A (Robust) Bayes Approach to Non-Linear Functions of Dynamic Causal Effects*

(Incomplete) Preliminary Draft

Martin Fankhauser†

September 1, 2025

Abstract

Modern tools for macroeconomic policy evaluation and causal inference often rely on sufficient statistics, which are non-linear functions of impulse responses. This paper extends these methods to the case where we are only able to set identify dynamic causal effects within a VAR framework. I examine the complications that arise when applying non-linear transformations—such as regressions in the impulse response space—under set identification, and I propose a robust Bayesian approach to address these issues. Further, by expressing parameters of interest as functions of the VAR's orthogonal reduced form, I introduce a novel class of identification strategies sharpening the identification of dynamic causal effects.

Keywords:

JEL Codes:

*I thank Massimiliano Marcellino, Toru Kitagawa, Geert Mesters, Giorgio Primiceri, Nicola Pavoni, Pamela Giustinelli, Efrem Castelnuovo, Amilcar Velez, Jiaming Huang and Tommaso Tornese for helpful discussions and comments. All mistakes are my own.

†PhD. Bocconi University

1 Introduction

A growing body of recent work derives structural parameters and policy-relevant statistics either directly as parameters of regressions in the impulse response space (henceforth, RIRS) (e.g., Barnichon and Mesters (2020), Barnichon and Mesters (2023b), Lewis and Mertens (2022), McKay and Wolf (2023)) or in general as other non-linear functions of impulse responses (e.g. Barnichon and Mesters (2021), Ramey and Zubairy (2018), Barnichon and Mesters (2023a)).[1] These studies demonstrate that impulse responses and forecasts serve as sufficient statistics for a wide range of macroeconomic questions, thereby reducing the need to fully specify a structural model in certain contexts. While these approaches are theoretically appealing, they typically require rich information: researchers must be able to point-identify impulse responses to multiple policy and non-policy shocks. This paper aims to relax that informational burden by exploring how to conduct inference on non-linear functions of impulse responses when impulse responses are only set-identified within a VAR framework.

The core econometric challenge posed by set-identified impulse responses in this context is intuitive: For instance, conducting a RIRS with set-identified impulse responses is analogous to estimating a linear regression when the data—either the regressors, the outcome, or both—are only known up to an interval. In such cases, a range of regression coefficients may be consistent with both the linear model and the observed data. As Figure 1 illustrates for the case where the outcome is perfectly observed but the regressors are interval-valued, even the sign of the regression coefficients can be indeterminate.

[1] see Barnichon and Mesters (2024) for an overview on policy statistics.

Figure 1: Regression with interval regressors (X) and outcomes ($y_1 > y_2 > y_3$)

In the microeconometric literature deriving the worst-case sharp identified set[2] for regression problems with interval regressors is still an open problem and existing solutions either introduce auxiliary assumptions (Manski and Tamer (2002)) or derive outer sets (see e.g. for the related problem of missing data Aucejo, Bugni, and Hotz (2017)). For the related problem of finding the best linear predictor see for example Horowitz et al. (2003) or Beresteanu, Molchanov, and Molinari (2011).[3] The RIRS setting considered here, however, offers more structure than the general interval regression problem. Specifically, rather than evaluating all possible combinations of point values for each observation, one only needs to consider all admissible trajectories of impulse responses. When the impulse responses are generated from a VAR, I show that the RIRS coefficients can be expressed in terms of the orthogonalized reduced-form parameters of the VAR, substantially simplifying the problem.

In macroeconomics, the predominant method for set-identifying impulse responses involves the use of sign restrictions within a Bayesian VAR framework, for early contributions see Faust (1998) and Uhlig (2005). The seminal work of Moon and Schorfheide (2012) demonstrates that the posterior distribution in set-identified models contains an unrevisable component that remains unaffected by the data. Consequently, inference within the identified set is entirely driven by prior beliefs. As Figure 1 illustrates, the parameters of a RIRS can be sensitive to any kind of imputation rule, i.e. in our case the location of the posterior mean of the impulse

[2] i.e. the smallest set of parameter values consistent with the data and the assumed model.

[3] In ongoing companion work we derive the sharp identifed set for the general linear regression problem with missing and interval data extending the ideas in Beresteanu, Molchanov, and Molinari (2011) to this problem. The tools derived in this work could potentially be also utilized for the RIRS problem at hand, if estimates for the bounds of the dynamic causal effects are available. An advantage of this procedure would be that it could also handle missing impulse responses. I leave this for future work.

responses within the identified set, suggesting that for RIRS parameters inference is even more sensitive to prior information than inference on the impulse responses themselves.[4]

The contribution of this paper is four-fold: First, to the best of my knowledge this paper is the first highlighting the above issue for the non-linear sufficient statistic approaches mentioned above. I further extend the above argument and show that inference can be potentially misleading via simulation studies and in applications. Second, I express the parameters of interest in terms of the orthogonalized reduced-form parameters of the VAR and study this mapping. Continuity and boundedness thereof is established for several examples of sufficient statistic approaches and I propose a general algorithm to verify path-connectedness of the respective domain of these mappings. These results allow the application of the robust Bayes procedure of Giacomini and Kitagawa (2021), delivering the sharp identified set for RIRS parameters, i.e. the worst-case bounds consistent with the data, thereby avoiding overreliance on arbitrary priors. Finally, the continuity of the mapping from VAR parameters to the parameters of interest enables the introduction of a new class of structural identification strategies. For example, researchers can impose sign restrictions on causal parameters such as those in a (causal) Phillips curve, or derive impulse responses under the assumption that specific policy institutions systematically under- or overreacted during historical episodes. The theoretical properties of such restrictions are studied for several examples and applied to sharpen identification of an optimal policy coefficients for the US.

This paper is structured as follows: After Section 2 briefly discusses the related literature, Section 3 proceeds by illustrating the main points of this work within a simple simulation study. Subsequently, the general theory is presented in Section 4, before Section 5 applies the theory to the identification of a US monetary policy shock as well as derive robust set identified bounds for a sequence of optimal policy perturbations. Section 6 concludes.

[4]For the latter, see discussions in Baumeister and Hamilton (2015),Arias and Waggoner (2024), etc.

2 Literature Review

To be completed.

3 Motivating Example

This section delineates the main contributions of this paper within a simple example before the subsequent sections develop the formal theory to justify the procedures used. Suppose the economy follows the following log-linearized solution to a standard 3-variable New Keynesian model:

$$\begin{aligned}
\pi_t &= \mathbb{E}_t[\pi_{t+1}] + \kappa x_t - \epsilon_t^S \\
x_t &= \mathbb{E}_t[x_{t+1}] - \frac{1}{\sigma}(i_t - \mathbb{E}_t[\pi_{t+1}]) + \epsilon_t^D \\
i_t &= \phi_\pi \pi_t + \phi_{\epsilon^S} \epsilon_t^S + \epsilon_t
\end{aligned}$$

with π_t the inflation gap, x_t the output gap, i_t the nominal interest rate set by the central bank, and ϵ_t^S and ϵ_t^D represents supply and demand shocks, respectively. The non-standard policy rule includes the true structural supply shock in order to better control how optimal a specific policy maker is reacting to different types of shocks.The model solution for the non-policy variables will be of the form:

$$\begin{pmatrix} \pi_t \\ x_t \end{pmatrix} = \Gamma \begin{pmatrix} \epsilon_t^S \\ \epsilon_t^D \end{pmatrix} + \mathcal{R}\epsilon_t$$

where Γ collects the impulse responses to the supply and demand shock and $\mathcal{R}$ collects the impulse responses with respect to the policy shock. Under a quadratic loss function weighting the output and inflation gap equally, Barnichon and Mesters (2023b) show that the optimal reaction adjustment (ORA) is given by:

$$\text{ORA} = -(\mathcal{R}'\mathcal{R})^{-1}\mathcal{R}'\Gamma$$

where implicitly the impulse responses are understood as those under the baseline rule. There are two subset ORA statistics, only measuring whether the response to a demand or supply shock has been optimal. For simplicity, in the following we I will only focus on the response to supply shock. In general, the ORA will be close to zero if the conduct of policy is optimal. A negative value for the ORA statistic indicates that the policy rate is too high, e.g. it increased too much or declined too slow, after a respective non-policy shock. Note that this is just one example of a statistic of interest, the theory developed in this paper is applicable more broadly as the subsequent sections will show.

Suppose a researcher estimates a VAR using a sample of 5000 observations of the triplet $y_t = (\pi_t, x_t, i_t)$ and imposing the following sign restrictions to identify the impulse responses:

$$\begin{pmatrix} u_t^x \\ u_t^\pi \\ u_t^i \end{pmatrix} = \begin{pmatrix} ? & + & - \\ ? & - & - \\ ? & - & + \end{pmatrix} \times \begin{pmatrix} \epsilon_t^D \\ \epsilon_t^S \\ \epsilon_t \end{pmatrix}$$

Thus, the monetary policy shock is the only one affecting inflation and the interest rate in the opposite direction and the supply shock is the only shock affecting the output gap and inflation in opposite directions. On top of these standard restrictions I also impose a negative output response to the policy shock and a negative response of the interest rate with respect to a supply shock. Importantly, the researcher does not know the true data generating process.

Figure 2: Top panel: ORA statistic and Impulse Responses

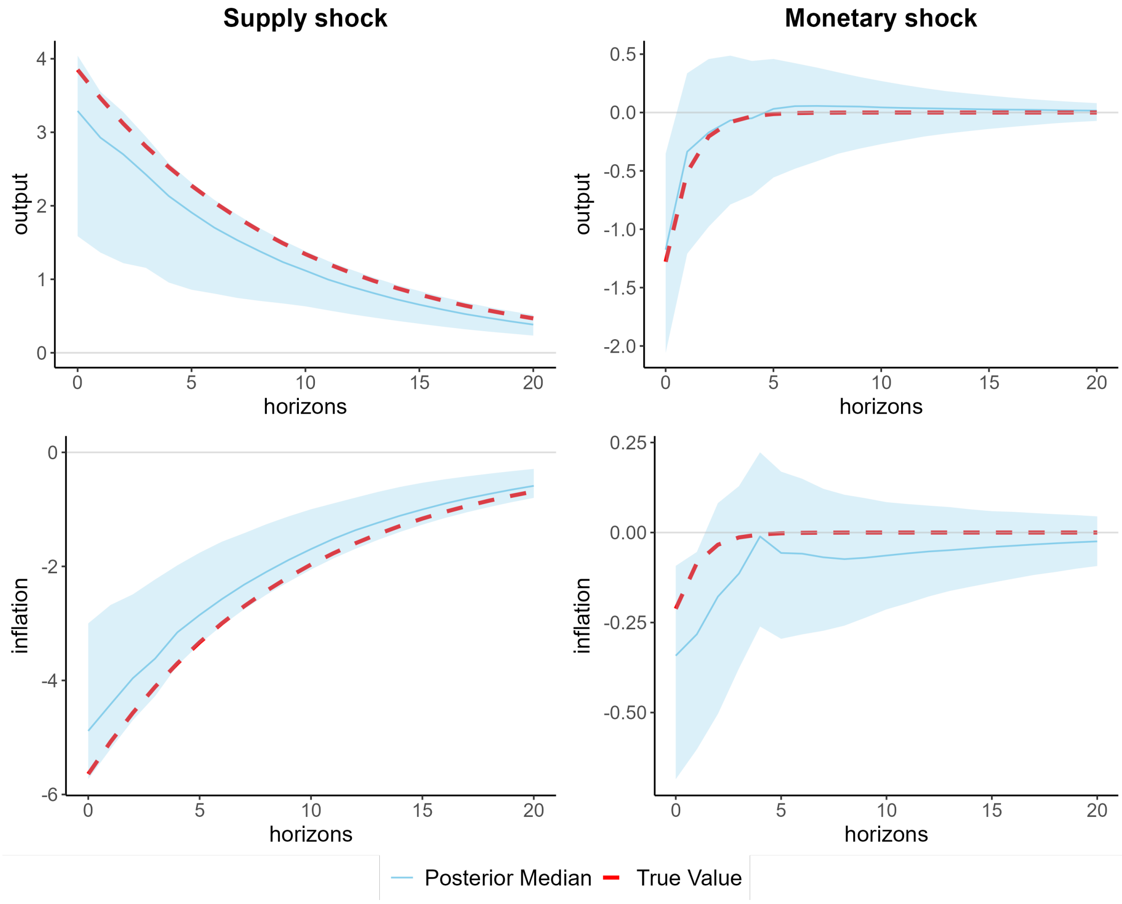

Note:

Figure 2 reports the estimated impulse responses. The blue solid lines represent the posterior medians, and the shaded regions the 95% pointwise credible intervals, obtained under the common uniform prior on the orthogonal matrix (the Haar prior).[5] Comparing these posterior distributions with the true impulse responses (red dashed lines), the sign restrictions overall recover the true effects well. Nevertheless, in a set-identified environment there is no reason to expect the posterior distribution to be centered on the true effect, that is, the posterior mean is not a consistent estimator of the true impulse response. Consequently, even with many observations, the posterior may occasionally place little mass on the true response. For the impulse responses themselves, this issue is relatively mild in the present example,

[5]For the reduced form, I adopt an uninformative prior. Given the large sample size, however, the choice of prior for the reduced form has no meaningful effect on the posterior.

since the true dynamics of the economy are still qualitatively well captured. However, Figure 2 shows that the problem becomes more severe when considering non-linear transformations of the impulse responses. The figure reports in blue the posterior distribution of the ORA, constructed from the posterior draws of the impulse responses above. The median[6] is negative, whereas the true ORA statistic is positive. Thus, even small inaccuracies in estimating the impulse responses can translate into misleading conclusions about policy-relevant statistics. Although substantial uncertainty remains, a researcher observing this evidence might conclude that the policymaker should have responded less, when in fact a stronger reaction would have been closer to the optimum.

Figure 3: Top panel: ORA statistic and Impulse Responses

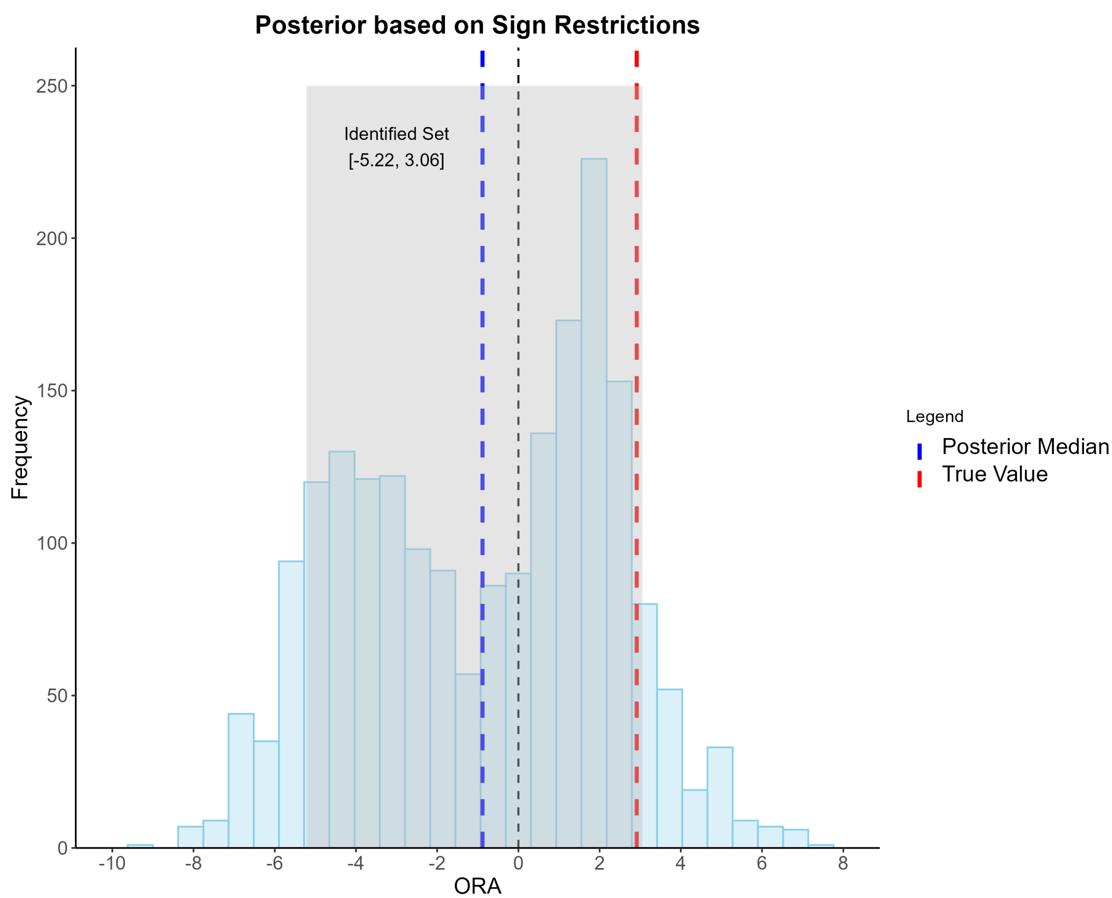

Note:

Figure 2 also reports the estimated identified set, shown in light grey, which represents

[6]The mean, not shown for simplicity, is also negative.

the smallest set of parameters consistent with the data, the VAR model, and the imposed sign restrictions. The following sections describe how this set is constructed and analyze its properties under several sufficient-statistic approaches. The key insight is that the set correctly reflects the ambiguity surrounding the true ORA statistic by encompassing both positive and negative values. Importantly, this set is obtained using the robust Bayes procedure of Giacomini and Kitagawa (2021), which considers a class of prior distributions that includes the Haar prior. Consequently, the posterior median based on the Haar prior always lies within the identified set. However, the fact that in this example the median happens to be centered within the set is purely coincidental.

It is well known that weak inequality restrictions are not very informative, typically yielding large identified sets for impulse responses, as in the example above. This has motivated a range of extensions, including augmenting sign restrictions with narrative information (Antolin-Diaz and Rubio-Ramirez (2018)), imposing restrictions on the coefficients of VAR equations (Arias, Caldara, and Rubio-Ramirez (2019)), and adding classical zero restrictions on short- or long-run impulse responses (see, e.g., the example in Giacomini and Kitagawa (2021)). These approaches have been successfully applied in practice and could likewise be used to shrink the identified set in the present example. However, as the following sections demonstrate, representing sufficient statistics directly in terms of the orthogonal reduced form of the VAR opens the door to a new class of sign restrictions applied directly to these statistics.

To illustrate, consider again the simple three-variable New Keynesian model described above, but now suppose the econometrician is only willing to impose standard sign restrictions:

$$\begin{pmatrix} u_t^x \\ u_t^\pi \\ u_t^i \end{pmatrix} = \begin{pmatrix} + & + & ? \\ + & - & - \\ ? & ? & + \end{pmatrix} \times \begin{pmatrix} \epsilon_t^D \\ \epsilon_t^S \\ \epsilon_t \end{pmatrix}$$

In particular, following Uhlig (2005), we do not impose that output must respond negatively to a contractionary monetary policy shock. However, suppose that, based on narrative evidence or economic reasoning, we are confident in imposing the condition that the policymaker, on average, underreacted to supply shocks during her term. This corresponds to the restriction ORA > 0. Thus, we add one additional (and in this simulation, correct) sign restriction.

Figure 4: Top panel: ORA statistic and Impulse Responses

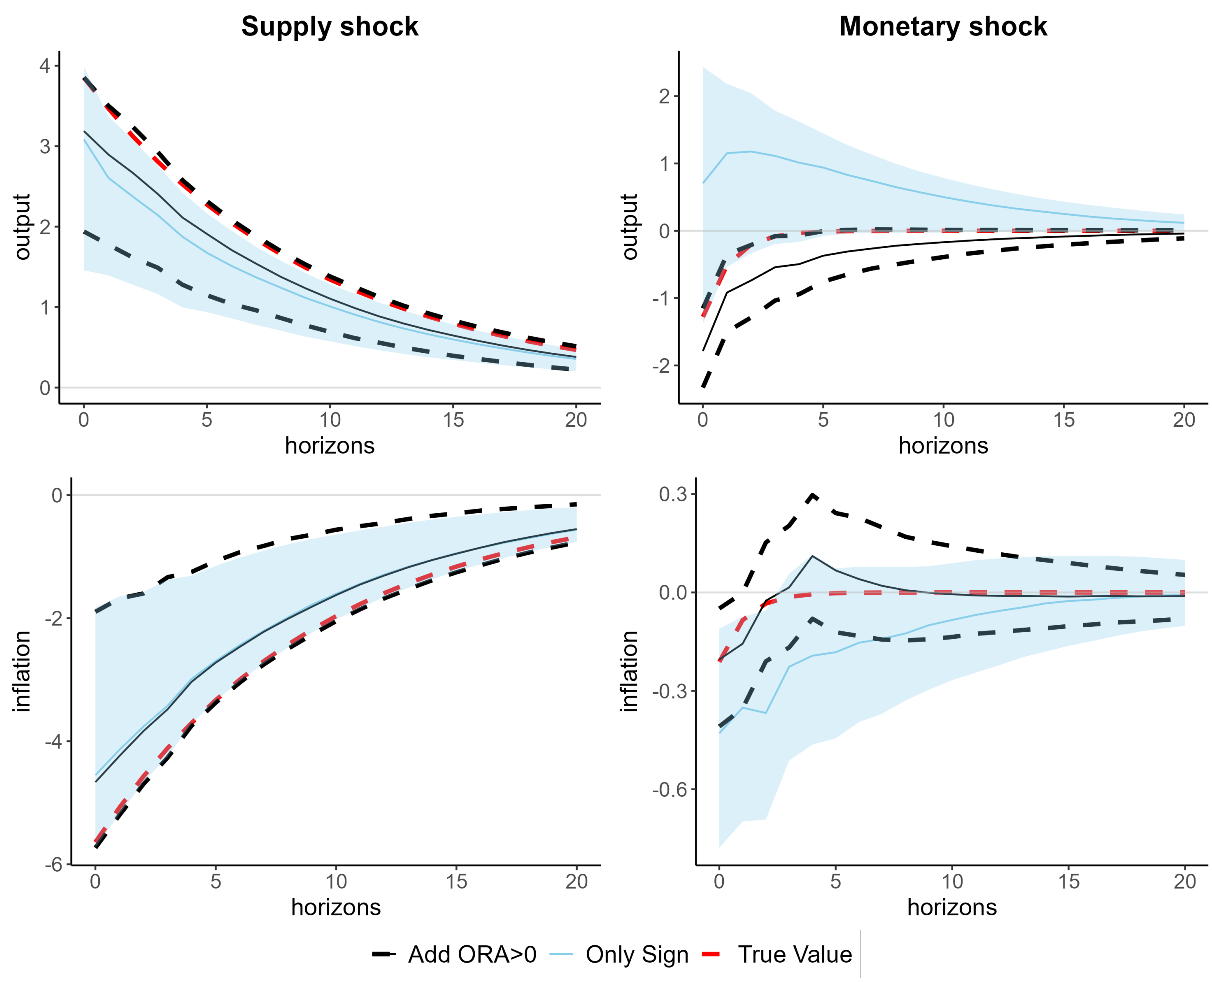

Note:

Figure 4 illustrates how adding a single additional restriction can improve the estimation of the impulse responses. The blue bands and solid line represent the credible intervals and posterior median based solely on classical sign restrictions, while the black dashed bands and solid line correspond to the same restrictions augmented with the condition that the ORA is positive. Focusing on the response of output to a monetary policy shock, the posterior based only on sign restrictions incorrectly assigns nearly all its mass to positive responses, a manifestation of the masquerading shock problem described by Wolf (2020). By contrast, once the ORA restriction is imposed, this problem disappears and the posterior aligns much more closely with the true impulse response.

The plausibility of such a restriction naturally depends on the context. However, as with the preceding discussion, the idea is not limited to the ORA statistic. For example, in Section 5 I apply the same logic to restrict the Phillips multiplier (Barnichon and Mesters (2021)), assuming the existence of an unemployment–inflation trade-off to sharpen the identification of monetary policy shocks. A related approach in the same section imposes a narrative restriction on the optimal policy perturbation (Barnichon and Mesters (2023a)) in April 2008, assuming that at the onset of the financial crisis the FED did not lower the interest rate enough. As in the simulation above, both types of restrictions substantially sharpen the identification of impulse responses and, by extension, of non-linear functions thereof.

4 Theory

For illustrative purposes thus far we focused on the ORA statistic which is one of many examples of a parameter which is defined as non-linear function of impulse responses. In the following we will give a general framework to study the theoretical properties of statistics of a more general class and give several specific examples which fit within this framework. The starting point is to express these statistics in terms of the VAR parameters directly, allowing to characterize the sharp identified set in terms of the latter.

A crucial requirement for both, theoretically justifying sign restrictions on non-linear functions of impulse responses, as well as studying the properties of a robust Bayes procedure, we need to ensure continuity of the functional of interest. As the next subsections shows, the parameters of interest considered can be represented as ratios of two continuous mappings, thus continuity depends on the denominator of this expression. A relevance assumption is imposed and discussed for several examples.

The second non-trivial extension of the framework in Giacomini and Kitagawa (2021) is that I consider cross-equation restrictions (in form of the regression in the impulse response space restrictions) and situation with sign restrictions on more than one column of the rotation matrices.[7] This alters the set of feasible orthogonal matrices, which in turn can yield non-convex identified sets and non-differentiable bounds thereof, both of which break the asymptotic equivalence between the robust Bayes approach of Giacomini and Kitagawa (2021) and frequentist approaches.

4.1 Representation in terms of VAR parameters

Consider a standard VAR(p) for the $n \times 1$ vector of observable variables y_t (all vectors in the following are column vectors):

$$y_t' A_0 = x_t' A_+ + \epsilon_t' \qquad \text{with} \quad \epsilon_t \sim N(0, I)$$

where $A_+ \equiv [A_1' \dots A_p' \quad c']$, $x_t' \equiv [y_{t-1}' \dots y_{t-p}' \quad 1]$ and c is a vector containing the coefficients on the intercept. We have $dim(A_+) = (1 + np) \times n$ and $dim(A_0) = n \times n$. This structural VAR has a reduced form given by:

$$y_t' = x_t' B + u_t' \qquad \text{where} \quad u_t' = \epsilon_t' A_0^{-1} \quad \Sigma = (A_0 A_0')^{-1} \quad B = A_+ A_0^{-1}$$

Given that the system of equations is Gaussian, it is well known that two structural parameter pairs (A_0, A_+) and $(\tilde{A}_0, \tilde{A}_+)$ are observational equivalent if and only if there exists $Q \in \mathcal{Q}$ such that $A_0 = \tilde{A}_0 Q$ and $A_+ = \tilde{A}_+ Q$. Where $\mathcal{Q}$ denotes the set of n-dimensional orthogonal matrices. This result motivates the orthogonalized reduced form represention of the model

[7] e.g. the ORA statistic above requires the identification of two shocks simultaneously. Further, in this case, we need to ensure that the restrictions yield convex domains for both columns jointly.

(see e.g. Arias and Waggoner (2018)):

$$y_t' = x_t' B + \epsilon_t' Q' h(\Sigma)$$

where $h(\Sigma)$ is any decomposition of Σ such that $h(\Sigma)'h(\Sigma) = \Sigma$. We collect the reduced form parameters in $\phi = (vec(B), vech(\Sigma))$ and note that $Q \in \mathcal{Q}$. We can learn ϕ consistently from the data, but identification of the structural parameters requires restricting $\mathcal{Q}$. Denote the restricted set by $\mathcal{Q}(\phi|F, S)$, where F and S are zero and sign restrictions, respectively, which are discussed in detail latter.

Let $\Phi_h(B, \Sigma)$ collect the reduced form impulse responses at horizon h. The structural impact matrix A_0 relates the reduced form with the structural impulse responses $\Gamma_h(A_0, A_+) = \Phi_h(B, \Sigma)(A_0^{-1})'$. Let us define $\gamma = vec(\Gamma_1, \dots \Gamma_H)$. The object of interest in this paper can be viewed as functions of this vector and has the general form:

$$\theta = \frac{\eta_1(\gamma, x)}{\eta_2(\gamma, x)}$$

where x denotes some external vector. Let us denote by $\gamma_h^{a,b}$ the subset of impulse responses collecting the response of variable a to shock b at horizon h, then several recent sufficient statistic approaches are nested within this framework:

Example 1. *Optimal Policy Adjustment (ORA): Let ϵ_p be a policy shock (e.g. monetary policy) and ϵ_{np} a non-policy shock (e.g. supply shock). Let y denote the output gap and π inflation. Define $\mathcal{R}_p \equiv (\gamma_{0:h}^{y,\epsilon_p}, \gamma_{0:h}^{\pi,\epsilon_p})$ and $\mathcal{R}_{np} \equiv (\gamma_{0:h}^{y,\epsilon_{np}}, \gamma_{0:h}^{\pi,\epsilon_{np}})$. Under a quadratic loss function weighting the output gap and inflation equally, Barnichon and Mesters (2023b) define the subset ORA θ as $\theta = (\mathcal{R}_p'\mathcal{R}_p)^{-1}\mathcal{R}_p'\mathcal{R}_{np}$, i.e. $\eta_1(\gamma, x) = \mathcal{R}_p'\mathcal{R}_{np}$ and $\eta_2(\gamma, x) = (\mathcal{R}_p'\mathcal{R}_p)^{-1}$.*

Example 2. *Optimal Policy Perturbation (OPP): Assuming the same loss function as in Example 1 above and denote by $Y = (y, \pi)$ Barnichon and Mesters (2023a) define the OPP $\theta = (\mathcal{R}_p'\mathcal{R}_p)^{-1}\mathcal{R}_p'\mathbb{E}[Y]$, i.e. $\eta_1(\gamma, x) = \mathcal{R}_p'\mathbb{E}[Y]$ and $\eta_2(\gamma, x) = (\mathcal{R}_p'\mathcal{R}_p)^{-1}$, with $x = \mathbb{E}[Y]$.*

Example 3. *Phillips-Multiplier / Fiscal Multiplier: Let* ϵ *be a monetary policy shock. Let* $\mathcal{R}_h^{\bar{\pi}} = \frac{1}{H}\sum_{j=0}^{H} \gamma_j^{\pi,\epsilon}$ *and* $\mathcal{R}_h^{\bar{u}} = \frac{1}{H}\sum_{j=0}^{H} \gamma_j^{u,\epsilon}$ *be the impulse responses of average inflation and unemployment to a one-unit policy shock. Barnichon and Mesters (2021) define the Phillips-Multiplier as* $\theta = \frac{\mathcal{R}_h^{\bar{\pi}}}{\mathcal{R}_h^{\bar{u}}}$*, i.e.* $\eta_1(\gamma, x) = \mathcal{R}_h^{\bar{\pi}}$ *and* $\eta_2(\gamma, x) = \mathcal{R}_h^{\bar{u}}$*. The fiscal multiplier of Ramey and Zubairy (2018) is structured similarly.*

Example 4. *Set-Identified Instruments: Let* $y_t = \theta' Y_t + u_t$ *be an endogenous macro equation with a scalar outcome* y_t *and a vector of covariates* Y_t*. Let* ϵ *be a monetary policy shock and let* $\gamma_{0:H}^{Y,\epsilon}$ *stacking the responses of all variables in* Y *to* ϵ *in a column vector. Barnichon and Mesters (2020) show that if appropriate economic shocks (or proxies thereof) are used as instruments for this equation,* θ *can be represented as a regression in the impulse response space:* $\theta = \left(\left(\gamma_{0:H}^{Y,\epsilon}\right)' \gamma_{0:H}^{Y,\epsilon}\right)^{-1} \left(\gamma_{0:H}^{Y,\epsilon}\right)' \gamma_{0:H}^{y,\epsilon}$ *where all responses are with respect to the appropriately chosen shock. Thus, we define* $\eta_1(\gamma, x) = \left(\gamma_{0:H}^{Y,\epsilon}\right)' \gamma_{0:H}^{y,\epsilon}, \quad$ *and* $\quad \eta_2(\gamma, x) = \left(\left(\gamma_{0:H}^{Y,\epsilon}\right)' \gamma_{0:H}^{Y,\epsilon}\right)^{-1}$.

Example 5. *Using shocks in a second stage: Often researchers use shocks (or proxies thereof) as inputs for a local projection, e.g.* $y_{t+h} = \theta\epsilon_t + e_{t+h}$*, where for simplicity I abstract from control variables and non-linearities. Then* $\theta = (\epsilon'\epsilon)^{-1}\epsilon' y$*. This does not exactly follow in the above framework, however, by expressing the structural errors in terms of the reduced form* $\epsilon_t' = u_t' A_0$*, we can define in a slight abuse of notation* $\eta_1(\gamma, x) = \epsilon' y$ *and* $\eta_2(\gamma, x) = \epsilon'\epsilon$

There is a one-to-one mapping between the structural impulse responses and the orthogonal reduced form parameters (ϕ, Q), thus equivalently we can represent the parameters of interest as functions of (ϕ, Q, x):

$$\theta = \tilde{\eta}(\phi, Q, x) = \frac{\tilde{\eta}_1(\phi, Q, x)}{\tilde{\eta}_2(\phi, Q, x)}$$

This allows us, together with the restrictions on Q to characterize the identified set, i.e. the set of parameters consistent with the data and all imposed assumptions, as:

$$\Theta_I \equiv \{\theta \in \mathbb{R} : \quad \theta = \tilde{\eta}(\phi_0, Q, x) \quad , \quad Q \in \mathcal{Q}(\phi_0 | F, S)\}$$

where ϕ_0 is the true value of ϕ.

A potential robust Bayes estimator of this set is given by the procedure described in Giacomini and Kitagawa (2021), however whether such an estimator exhibits the frequentist properties derived in the latter, depends on the properties of the objective function, i.e. the mappings $\tilde{\eta}_1$ as well as $\tilde{\eta}_2$ and on the properties of the constraint space $\mathcal{Q}(\phi|F, S)$. The latter are discussed in the following subsection.

In terms of objective, a key requirement for convergence of the estimator to the true identified set defined above as well as coverage properties is continuity of the functional of interest. The following assumption is a necessary condition to ensure continuity:

Assumption 1. *(Relevance)* $|\tilde{\eta}_2(\phi, Q, x)| > \kappa$ *for some, possibly data dependent, constant* κ, π_ϕ*-a.s.*

Remark 1. In examples 3 and 4 this is literally equivalent to the instrument relevance assumptions in the two papers, whereas in the remaining examples it ensures that there is *enough* variation in the explanatory variables.

4.2 Constraint Sets

In the cases where we only have zero or sign restrictions on one column, i.e. our parameter of interest is only a function of impulse responses to one shock (Examples 2-4 above), all the results of Giacomini and Kitagawa (2021) are applicable. This is because above basically only changes the parameter of interest from a scalar impulse response at horizon h, to a scalar-valued continuous function of such impulse responses.

However, for statistics which are functions of impulse responses of more than one shock (Example 1), the domain of the mapping $\tilde{\eta}(\phi, Q, x)$ is not necessarily path-connected and thus non-convex identified sets can arise. In these cases, the results of Giacomini and Kitagawa (2021) are only valid for the convex hull of the identified set. Thus, next I discuss this issue. Currently I am working on an algorithm checking whether a set of restrictions yields a path-connected domain.

In general as Giacomini and Kitagawa (2021) point out, if sign restrictions are imposed on more than one column, the identified set of impulse responses can be non-convex. First, let us reconsider their simple example.

$$y_t = \begin{pmatrix} y_{1,t} \\ y_{2,t} \end{pmatrix} = A_0^{-1} \begin{pmatrix} \epsilon_{1,t} \\ \epsilon_{2,t} \end{pmatrix} \qquad \Sigma_{\text{tr}} = \begin{pmatrix} \sigma_{11} & 0 \\ \sigma_{21} & \sigma_{22} \end{pmatrix} \qquad Q = \begin{pmatrix} q_{11} & q_{12} \\ q_{21} & q_{22} \end{pmatrix}$$

The impulse responses at horizon zero are therefore given by

$$IR^0 = \begin{pmatrix} ir_{11} & ir_{12} \\ ir_{21} & ir_{22} \end{pmatrix} = \begin{pmatrix} \sigma_{11}q_{11} & \sigma_{11}q_{12} \\ \sigma_{21}q_{11} + \sigma_{22}q_{21} & \sigma_{21}q_{12} + \sigma_{22}q_{22} \end{pmatrix}$$

Suppose we combine the economic restriction $ir_{12} \geq 0$ together with the sign normalizations

$$\sigma_{22}q_{11} - \sigma_{21}q_{21} \geq 0 \qquad \sigma_{11}q_{22} \geq 0$$

How does the identified set for column 1 look in this case, specifically what about the element ir_{11}?

The restrictions $\sigma_{11}q_{22} \geq 0$ and $ir_{12} = \sigma_{11}q_{12} \geq 0$ imply $q_2 = (q_{12}, q_{22}) \geq 0$ (elementwise), thus to make $q_1' q_2 = 0$ possible we need $\{q_1 : q_{11}q_{21} \leq 0\}$, i.e. $\{q_{11} \geq 0, q_{12} \leq 0\} \cup \{q_{11} \leq 0, q_{12} \geq 0\}$. Geometrically, this restricts the set to the upper-left and lower-right quadrant of the following figure.

Figure 5: Illustration Sign Restrictions

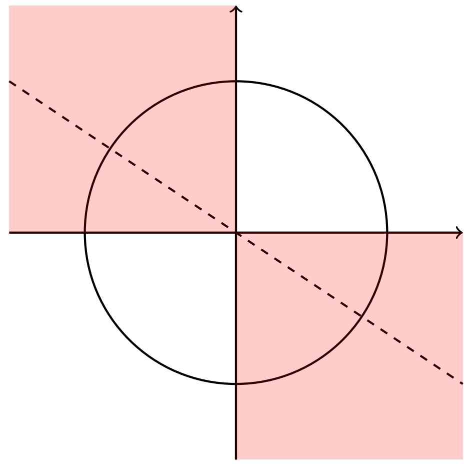

The light red areas are the feasible regions ignoring the $\sigma_{22}q_{11} - \sigma_{21}q_{21} \geq 0$, which defines the half space to the right of the dashed line. The non-convexity problem arises as the two arcs are separated, thus depending on the slope of the dashed line the resulting set are two disconnected arcs. If the slope $\frac{\sigma_{22}}{\sigma_{21}}$ is positive, the set is convex, if not its non-convex. As $\sigma_{21} \leq 0$ is of positive measure in the set of reduced form parameters we cannot rule out non-convex identified sets.

However, now assume the additional restriction on an regression in the impulse response space parameter, e.g. that $\frac{ir_{11}}{ir_{12}} \geq 0$, implying $q_{11}q_{12} \geq 0$, which reduces the feasible set to the lower right quadrant $\{q_{11} \geq 0, q_{12} \leq 0\}$. Thus now if the set is non-empty, its always convex. An even easier example when restrictions on two columns are imposed which can lead to convex sets is if instead of $ir_{12} \geq 0$ we assume $ir_{12} \leq 0$ which is not restricting the first column of Q to any subspace. Thus, restrictions on more than one column of Q can but do not have to yield non-convex domains.

Now suppose we restrict two columns of an n-dimensional Q. The above example suggest that what matters is how the sign restrictions on column 2 restrict the space of feasible q_1 signs. Excluding the origin, we have to make sure that the feasible space is path-connected. The following proposition gives a first attempt on an algorithm to check whether imposing restrictions on two columns yield to convex sets for their respective impulse responses. This are only necessary conditions, I am currently working on a more general version.

Proposition 1. *Fix a reduced form draw ϕ. Suppose columns q_i and q_j of Q are restricted by linear inequality restrictions encoded in the functions $S_{\cdot}(\phi)$ (and no zero restrictions are imposed.):*

1. *For all sign patterns (s_i, s_j) solve the following problem*

$$\begin{aligned}
&\min_{q_i, q_j} \; q_i' q_j \\
&S_i(\phi) q_i \geq 0 \quad , \quad S_j(\phi) q_j \geq 0 \\
&sign(q_i) = s_i \quad , \quad sign(q_j) = s_j \\
&||q_i|| = ||q_j|| = 1
\end{aligned}$$

   Similarly for $\max_{q_i, q_j} \; q_i' q_j$. Denote the respective objective values by p_{min} and p_{max}. If the problem's feasibility set is non-empty and $p_{min} \leq 0 \leq p_{max}$, keep sign pattern. Store all surviving sign patterns in $\mathcal{S} = (\mathcal{S}_i, \mathcal{S}_j)$, where $\mathcal{S}_{\cdot}$ are matrices with columns $s_{\cdot}$.

2. *Built a graph for $\mathcal{S}_i$ and $\mathcal{S}_j$ respectively. Each node is a sign pattern and nodes are connected if they share $n-1$ signs. Denote this graphs by G_i and G_j respectively. If G_i is connected, the domain of the continuous function mapping to the impulse response is path-connected thus the identified set for the impulse response is convex conditional on ϕ. Same for G_j.*

Remark 2. Above can be also used to just check whether one of the two yield an convex identified set and is thus of general interest.

Remark 3. If we cannot restrict some signs ex-ante we have $2^n \times 2^n$ sign patterns to check. For a six variable VAR this involves solving 2×4096 optimization problems for each reduced form draw.

Remark 4. If one of the two columns is point identified (i.e. through proxy variables) the above simplifies considerably. First, the set of possible sign patterns for the point identified shock reduces to one pattern. Second, the minimization problem reduces to a linear program.

Remark 5. Above only works if all constraints are homogeneous, e.g. they go through the origin, otherwise cones can be sliced into disjoint sets.

Remark 6. Above does not seem to work if one also includes zero restrictions, I have to develop this further. Given that Giacomini and Kitagawa (2021) already consider zero restrictions on multiple columns, there should be a possible combination that works. The problem is that they only consider the identified set for the response associated to one column!

4.3 Estimation

Theorem 1. *Consider* $\theta = \tilde{\eta}(\phi, Q, x)$ *subject to* $S(\phi, Q) \geq 0$ *and* $F(\phi, Q) = 0$ *and let* $\Theta_I(\phi_0)$ *be the true identified set. Let* $l(\phi) \equiv \inf_Q \{\tilde{\eta}(\phi, Q, x) : S(\phi, Q) \geq 0 \quad F(\phi, Q) = 0\}$ *and* $u(\phi) \equiv \sup_Q \{\tilde{\eta}(\phi, Q, x) : S(\phi, Q) \geq 0 \quad F(\phi, Q) = 0\}$*. Than* $\hat{\Theta}_I = [\bar{l}(\phi), \bar{u}(\phi)] \xrightarrow{p} \Theta_I(\phi_0)$ *if one of the following conditions hold:*

1. $\theta(\phi, Q) = \theta(\phi, q_j)$*,* $S(\phi, Q) = S(\phi)q_j$ *and the zero restrictions* $F(\phi, Q) = 0$ *follow the conditions in Giacomini and Kitagawa (2021)*

2. $\theta(\phi, Q) = \theta(\phi, q_j)$ *and* $S(\phi, Q) \geq 0$ *are such that Proposition 1 yields convex sets for column* j*.*

3. $\theta(\phi, Q) = \theta(\phi, [q_j, q_i])$ *and* $S(\phi, Q) \geq 0$ *are such that Proposition 1 yields convex sets for columns* j *and* i*.*

4.4 Imposing Restrictions on Sufficient Statistics

This section discusses whether we can impose sign and zero restrictions on the regression in the impulse response space parameters and sample from the associated constrained posterior distribution.

For sign restrictions only, notice that the theory in Arias and Waggoner (2018) applies to any continuous function $S(A_0, A_+)$, thus continuity of the regression in the impulse response space mappings is already sufficient for imposing sign restrictions on them.

For zero restrictions the following three conditions from Rubio-Ramirez, Waggoner, and Zha (2010) have to hold: 1.) $F(A_0Q, A_+Q) = F(A_0, A_+)Q$, 2.) $F(.)$ if its domain is open and $F(.)$ is continuously differentiable with $F'(.)$ of rank kn for all (A_0, A_+). and 3.) $F(\mathcal{A})$ is dense in the set of $k \times n$ matrices. Where for scalar regression in the impulse response space parameters $k = 1$. I am currently working on the verification of these conditions for the parameters of interest in examples 1-4 above.

Although, given a specific prior on the rotation matrix, above is already enough to justify the use of restrictions on impulse response space parameters, this could result in non-convex sets. The basic issue is that these class of restrictions, even if imposed only on parameters depending on a single column, are non-linear. In this cases Proposition 1 does not work as these non-linear restrictions can slice a cone into two disjoint sets.

Here I discuss results on two special cases used in the empirical applications below. First, the following ensures convexity of the identified set restricting the Multipliers of example 3:

Proposition 2. *Let $\phi = (B, \Sigma)$, $Q = [q_1, \dots, q_n]$ and $\Omega_h(\phi) = [\omega(\phi)_{ih}]_{i=1,\dots,n} \in \mathbb{R}^{n\times n}$. Consider restricting the sign of ratios of functions of impulse responses to one shock (e.g. example 3), i.e. $\theta_h(\phi) = \frac{h^{-1}\sum_{k=1}^{h}\omega'_{ik}q_1}{h^{-1}\sum_{k=1}^{h}\omega'_{jk}q_1}$. Let $C_h(\phi) = c'_j(\phi)c_i(\phi)$ with $c_l(\phi) = \left(\sum_{k=1}^{h}\omega_{1,lk}, \cdots \sum_{k=1}^{h}\omega_{n,lk}\right)$ and λ_{min} and λ_{max} be its smallest and largest eigenvalue. Than $\theta_h(\phi) \geq 0$ (respectively, $\theta_h(\phi) \leq 0$) iff $\lambda_{min} \geq 0$ (respectively, $\lambda_{max} \leq 0$).*

1. *$\theta_h(\phi) \geq 0$ (respectively, $\theta_h(\phi) \leq 0$) iff $\lambda_{min} \geq 0$ (respectively, $\lambda_{max} \leq 0$).*
2. *Let $\Gamma_h = [\gamma_{hk}]_{k=1,\dots,n}$ and consider additional sign restrictions q_1 of the form $S_1(\phi)q_1 \geq 0$. Then the identified set for γ_{h1} is convex and the identified set for θ_h is convex if $h^{-1}\sum_{k=1}^{h}\omega'_{jk}q_1 \neq 0$ π_ϕ almost surely.*

Proof. Consider $\theta_h(\phi) \geq 0$, the opposite case is proven analogously. Note

$$\theta_h(\phi) \geq 0 \qquad \iff \qquad \left(\sum_{k=1}^{h} \omega'_{ik} q_1\right)\left(\sum_{k=1}^{h} \omega'_{jk} q_1\right) \geq 0$$

Basic linear algebra shows this is equivalent to $q_1' C_h(\phi) q_1 \geq 0$. The condition on this quadratic form[8] to be positive is positive semi-definiteness of $C_h(\phi)$, i.e. $\lambda_{min} \geq 0$ because $||q_1|| = 1$. Note that the latter is a restriction which does not involve the rotation matrix. Thus, sampling ϕ conditional $\lambda_{min} \geq 0$ ensures $\theta_h(\phi) \geq 0$ for all Q. Convexity of the identified set than follows from the arguments in Giacomini and Kitagawa (2021) together with continuity of $\theta_h(\phi)$ under the conditions of this proposition. □

Above clarifies the statistical properties and justification of restricting these non-linear functions of impulse response, but what is the economic justification? In terms of the phillips-multiplier of Barnichon and Mesters (2021) restricting the statistic to be (weakly) positive is assuming that there is a trade-off between unemployment and inflation at some horizons. The assumption does not postulate the existence of a parametric phillips-curve but can be surely contested, especially at shorter horizons due to indeterminancies.

A second potentially interesting restrictions pertains the Optimal Policy Perturbation (OPP) statistic of Barnichon and Mesters (2023a). Economically, under certain conditions, the OPP characterizes the optimal adjustment to the baseline policy rule conditional on a loss function. Restricting the sign of such an assumption for certain historical episodes introduces a new class of narrative restrictions. For instance, a researcher might impose that the Federal Reserve lowered the short term interest too slow at the onset of the great recession. The following proposition characterizes such an assumption statistically and shows that any such sign restriction does not induce non-convexities to the identified sets.

Proposition 3. *Let $\phi = (B, \Sigma)$, $Q = [q_1, \ldots, q_n]$, and $\Omega_h(\phi) = [\omega_{ih}(\phi)]_{i=1,\ldots,n} \in \mathbb{R}^{n \times n}$. Let the parameter of interest be the optimal policy perturbation $\theta = (\mathcal{R}_p' \mathcal{R}_p)^{-1} \mathcal{R}_p' \mathbb{E}[Y]$, where*

[8] Note, $C_h(\phi)$ is not symmetric but due to standard arguments we can consider $\tilde{C}_h(\phi) = (C_h(\phi) + C_h'(\phi))/2$

$\mathbb{E}[Y] = (y_1, \dots, y_h, \pi_1, \dots, \pi_h)$. *Define the* $1 \times n$ *vector*

$$P_{y\pi} = y_1\,\omega'_{y1} + \cdots + y_h\,\omega'_{yh} + \pi_1\,\omega'_{\pi 1} + \cdots + \pi_h\,\omega'_{\pi h} \tag{1}$$

Then $\theta_h(\phi) \geq 0$ *(respectively,* $\theta_h(\phi) \leq 0$*) if and only if* $P_{y\pi}q_1 \geq 0$ *(respectively,* $P_{y\pi}q_1 \leq 0$*). Imposing this restriction on top of additional sign restrictions on* q_1*, the resulting identified sets for both* γ_{h1} *and* θ_h *remain convex.*

Proof. Follows trivially from the condition $\mathcal{R}'_p\mathbb{E}[Y] \geq 0$. □

Remark 7. Inference on the restricted OPP in this case might be not of primary interest, however it sharpens identification of the impulse response itself and thus on OPP for other periods in question. For instance, a policy institution might use insights on past OPP to estimate the current OPP in an online fashion (see Einarsson (2024)). Specifically, let $P_{y\pi}^{\text{April,2008}}$ denote the OPP related matrix above for April 2008. In retrospect we might feel confident in claiming that the Federal Reserve underreacted to the onset of the great recession. Imposing this condition sharpens identification of the associated impulse responses, thus reduces uncertainty while estimating the current OPP.

5 Empirical Application

This section presents two applications. The first demonstrates how the new class of sign restrictions on sufficient statistics can sharpen the identification of monetary policy shocks. Specifically, I compare a set of sign restrictions with a proxy BVAR in the context of the classic study by Gertler and Karadi (2015). The aim is to highlight the value of the proposed approach, even when the sufficient statistic itself is not the primary object of interest.

The second application integrates the methods discussed above to conduct inference using only sign restrictions on the optimal policy perturbation from Barnichon and Mesters (2023a) for the US, the set of weak sign restrictions allows (arguably) correct inference on the sign of the optimal policy perturbation in nearly all periods.

5.1 Restricting Phillips-Multiplier

The seminal study by Gertler and Karadi (2015) employs high-frequency monetary policy instruments to examine the transmission of monetary policy shocks across various specifications. Here, I focus on their simple VAR, which includes log industrial production, the log consumer price index, the one-year government bond rate as the policy indicator, and the GZ excess bond premium.

We consider three specifications. The first replicates the results of Gertler and Karadi (2015) by using their proxy variable within a proxy BVAR, estimated as described in Arias and Waggoner (2021) with uninformative priors. The second specification imposes classical Uhlig (2005)-type sign restrictions to identify the monetary policy shock, assuming that the monetary policy shock is the only shock affecting the interest rate and inflation in opposite directions. The third specification augments these sign restrictions with the additional condition that the Phillips multiplier of Barnichon and Mesters (2021), defined using the responses of industrial production and the consumer price index, is positive at all horizons, that is, that a trade-off between output and inflation exists in the economy. The latter two models are estimated using the algorithm in Arias and Waggoner (2018).

Figure 5 shows the results for all three models in the respective columns. The replicated proxy BVAR results display a delayed negative response of inflation and a more immediate decline in industrial production, while the excess bond premium rises on impact and dissipates after about 12 months. Imposing only classical Uhlig (2005) sign restrictions yields no significant effects in most periods, except for an immediate drop in inflation. Most notably, consistent with prior literature, the effect on output, industrial production in this example, is ambiguous at all horizons, implying that output could even increase following a contractionary monetary policy shock. This ambiguity is resolved once the Phillips-multiplier restriction is added, as shown in the final column. Inflation and industrial production both decline on impact, and the estimated trajectory for industrial production closely tracks the proxy BVAR results. The effect on the excess bond premium is similarly well captured. Overall, assuming the existence of an inflation–output trade-off appears sufficient to qualitatively reproduce the dynamics obtained via high-frequency identification.

Figure 6: Impulse Responses

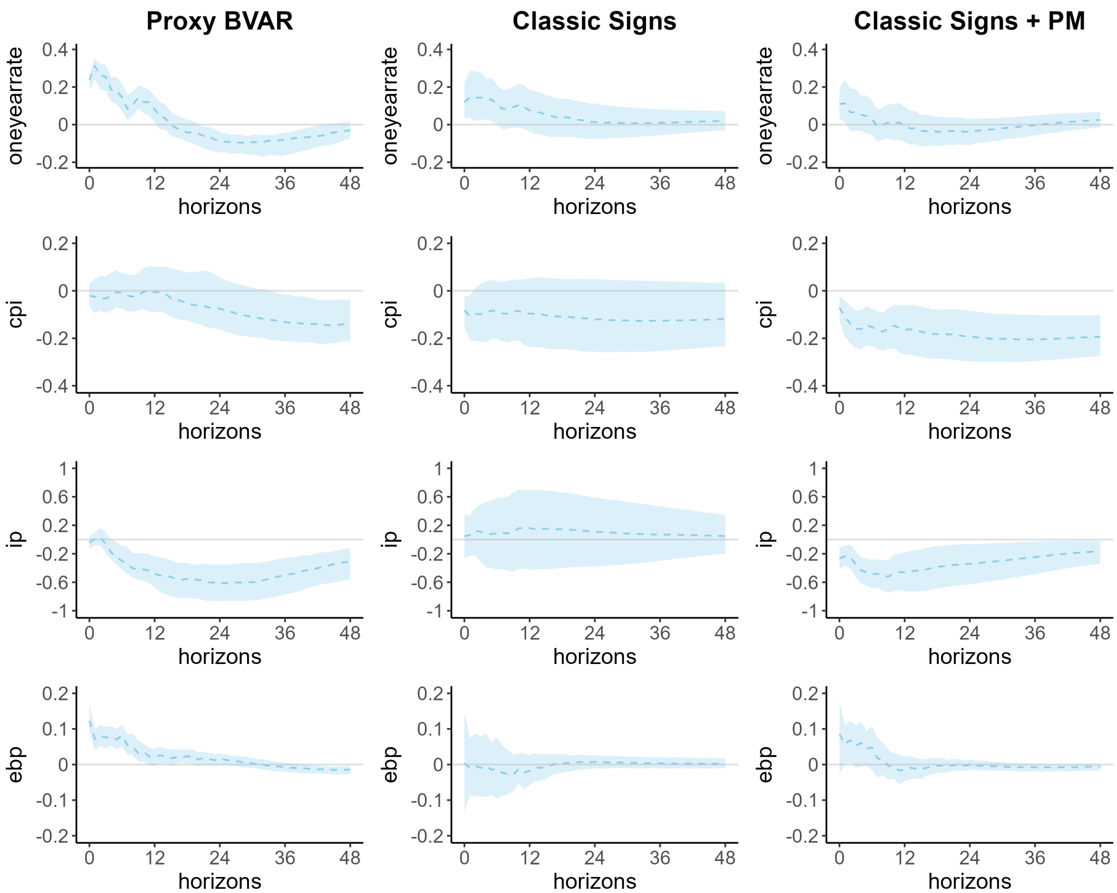

Note:

An important caveat is that, for illustrative purposes, the results above are conditional on the Haar prior for the orthogonal matrices. As discussed, this prior can have a substantial impact on the impulse responses, particularly for non-linear functions thereof. Nonetheless, its use for the impulse responses themselves can be justified by Arias and Waggoner (2024), who show that the Haar prior is the only prior that induces a uniform joint distribution over the space of all impulse responses.[9] Reporting results conditional on this prior is also

[9]Interestingly, this result implies that the prior on non-linear transformations of the impulse responses is

practically useful, as it remains widely employed in applied work. Figure 6 illustrates the same impulse responses, comparing estimators of the identified set, i.e., the region between the means of the upper and lower bounds, for the classical sign restriction case (blue area) and the case with the Phillips-multiplier restriction added (dashed black lines).

Figure 7: Identified Sets under Classic and PM Restrictions

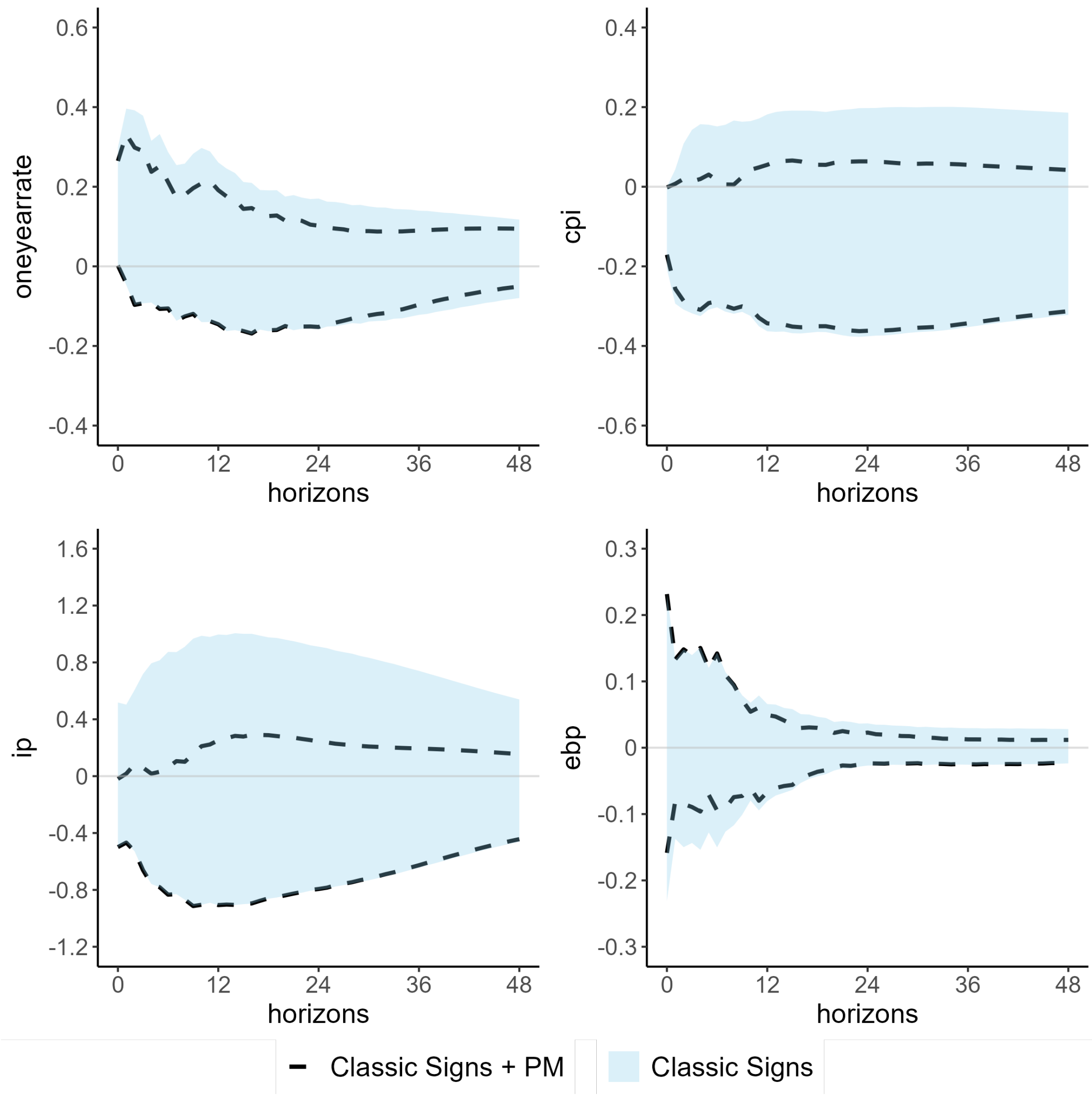

Note:

For the log consumer price index and the log of industrial production, the Phillips-multiplier restrictions provide substantial identification power: positive responses are largely ruled

not uniform. Moreover, a uniform prior is not necessarily uninformative, so it is not clear whether such a prior is desirable in general.

out and any remaining responses are modest. This suggests that the new restrictions are informative independently of the chosen prior on the orthogonal matrices. For the excess bond premium, however, the robust estimators indicate that the improvement observed in the previous figure arises from the combined effect of the Haar prior and the Phillips-multiplier restriction.

5.2 Optimal Monetary Policy US

This section integrates the previous discussion to estimate the short-rate optimal policy perturbation from Barnichon and Mesters (2023a) for the US, using only sign restrictions. Details on the dataset are provided in Barnichon and Mesters (2023a). I estimate a VAR with the same three variables—excluding the shock series, namely inflation, unemployment, and the federal funds rate. Both the classical sign restrictions and the Phillips-multiplier restriction are imposed as before, with the key difference that the Phillips multiplier is now defined using unemployment and is therefore restricted to be negative.

To further sharpen identification, we apply the new class of narrative restrictions to a specific period for the optimal policy perturbation. In particular, the OPP in April 2008 is restricted to be negative, reflecting the assumption that the Federal Reserve set the interest rate too high at the onset of the financial crisis. For a discussion of why this assumption might be reasonable based on narrative evidence, see Barnichon and Mesters (2023a).

Figure 7 shows the estimated identified set of the OPP from 1990 to 2022 based on these sign restrictions (blue), while the black line replicates the results from Barnichon and Mesters (2023a), ignoring the zero lower bound.[10] Consequently, results for 2008–2015 should be interpreted as theoretical adjustments. Nevertheless, in all periods the identified set in blue correctly captures the sign of the optimal policy perturbation relative to the replication mean. Although substantial uncertainty remains regarding the magnitude, this illustrates that even weak sign restrictions allow a researcher to infer the direction of the adjustment needed for

[10] I am working on estimating the constrained OPP that respects the zero lower bound; this requires studying the properties of a new target statistic, which has no closed-form solution.

optimal monetary policy.

Figure 8: Impulse Responses

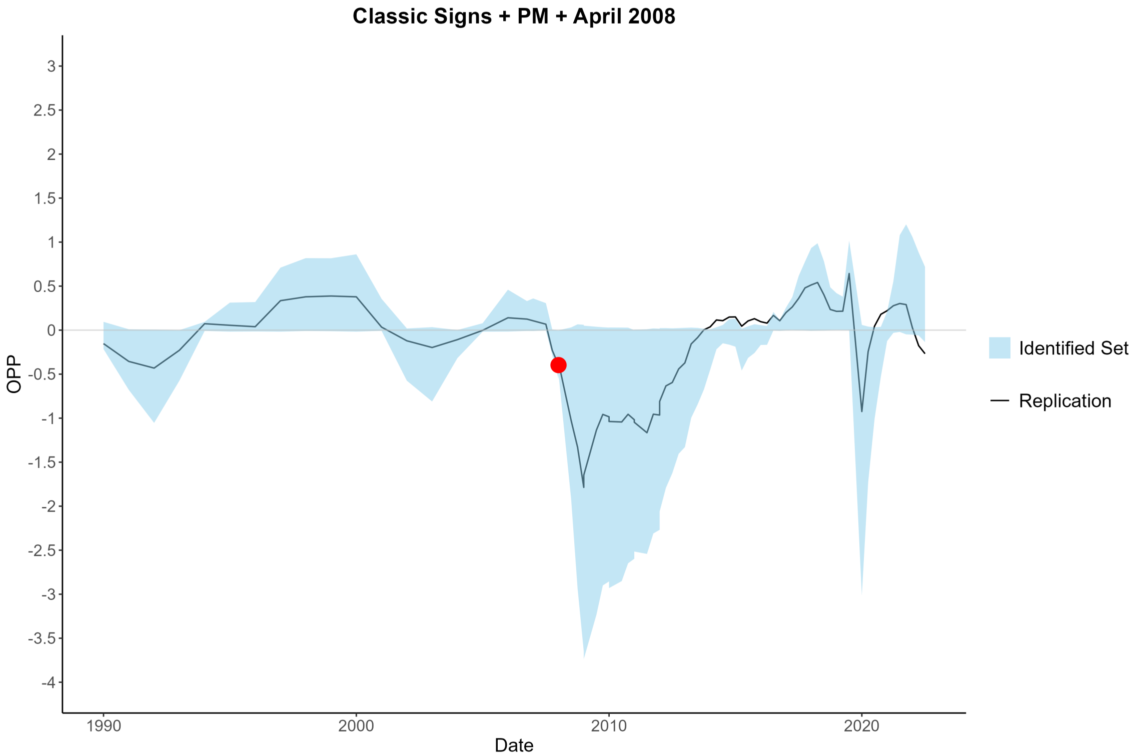

Note:

Finally, returning to the main point of this paper, Figure 8 illustrates that using the Haar prior on the VAR rotation matrices to compute non-linear statistics can severely distort inference. Here, the OPP is estimated using only classical sign restrictions. The blue area represents the robust identified set, while the black line shows the original results from Barnichon and Mesters (2023a). The blue dashed line indicates the posterior median based solely on the Haar prior. In nearly every period, this estimate disagrees with the replication and, for example, would incorrectly suggest that the Federal Reserve should have raised the interest rate just before the financial crisis.

Figure 9: Impulse Responses

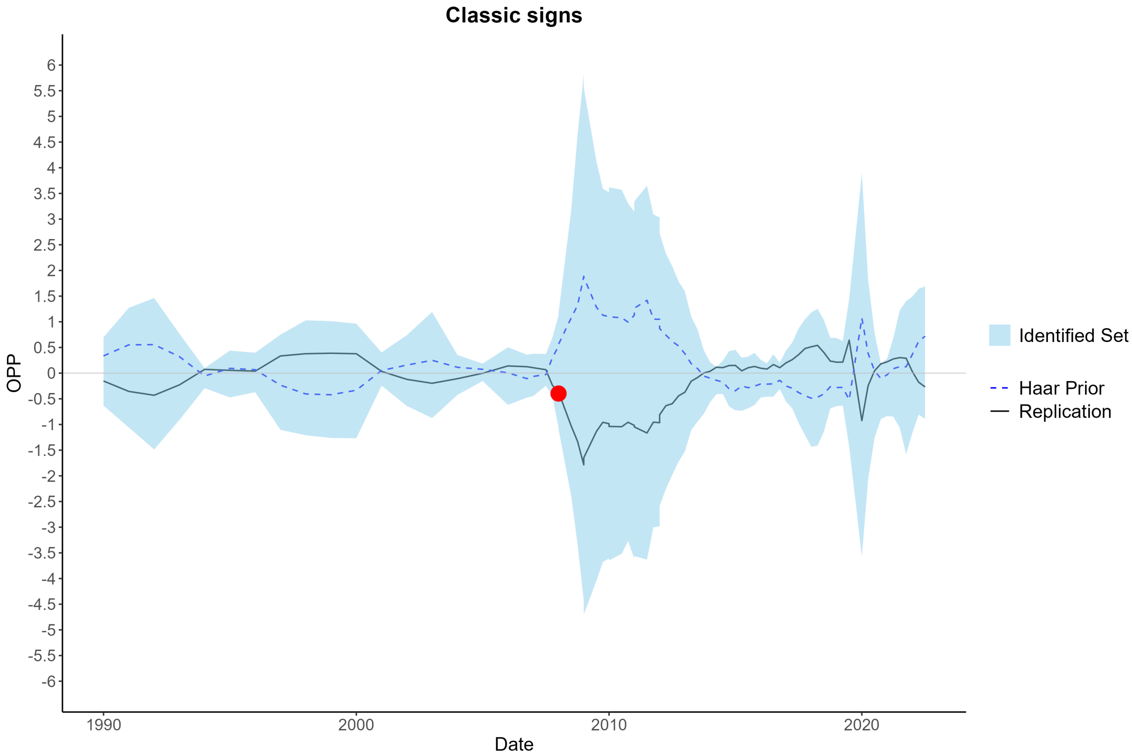

Note:

6 Conclusion

This paper demonstrates that widely used non-linear sufficient statistics of impulse responses, such as RIRS coefficients, the Phillips multiplier, and optimal-policy perturbations, can be highly sensitive to the prior under set identification. Simple (single prior) posterior summaries, such as medians based on the Haar prior, may even reverse signs and mislead policy inference, despite the underlying impulse responses appearing reasonable. Building on a mapping from the orthogonalized reduced form to the sufficient statistics and under transparent relevance and continuity conditions, I adapt robust Bayes methods to deliver sharp identified sets for these statistics. Crucially, I also introduce a new class of identification restrictions applied directly to the statistics themselves. Under tractable convexity and path-connectedness conditions, these restrictions preserve a well-behaved geometry and restore the frequentist validity of the robust envelope.

Two applications illustrate the practical value of this approach. In the Gertler and Karadi (2015) setting, augmenting classical sign restrictions with a Phillips-multiplier sign restriction sharpens identification and reproduces the qualitative dynamics of a proxy-SVAR without relying on external instruments. In a US policy exercise, combining the Phillips-multiplier restriction with a narrative OPP sign in April 2008 produces identified sets that track the direction of optimal rate adjustments from 1990 to 2022, while highlighting how Haar-based medians alone could provide misleading guidance. Future work could build on these insights in an online monitoring framework, providing real-time feedback for policymakers, as in Einarsson (2024).

References

Antolin-Diaz, J. and J. Rubio-Ramirez (2018). "Narrative Sign Restrictions for SVARs". In: *The American Economic Review* 108.10, pp. 2802–2829.

Arias, J., D. Caldara, and J. Rubio-Ramirez (2019). "The Systematic Component of Monetary Policy in SVARs: An Agnostic Identication Procedure". In: *Journal of Monetary Economics* 101, pp. 1–13.

Arias J. Rubio-Ramirez, J. and D. Waggoner (2018). "Inference based on Structural Vector Autoregressions identified with Sign and Zero Restrictions: Theory and Applications". In: *Econometrica* 86, pp. 685–720.

— (2021). "Inference in Bayesian Proxy-SVARs". In: *Journal of Econometrics* 225, pp. 88–106.

— (2024). "Uniform Priors for Impulse Responses". In: *forthcoming.*

Aucejo, E., F. Bugni, and J. Hotz (2017). "Identification and Inference on Regressions with Missing Covariate Data". In: *Econometric Theory* 33, pp. 196–241.

Barnichon, R. and G. Mesters (2020). “Identifying Modern Macro Equations with Old Shocks”. In: *Quarterly Journal of Economics* 2255-2298.

— (2021). “The Phillips multiplier”. In: *Journal of Monetary Economics* 117, pp. 689–705.

— (2023a). “A Sufficient Statistic Approach for Macro Policy”. In: *American Economic Review* 113.11.

— (2023b). “Evaluating Policy Institutions - 150 Years of US Monetary Policy”. In.

— (2024). “Policy Evaluation with Sufficient Macro Statistics”. In.

Baumeister, C. and J. Hamilton (2015). “Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information”. In: *Econometrica* 83.5, pp. 1963–1999.

Beresteanu, A., I. Molchanov, and F. Molinari (2011). “Sharp Identification Regions in Models with Convex Moment Predictions”. In: *Econometrica* 79.6, pp. 1785–1821.

Einarsson, B. (2024). “Online Monitoring of Policy Optimality”. In.

Faust, J. (1998). “The robustness of identified VAR conclusions about money”. In: *Carnegie-Rochester Conference Series on Public Policy* 49, pp. 207–244.

Gertler, M. and P. Karadi (2015). “Monetary Policy Surprises, Credit Costs, and Economic Activity”. In: *American Economic Journal: Macroeocnomics* 7.1, pp. 44–76.

Giacomini, R. and T. Kitagawa (2021). “Robust Bayesian Inference for Set-Identified Models”. In: *Econometrica* 89, pp. 1519–1556.

Horowitz, J. et al. (2003). “Computation of Bounds on Population Parameters When the Data Are Incomplete”. In: *Reliable Computing* 9, pp. 419–440.

Lewis, D. and K. Mertens (2022). “Dynamic Identification Using System Projections and Instrumental Variables”. In.

Manski, C. and E. Tamer (2002). “Inference on Regressions with Interval Data on a Regressor or Outcome”. In: *Econometrica* 70, pp. 519–546.

McKay, A. and C. Wolf (2023). “What can time-series regressions tell us about policy counterfactuals”. In: *Econometrica* 91.5.

Moon, H. and F. Schorfheide (2012). “Bayesian and Frequentist Inference in Partially Identified Models”. In: *Econometrica* 80.2, pp. 755–782.

Ramey, V. and S. Zubairy (2018). “Government spending multipliers in good times and in bad: evidence from U.S. historical data.” In: *Journal of Political Economy* 126.

Rubio-Ramirez, J., D. Waggoner, and T. Zha (2010). “Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference”. In: *The Review of Economic Studies* 77.2.

Uhlig, H. (2005). “What are the effects of monetary policy on output? Results from an agnostic identification procedure”. In: *Journal of Monetary Economics* 52, pp. 381–419.

Wolf, C. (2020). “SVAR (Mis)Identification and the Real Effects of Monetary Policy Shocks”. In: *American Economic Journal: Macroeconomics* 12.4, pp. 1–32.