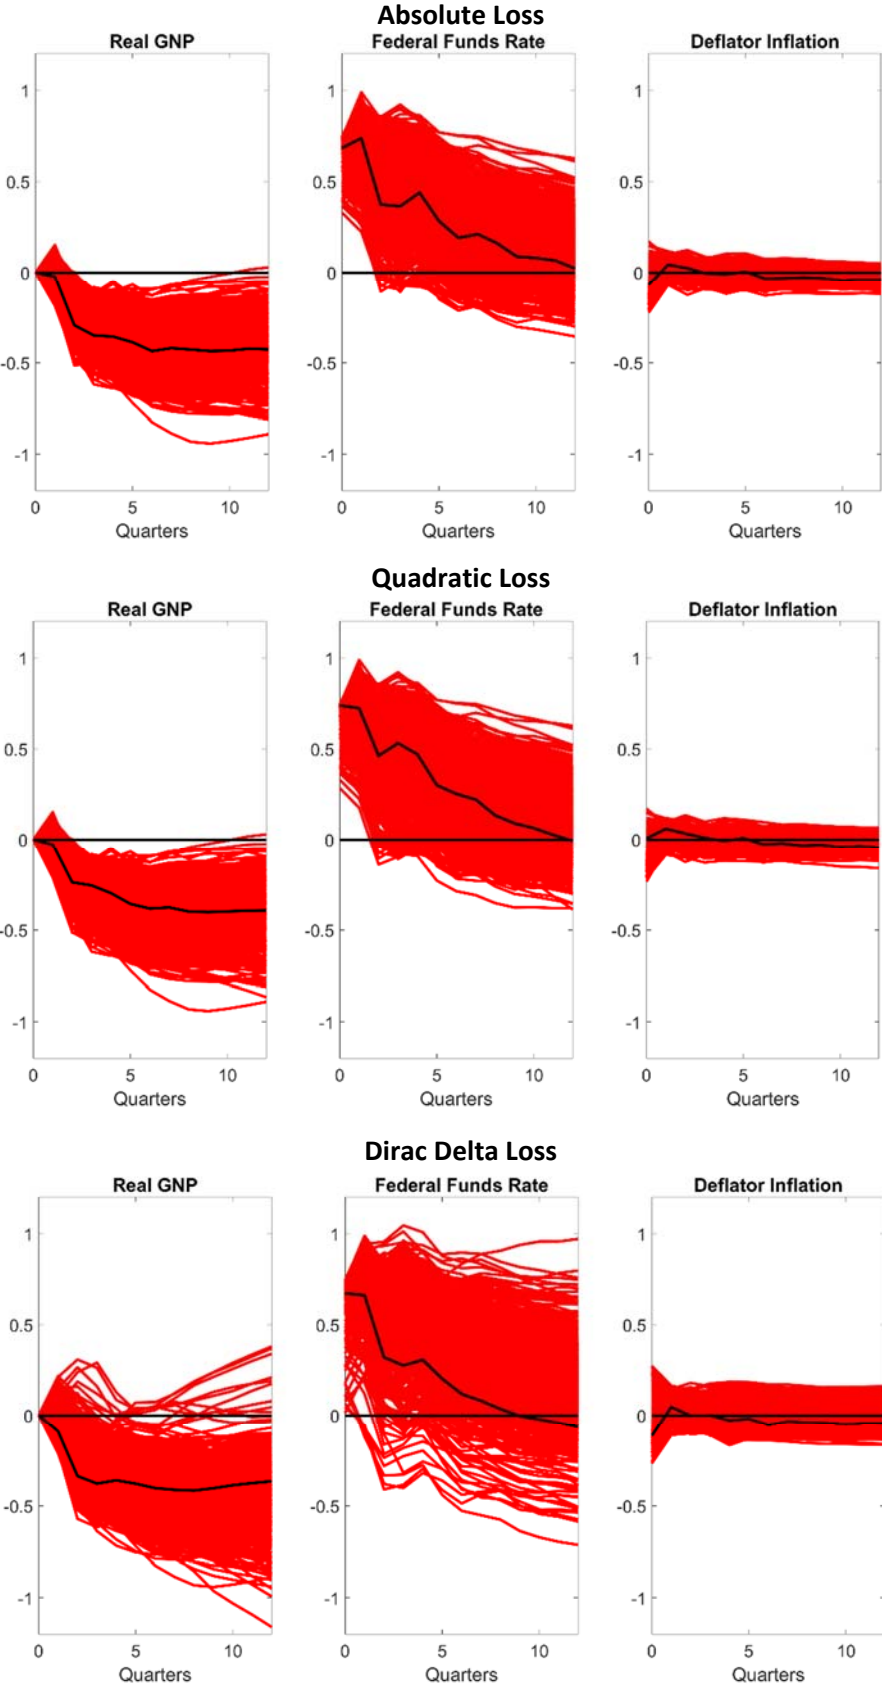29. Mumtaz, H., Surico, P., 2009. The transmission of international shocks: A factor-augmented VAR approach. Journal of Money, Credit and Banking 41, 71-100. https://doi.org/10.1111/j.1538-4616.2008.00199.x

30. Nam, D.,Wang, J., 2018. Mood swings and business cycles: Evidence from sign restrictions. Journal of Money, Credit, and Banking 51, 1623-1649. https://doi.org/10.1111/jmcb.12568

31. Plagborg-Møller, M., 2019. Bayesian inference on structural impulse response functions. Quantitative Economics 10, 145-184. https://doi.org/10.3982/QE926

32. Rubio-Ramirez, J.F., Waggoner, D., Zha, Z., 2010. Structural vector autoregressions: Theory of identification and algorithms for inference. Review of Economic Studies 77, 665-696. https://doi.org/10.1111/j.1467-937X.2009.00578.x

33. Sims, C.A., Zha, T., 1999. Error bands for impulse responses. Econometrica 67, 1113-1156. https://doi.org/10.1111/1468-0262.00071

34. 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. https://doi.org/10.1016/j.jmoneco.200

35. Waggoner, D., Zha, T., 2012. Confronting model misspecification in macroeconomics. Journal of Econometrics 171, 167-184. https://doi.org/10.1016/j.jeconom.2012.06.013

36. Woodford, M., 2003. Interest and prices: Foundations of a theory of monetary policy. Princeton University Press, Princeton.

37. Zha, T., 1999. Block recursion and structural vector autoregressions. Journal of Econometrics 90, 291-316. https://doi.org/10.1016/S0304-4076(98)00045-1

38. Zhou, X., 2020. Refining the workhorse oil market model," Journal of Applied Econometrics, 35, 130-140. https://doi.org/10.1002/jae.2743

Figure 1: AR(1) example of how median response functions may distort the shape of the impulse response function
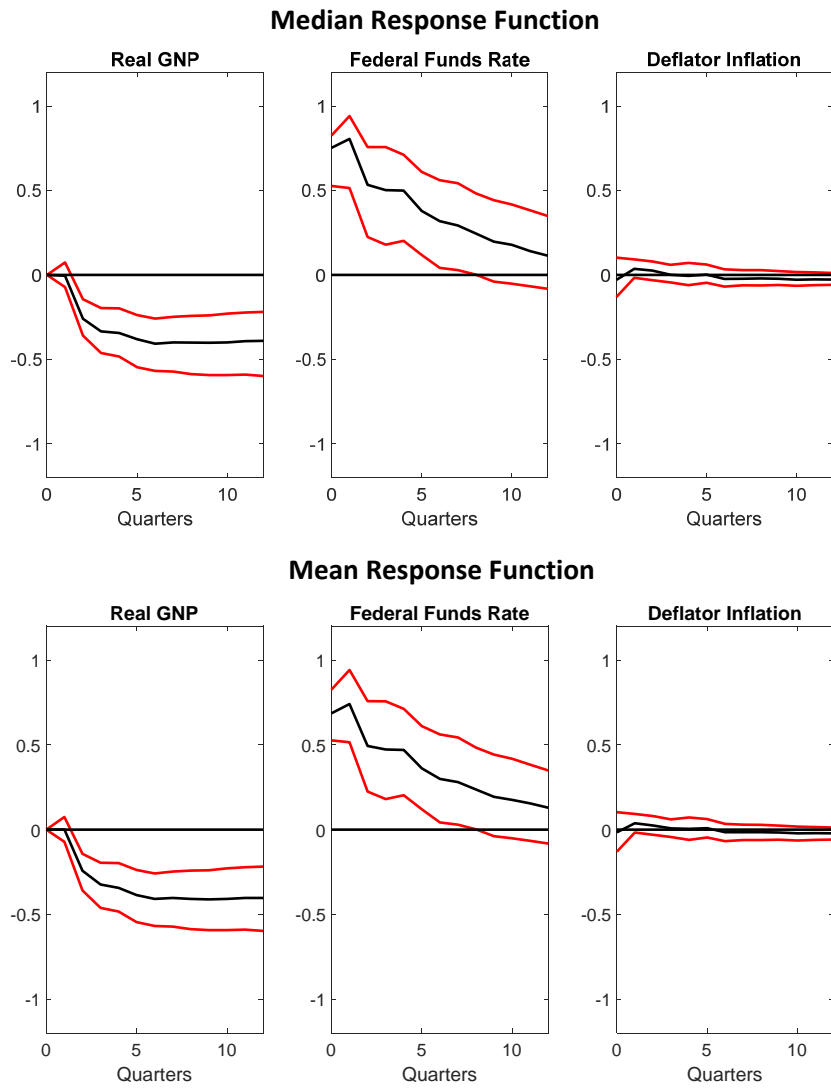

NOTES: Responses generated from slope parameter in AR(1) model with posterior $\rho \sim (\rho_{mean}, 1/n_T)$ with $\rho_{mean} = 0.7$ and $n_T = 0.769$. All estimates are based on 10,000 random draws from the posterior.
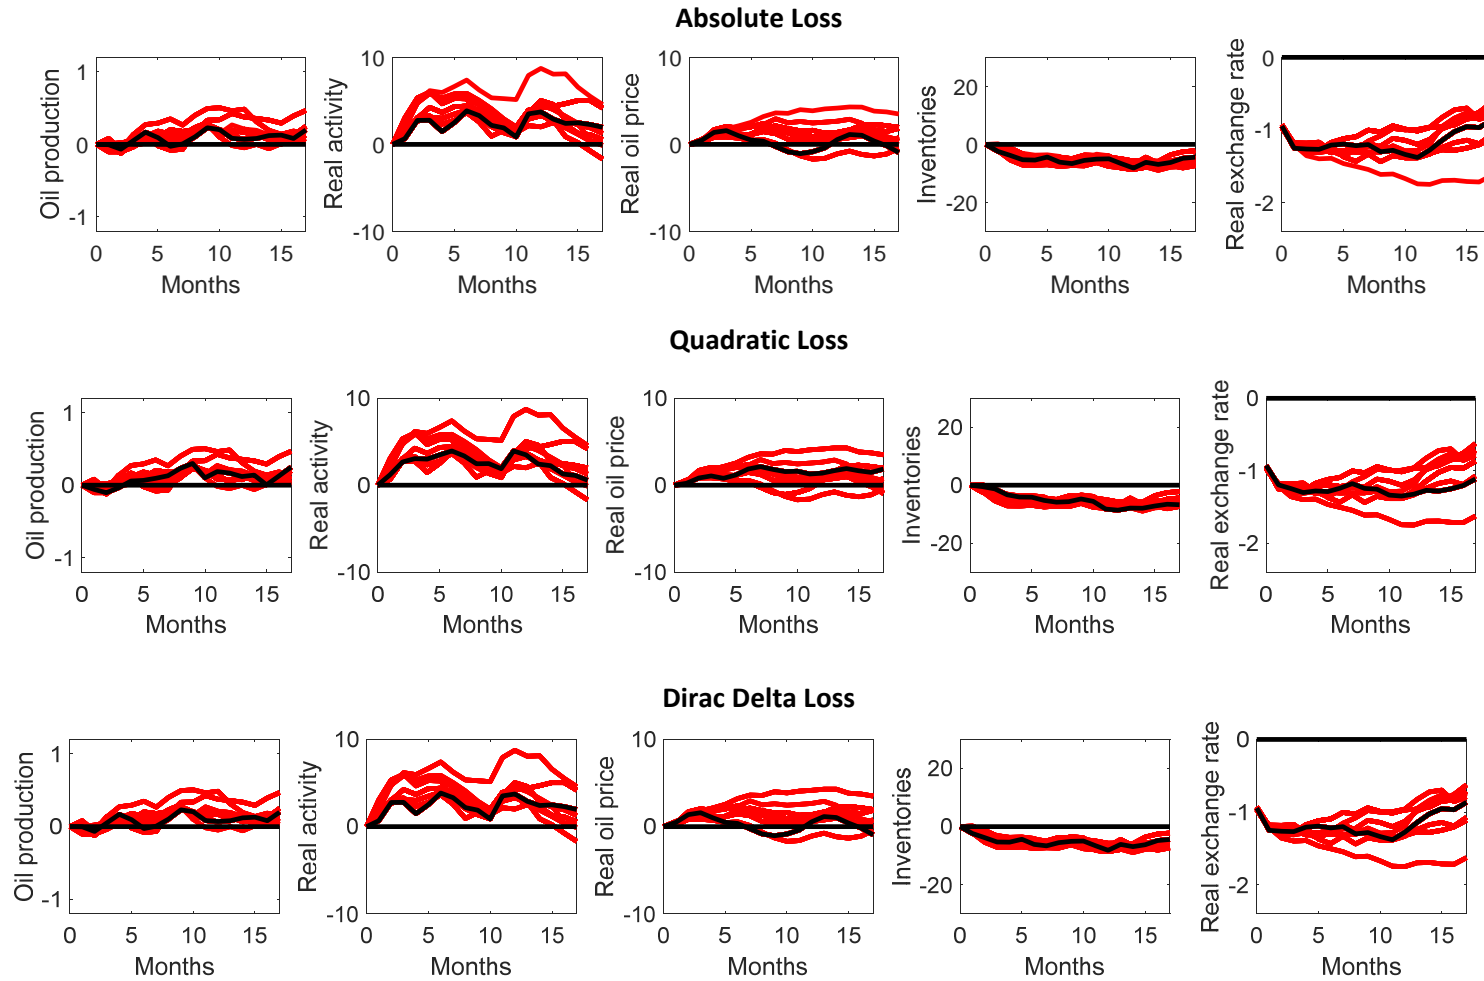
Figure 2: Responses to exogenous U.S. real interest rate shock: 68% joint credible set under alternative loss functions with preferred estimate highlighted in black

Figure 3: Responses to exogenous U.S. real interest rate shock: Pointwise response functions with 68% pointwise error bands
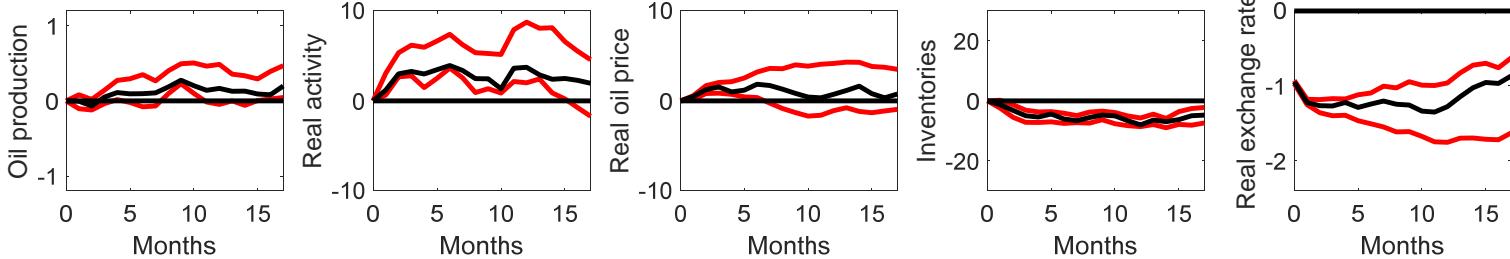
Median Response Function

Real GNP   Federal Funds Rate   Deflator Inflation

Mean Response Function

Real GNP   Federal Funds Rate   Deflator Inflation

Figure 4: Responses to exogenous U.S. real interest rate shock:
95% joint credible set under alternative loss functions with preferred model highlighted in black

46

Figure 5: Responses to exogenous U.S. real interest rate shock:
Pointwise response functions with 95% pointwise error bands

Median Response Function

Mean Response Function