

Uncertain Identification*

Raffaella Giacomini[†], Toru Kitagawa[‡] and Alessio Volpicella[§]

This draft: April 2017

Abstract

Uncertainty about the choice of identifying assumptions is common in causal studies, but is often ignored in empirical practice. This paper considers uncertainty over models that impose different identifying assumptions, which, in general, leads to a mix of point- and set-identified models. We propose performing inference in the presence of such uncertainty by generalizing Bayesian model averaging. The method considers multiple posteriors for the set-identified models and combines them with a single posterior for models that are either point-identified or that impose non-dogmatic assumptions. The output is a set of posteriors (*post-averaging ambiguous belief*) that are mixtures of the single posterior and any element of the class of multiple posteriors, with weights equal to the posterior model probabilities. In practice, we suggest reporting the range of posterior means and the associated credible region, and provide a simple algorithm to compute them. We establish that the prior model probabilities are updated when the models are “distinguishable” and/or they specify different priors for reduced-form parameters, and characterize the asymptotic behavior of the posterior model probabilities. The method provides a formal framework for conducting sensitivity analysis of empirical findings to the choice of identifying assumptions. In a standard monetary model, for example, we show that one would need to attach a prior probability greater than 0.89 to the validity of the assumption that prices do not react contemporaneously to a contractionary monetary policy shock in order to obtain a negative output response to such a shock.

Keywords: Partial Identification, Sensitivity Analysis, Model Averaging, Bayesian Robustness, Ambiguity.

*We would like to thank Frank Kleibergen, José Montiel-Olea, John Pepper, James Stock, and seminar participants at Monash University, University of Bonn, University of North Carolina, and University of Virginia for valuable comments. We also thank participants at IAAE 2016 (Milan), RES 2016 (Brighton), AMES 2016 (Kyoto), MMF 2016 (Bath), 2016 University of Leicester Conference, EWM of Econometric Society 2016 (Edinburgh) and the workshop of “New Identification Approaches to Macroeconometrics” (Oxford) for beneficial discussions. Financial support from the ESRC through the ESRC Centre for Microdata Methods and Practice (CeMMAP) (grant number RES-589-28-0001) is gratefully acknowledged. Alessio Volpicella acknowledges the financial support of the School of Economics and Finance, Queen Mary University of London and from the “Ermenegildo Zegna Founder’s Scholarship”.

[†]University College London, Department of Economics. Email: r.giacomini@ucl.ac.uk

[‡]University College London, Department of Economics. Email: t.kitagawa@ucl.ac.uk

[§]Queen Mary University of London, School of Economics and Finance. Email: a.volpicella@qmul.ac.uk

1 Introduction

The choice of identifying assumptions is the crucial step that allows researchers to draw causal inferences using observational data. This is often a controversial choice, and there can be uncertainty about which assumptions to impose from a menu of plausible ones, but this uncertainty and its effects on inference have been typically ignored in empirical work. We propose a formal method for drawing inferences about causal effects in the presence of uncertainty about identifying assumptions, which we characterize as uncertainty over a class of models that impose different sets of assumptions. The method can be viewed as a generalization of Bayesian model averaging to include set-identified models, which commonly arise when the assumptions are under-identifying or take the form of inequality restrictions.

There are several examples in economics where empirical researchers face uncertainty about identifying assumptions that lead to point- or set-identification of a common causal estimand of interest. The first is macroeconomic policy analysis based on structural vector autoregressions (SVARs), where common identifying assumptions include causal ordering restrictions (Bernanke (1986) and Sims (1980)), long-run neutrality restrictions (Blanchard and Quah (1993)), and Bayesian prior mean restrictions implied by a structural model (Del Negro and Schorfheide (2004)). Subsets of these assumptions deliver set-identified impulse-responses, as does the use of sign restrictions (Canova and Nicolo (2002), Faust (1998), and Uhlig (2005)). The second example is microeconomic causal effect studies with identifying assumptions such as selection on observables (Ashenfelter (1978) and Rosenbaum and Rubin (1983)), selection on observables and unobservables (Altonji, Elder, and Taber (2005)), exclusion and monotonicity restrictions in instrumental variables methods (Imbens and Angrist (1994), yielding set-identification of the average treatment effect), and monotone instrument assumptions (Manski and Pepper (2000), also yielding set-identification). The third example is missing data with identifying assumptions such as missing at random, Bayesian imputation (Rubin (1987)), and unknown missing mechanism (Manski (1989), yielding set-identification). Finally, estimation of structural models with multiple equilibria relies on assumptions about the equilibrium selection rule, with different assumptions (or lack thereof) delivering point- or set-identified models (e.g., Bajari, Hong, and Ryan (2010), Beresteanu, Molchanov, and Molinari (2011), and Ciliberto and Tamer (2009)).

The common practice in empirical work is to report results based on what is deemed the most credible set of identifying assumptions, or, sometimes, based on a number of alternative assumptions, viewed as an informal sensitivity analysis. Our proposed method provides a formal framework for conducting sensitivity analysis and for aggregating results based on point- and set-identifying assumptions.

The idea of model averaging has a long history in econometrics and statistics since the pioneering works of Bates and Granger (1969) and Leamer (1978). The literature has considered

Bayesian approaches (see, e.g., Hoeting, Madigan, Raftery, and Volinsky (1999) and Claeskens and Hjort (2008)), frequentist approaches (Hansen (2007, 2014), Hjort and Claeskens (2003), Hansen and Racine (2012), Liu (2015), Liu and Okui (2013), and Zhang and Liang (2011)), and hybrid approaches (Hjort and Claeskens (2003), Kitagawa and Muris (2016), and Magnus, Powell, and Prüfer (2010)), but none of them allows for set-identification in any candidate models.

We tackle this problem from the angle of Bayesian model averaging. Standard Bayesian model averaging delivers a single posterior that is a mixture of the posteriors of the candidate models with weights equal to the posterior model probabilities. This approach could in principle be extended to our context if one could obtain a single posterior for every set-identified model. Assuming a single prior is however problematic from the robustness viewpoint as the choice of a single prior, even an apparently uninformative one, can lead to spuriously informative posterior inference for the object of interest. The severity of the problem is magnified by the fact that the effect of the prior choice persists asymptotically, unlike in the case of point-identified models (Baumeister and Hamilton (2015), Moon and Schorfheide (2012), Poirier (1998), among others).

The key innovation of our approach to Bayesian model averaging is that we do not assume availability of a single posterior for the set-identified models. Rather, we allow for multiple priors (*an ambiguous belief*) within the set-identified models, and then combine the corresponding multiple posteriors with single posteriors for models that are either point-identified or that impose non-dogmatic identifying assumptions in the form of a Bayesian prior for the structural parameters (as in (Baumeister and Hamilton (2015))). The output of the procedure is a set of posteriors (*post-averaging ambiguous belief*), that are mixtures of the single posteriors and any element of the set of multiple posteriors, with weights equal to the posterior model probabilities. To summarize and visualize the post-averaging ambiguous belief, we recommend reporting the range of posterior quantities (e.g., the mean or median) and the associated credible region (an interval to which any posterior in the class assigns a certain credibility level). We show that these quantities have analytically simple expressions and are easy to compute in practice.

This paper contributes to the growing literature on Bayesian inference for partially identified models (Giacomini and Kitagawa (2015), Kitagawa (2012), Kline and Tamer (2016), Moon and Schorfheide (2012), Norets and Tang (2014), Liao and Simoni (2013)). We follow the multiple-prior approach to model the lack of knowledge within the identified set as in Giacomini and Kitagawa (2015) and Kitagawa (2012). When the set-identified model is the only model considered, the range of posteriors generated by the approach leads to the posterior inference for the identified set proposed in Kline and Tamer (2016), Liao and Simoni (2013), and Moon and Schorfheide (2011). When there is uncertainty about the identifying assumptions, however, the usual definition of identified set is not available without conditioning on the model. The multiple prior viewpoint has an advantage in this case since the range of posteriors has a

well-defined subjective interpretation even in the presence of model uncertainty.

The method proposed in this paper provides a formal framework for conducting sensitivity analysis of causal inferences to the choice of identifying assumptions. First, when the set-identified model nests the point-identified model, the method can be used to assess the posterior sensitivity in the point-identified model with respect to perturbations of the prior in the direction of relaxing some of the point-identifying assumptions. In this case, we can formally interpret our averaging method as an example of the ϵ -contamination sensitivity analysis developed in Huber (1973) and Berger and Berliner (1986) with a particular construction of the prior class. Second, if the point-identified model can be considered a reasonable benchmark, the method offers a simple and flexible way to add non-dogmatic identifying information to the set-identified model, which results in increasing informativeness of the conclusions in a transparent manner. Third, the method can be used to perform reverse engineering exercises that compute the prior probability one would need to attach to a set of identifying assumptions in order for the averaging to preserve a given empirical conclusion (e.g., the so-called price and liquidity puzzles in monetary SVARs, respectively discussed by (Sims, 1992) and (Reichenstein, 1987)).

Our proposed method can also be viewed as bridging the gap between point- and set-identification. When focusing solely on a point-identified model, a researcher who is not fully confident about the choice of identifying assumptions may doubt the robustness of the conclusions. On the other hand, discarding some of the point-identifying assumptions and reporting estimates of the identified set may appear “excessively agnostic”, and often results in uninformative conclusions. Our averaging procedure reconciles these two extreme representations of the posterior beliefs by exploiting the prior weights that one can assign to alternative sets of identifying assumptions. The output of the procedure is a weighted average of the posterior mean in the point-identified model and the range of posterior means in the set-identified model. When the identified set is a connected interval, the range of posterior means can be viewed as an estimate of the identified set (Giacomini and Kitagawa (2015)), and thus our averaging procedure effectively shrinks the identified set estimate toward the point estimate from the point-identified model, with the degree of shrinkage governed by the posterior model probabilities.

In addition to our novel proposal of averaging, this paper makes other analytical contributes to the literature of Bayesian model selection and averaging. First, we clarify under which conditions the prior model probabilities can be updated by data. We show that the updating occurs if some models are “distinguishable” for some distribution of data and/or the priors for the reduced-form parameters differ across models. Second, we investigate the asymptotic properties of the posterior model probabilities and our averaging method. We show that when only one model is consistent with the true distribution of the data our method asymptotically

assigns probability one to it. When multiple models are observationally equivalent and “not falsified” at the true data generating process, the posterior model probabilities asymptotically assign nontrivial weights to them. We clarify what part of the prior input determines the asymptotic posterior model probabilities in such case. The consistency property of Bayesian model selection has been well-studied in the statistics literature (see, e.g., Claeskens and Hjort (2008) and references therein), but little is known about the asymptotic behavior of posterior model probabilities when the models differ in terms of the identifying assumptions but can be observationally equivalent in terms of their reduced form representations. These new results therefore could be of separate interest.

The empirical application in this paper considers SVAR analysis with uncertainty over the classes of identifying assumptions typically used in empirical work: causal ordering restrictions (Bernanke (1986) and Sims (1980)), sign restrictions (Canova and Nicolo (2002), Faust (1998), and Uhlig (2005)), and restrictions implied by a Dynamic Stochastic General Equilibrium (DSGE) model. The choice of identifying assumptions has often been a source of controversy in this literature, given that researchers have differing opinions about their credibility. One popular choice is the use of sign restrictions. Although the resulting model is set-identified and the approach therefore raises serious robustness concern as we discussed above, the common practice is to consider single-prior Bayesian inference in set-identified SVARs. The large body of the empirical literature adopting this approach includes Canova and Nicolo (2002), Faust (1998), Mountford (2005), Rafiq and Mallick (2008), Scholl and Uhlig (2008), Uhlig (2005), and Vargas-Silva (2008) for applications to monetary policy, Dedola and Neri (2007), Fujita (2011), and Peersman and Straub (2009) for applications to business cycle model, Mountford and Uhlig (2009) for applications to fiscal policy, Kilian and Murphy (2012) for applications to oil prices. Alternative approaches that do not suffer from the pitfalls of single-prior Bayesian inference are Moon, Schorfheide, and Granziera (2013) and Gafarov, Meier, and Montiel-Olea (2016a,b), who consider frequentist inference for the identified set and Giacomini and Kitagawa (2015), who propose a robust Bayesian approach. To our knowledge, little work has been done on multi-model inference in the SVAR literature, and the methods proposed in this paper could therefore prove helpful in reconciling the controversies about the identifying assumptions that are widespread in this literature. As an example, the empirical application documents the high sensitivity of the conclusion in standard monetary SVARs that output decreases after a contractionary monetary policy shock to the choice of identifying assumptions.

The remainder of the paper is organized as follows. Section 2 illustrates the motivation and the implementation of the averaging method in the context of a simple model. Section 3 presents the formal analysis in a general framework and provides a computational algorithm to implement the procedure. Section 4 discusses the relationship between our method and existing Bayesian methods, and discusses elicitation of model probabilities. Section 5 applies

our method to impulse response analysis in monetary SVARs. The Appendix contains proofs and a microeconomic application.

2 Illustrative Example

We present the key ideas and the implementation of the method in a static model of labor supply and demand, subject to common types of identifying assumptions.¹ The model is:

$$A \begin{pmatrix} \Delta n_t \\ \Delta w_t \end{pmatrix} = \begin{pmatrix} \epsilon_t^d \\ \epsilon_t^s \end{pmatrix}, \quad A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}, \quad t=1, \dots, T, \quad (2.1)$$

where $(\Delta n_t, \Delta w_t)$ are the growth rates of employment and wages and $(\epsilon_t^d, \epsilon_t^s)$ is an i.i.d. normally distributed vector of demand and supply shocks with variance-covariance the identity matrix. A is the structural parameter and the contemporaneous impulse responses are elements of A^{-1} .

The reduced-form model is indexed by Σ , the variance-covariance matrix of $(\Delta n_t, \Delta w_t)$, which satisfies $\Sigma = A^{-1}(A^{-1})'$. Denote its lower triangular Cholesky decomposition with nonnegative diagonal elements by $\Sigma_{tr} = \begin{pmatrix} \sigma_{11} & 0 \\ \sigma_{21} & \sigma_{22} \end{pmatrix}$ with $\sigma_{11} \geq 0$ and $\sigma_{22} \geq 0$, and define the reduced form parameter as $\phi = (\sigma_{11}, \sigma_{21}, \sigma_{22}) \in \Phi = \mathbb{R}_+ \times \mathbb{R} \times \mathbb{R}_+$.² Let the mapping from the structural parameter to the reduced-form parameter be denoted by $\phi = g(A)$.

Suppose the object of interest is the response of the first variable to a unit positive shock in the first variable, $\alpha \equiv (1,1)$ -element of A^{-1} . Without identifying assumptions, the structural parameter is set-identified since knowledge of the reduced-form parameter ϕ cannot uniquely pin down the structural parameter ($\phi = g(A)$ is a many-to-one mapping). Imposing assumptions can lead to a set or a point for α , depending on the type and number of assumptions.

We now illustrate our proposal for two different types of identifying assumptions.

2.1 Dogmatic Identifying Assumptions

First consider dogmatic identifying assumptions, which are equality or inequality restrictions on (functions of) the structural parameter that hold with probability one.

Scenario 1: Candidate Models

- *Model M^p (point-identified)*: The labor demand is inelastic to wage, $a_{12} = 0$.
- *Model M^s (set-identified)*: The wage elasticity of demand is non-positive, $a_{12} \geq 0$, and the wage elasticity of supply is non-negative, $a_{21} \leq 0$.

¹See Appendix A.2 for a microeconomic application to a treatment effect model with noncompliance.

²The positive semidefiniteness of Σ does not constrain the value of ϕ other than $\sigma_{11} \geq 0$ and $\sigma_{22} \geq 0$.

Model M^p restricts A to be lower-triangular, as in the classical causal ordering assumptions of Sims (1980) and Bernanke (1986). Combined with the sign normalization restrictions requiring the diagonal elements of A to be nonnegative, the assumption implies that the contemporaneous impulse responses can be identified by $A^{-1} = \Sigma_{tr}$. The parameter of interest can be expressed as $\alpha = \alpha_{M^p}(\phi) \equiv \sigma_{11}$.

Model M^s imposes sign restrictions that only set-identify α . Appendix A shows that the identified set for α is:

$$IS_\alpha(\phi) \equiv \begin{cases} \left[\sigma_{11} \cos \left(\arctan \left(\frac{\sigma_{22}}{\sigma_{21}} \right) \right), \sigma_{11} \right], & \text{for } \sigma_{21} > 0, \\ \left[0, \sigma_{11} \cos \left(\arctan \left(-\frac{\sigma_{21}}{\sigma_{22}} \right) \right) \right], & \text{for } \sigma_{21} \leq 0. \end{cases} \quad (2.2)$$

Note that the identified set is non-empty for any ϕ . Hence, models M^p and M^s are observationally equivalent at any $\phi \in \Phi$ and neither of them is falsifiable, i.e., for any $\phi \in \Phi$ in both models there exist a structural parameter A that satisfies the identifying assumptions.³

Our method specifies a prior for the reduced-form parameter in each model. This prior is updated by the data and thus such a choice does not asymptotically affect the conclusions about the parameter of interest within a given model. However, as we show in Section 3.5, the choice of priors for the reduced-form parameter can influence the posterior model probabilities, even asymptotically. In the example, given the observational equivalence of the two models, it might be reasonable to specify the same prior for ϕ :

$$\pi_{\phi|M^p} = \pi_{\phi|M^s} = \tilde{\pi}_\phi \quad (2.3)$$

where $\tilde{\pi}_\phi$ is a *proper* prior, such as the one induced by a Wishart prior on Σ . The same prior for ϕ in observationally equivalent models leads to the same posterior:

$$\pi_{\phi|M^p,Y} = \pi_{\phi|M^s,Y} = \tilde{\pi}_{\phi|Y}, \quad (2.4)$$

where Y denotes the sample.

In model M^p , the posterior for ϕ implies the unique posterior for α , $\pi_{\alpha|M^p,Y}$, via the mapping $\alpha = \alpha_{M^p}(\phi)$.

In model M^s , on the other hand, the posterior for ϕ does not yield a unique posterior for α , since the mapping in (2.2) is generally set-valued. Following Giacomini and Kitagawa (2015) and Kitagawa (2012), we formulate the lack of prior knowledge by considering multiple priors (ambiguous belief). Formally, given the prior for the reduced form parameter $\pi_{\phi|M^s}$, we form

³When $\sigma_{21} > 0$, the point-identified α in model M^p is the upper-bound of the identified set in model M^s , whereas when $\sigma_{21} < 0$, the identified set in model M^s does not contain the point-identified α . This is because in model M^p we have $a_{12} = -\frac{\sigma_{21}}{\sigma_{11}\sigma_{22}}$, which is positive if $\sigma_{21} < 0$, meaning that the point-identifying assumptions $a_{12} = 0$ and $\sigma_{21} < 0$ are not compatible with the restriction $a_{21} \leq 0$.

the class of priors for A by admitting arbitrary conditional priors for A given ϕ , as long as they are consistent with the identifying assumptions:

$$\Pi_{A|M^s} \equiv \left\{ \pi_{A|M^s} = \int_{\Phi} \pi_{A|M^s, \phi} d\pi_{\phi|M^s} : \pi_{A|M^s, \phi}(\mathcal{A}_{sign} \cap g^{-1}(\phi)) = 1, \pi_{\phi|M^s}\text{-a.s.} \right\},$$

where $\mathcal{A}_{sign} = \{A : a_{12} \geq 0, a_{21} \leq 0, \text{diag}(A) \geq 0\}$ is the set of structural parameters that satisfy the sign restrictions and the sign normalizations and $g^{-1}(\phi)$ is the set of observationally equivalent structural parameters given the reduced-form parameter ϕ .

Since the likelihood depends on the structural parameter only through the reduced-form parameter, applying Bayes' rule to each prior in the class only updates the prior for ϕ , and thus leads to the following class of posteriors for A :

$$\Pi_{A|M^s, Y} \equiv \left\{ \pi_{A|M^s, Y} = \int_{\Phi} \pi_{A|M^s, \phi} d\pi_{\phi|M^s, Y} : \pi_{A|M^s, \phi}(\mathcal{A}_{sign} \cap g^{-1}(\phi)) = 1, \pi_{\phi|M^s}\text{-a.s.} \right\}. \quad (2.5)$$

Marginalizing the posteriors in $\Pi_{A|M^s, Y}$ to α leads to the class of α -posteriors:

$$\Pi_{\alpha|M^s, Y} \equiv \left\{ \pi_{\alpha|M^s, Y} = \int_{\tilde{\Phi}} \pi_{\alpha|M^s, \phi} d\pi_{\phi|M^s, Y} : \pi_{\alpha|M^s, \phi}(IS_{\alpha}(\phi)) = 1, \pi_{\phi|M^s}\text{-a.s.} \right\}. \quad (2.6)$$

We view this class as a representation of the posterior uncertainty about α in the set-identified model. The class contains any α -posterior that assigns probability one to the identified set, and it represents the lack of belief therein in terms of Knightian uncertainty (ambiguity). This is a key departure from the standard approach to Bayesian model averaging, which requires a single posterior for all models, including those where the parameter is set-identified.

Suppose that the researcher's prior uncertainty over the two models can be represented by prior probabilities $\pi_{M^p} \in [0, 1]$ for model M^p and $(1 - \pi_{M^p})$ for model M^s .⁴

Our proposal is to combine the single posterior for α in model M^p and the set of posteriors for α in model M^s according to the posterior model probabilities $\pi_{M^p|Y}$ and $\pi_{M^s|Y}$ (the posterior model probability for model M^s depends only on the single prior for the reduced-form parameter, so it is unique in spite of the multiple priors for the structural parameter). The combination delivers a class of posteriors $\Pi_{\alpha|Y}$, the *post-averaging ambiguous belief*:

$$\Pi_{\alpha|Y} = \{\pi_{\alpha|M^p, Y} \pi_{M^p|Y} + \pi_{\alpha|M^s, Y} \pi_{M^s|Y} : \pi_{\alpha|M^s, Y} \in \Pi_{\alpha|M^s, Y}\}. \quad (2.7)$$

As we show in Section 4.1, our proposal can be interpreted as applying Bayes' rule to each prior in a class that has the form of an ϵ -contaminated class of priors (see Berger and Berliner (1986)).

⁴We discuss interpretation and elicitation of the prior model probabilities in Section 4.3.

A key result of the paper is to establish conditions under which the prior model probabilities are updated by the data, which we show occurs when the models are “distinguishable” for some reduced-form parameter values and/or they specify different priors for ϕ (see Lemma 3.1 below). In the current scenario, the two models are indistinguishable, so the prior model probabilities are not updated if they use a common ϕ -prior.

In practice, we recommend reporting as the output of the procedure the post-averaging range of posterior means or quantiles of $\Pi_{\alpha|Y}$ and its associated *robust credible region* with credibility $\gamma \in (0, 1)$, defined as the shortest interval that receives posterior probability at least γ for every posterior in $\Pi_{\alpha|Y}$. Proposition 3.1 shows that the range of posterior means is the weighted average of the posterior mean in model M^p and the range of posterior means in model M^s :

$$\begin{aligned} & \left[\inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha), \sup_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha) \right] \\ &= \pi_{M^p|Y} E_{\alpha|M^p,Y}(\alpha) + \pi_{M^s|Y} [E_{\phi|M^s,Y}(l(\phi)), E_{\phi|M^s,Y}(u(\phi))], \end{aligned} \quad (2.8)$$

where $(l(\phi), u(\phi))$ are the lower and upper bounds of the nonempty identified set for α shown in (2.2), $a + b[c, d]$ stands for $[a + bc, a + bd]$, and $E_{\phi|M^s,Y}(\cdot)$ denotes the posterior mean with respect to $\pi_{\phi|M^s,Y} = \tilde{\pi}_{\phi|Y}$. Since the range of posterior means can be viewed as an estimator for the identified set in model M^s , our procedure effectively shrinks the estimate of the identified set in the set-identified model toward the point estimate in the point-identified model, with the amount of shrinkage determined by the posterior model probabilities.

The robust credible region for α with credibility γ can be computed as follows. We first draw z_1, \dots, z_G randomly from a Bernoulli distribution with mean $\pi_{M^p|Y}$ and then generate $g = 1, \dots, G$ random draws of the “mixture identified set” for α according to

$$IS_{\alpha}^{mix}(\phi_g) = \begin{cases} \{\alpha(\phi_g)\}, & \phi_g \sim \pi_{\phi|M^p,Y} = \tilde{\pi}_{\phi|Y}, & \text{if } z_g = 1 \\ [l(\phi_g), u(\phi_g)], & \phi_g \sim \pi_{\phi|M^s,Y} = \tilde{\pi}_{\phi|Y} & \text{if } z_g = 0. \end{cases} \quad (2.9)$$

Intuitively, with probability $\pi_{M^p|Y}$, a draw of the mixture identified set is a singleton corresponding to the point-identified value of α , and with probability $\pi_{M^s|Y}$ it is a nonempty identified set for α . The robust credible region with credibility level γ is approximated by an interval that contains the γ -fraction of the drawn $IS_{\alpha}^{mix}(\phi)$'s. The minimization problem in Step 5 of Algorithm 4.1 in Giacomini and Kitagawa (2015) is solved to obtain the shortest-width robust credible region.

2.2 Non-dogmatic Identifying Assumptions

Our method allows for identifying assumptions that are expressed as a non-dogmatic prior for the structural parameter.

Scenario 2: Candidate Models

- *Model M^B (single prior)*: A prior for the structural parameter A .
- *Model M^s (multiple priors)*: Same as the set-identified model in Scenario 1.

Model M^B assumes availability of a prior for the whole structural parameter. This prior can reflect Bayesian probabilistic uncertainty about identifying assumptions expressed as equalities (see, e.g., Baumeister and Hamilton (2015), who propose a prior for a dynamic version of the current model based on a meta-analysis of the literature). Another key example of a model that implies a single prior for the structural parameter is a Bayesian DSGE model.

Model M^B always yields a single posterior for α . However, the influence of prior choice does not vanish asymptotically due to the lack of identification. In principle, if the researcher were confident about the prior specification in model M^B , she could perform standard Bayesian inference and obtain a credible posterior, despite the identification issues. In practice, this is rather rare. For instance, the prior considered by Baumeister and Hamilton (2015) is based on the elicitation of first and second moments and the remaining characteristics of the distribution are chosen for analytical or computational convenience. Further, eliciting dependence among structural parameters is challenging, and an independent prior could lead to unintended or counter-intuitive effects on posterior inference.⁵ These robustness concerns can be addressed by averaging the Bayesian model M^B with the set-identified model M^s , which accommodates the lack of prior knowledge about the structural parameter (beyond the inequality restrictions).

One important consideration in this scenario is that the single prior for A in model M^B implies a single prior for the reduced form parameter. Here we thus allow the prior for ϕ in model M^s to differ from that in model M^B . This, in turn, affects the posterior model probabilities, which are given by:

$$\begin{aligned}\pi_{M^B|Y} &= \frac{p(Y|M^B) \cdot \pi_{M^B}}{p(Y|M^B) \cdot \pi_{M^B} + p(Y|M^s) \cdot (1 - \pi_{M^B})}, \\ \pi_{M^s|Y} &= \frac{p(Y|M^s) \cdot (1 - \pi_{M^B})}{p(Y|M^B) \cdot \pi_{M^B} + p(Y|M^s) \cdot (1 - \pi_{M^B})},\end{aligned}\tag{2.10}$$

where π_{M^B} is the prior weight assigned to model M^B , $p(Y|M) \equiv \int_{\Phi} p(Y|\phi, M) d\pi_{\phi|M}(\phi)$, $M = M^B, M^s$, are the marginal likelihoods of model M with $p(Y|\phi, M)$ the likelihood of the reduced form parameters. In this scenario the different priors for ϕ imply $p(Y|M^B) \neq p(Y|M^s)$, and therefore the prior model probabilities can be updated by the data.

Given these posterior model probabilities, the construction of the post-averaging ambiguous belief proceeds as in (2.7). The range of posterior means for α can be obtained similarly to

⁵“Knowing no dependence” among the parameters differs from “not knowing their dependence.”

(2.8), where M^B replaces M^P . The robust credible region can be constructed as in Scenario 1, by drawing iid draws $z_1, \dots, z_G \sim \text{Bernoulli}(\pi_{M^B|Y})$ and letting

$$IS_{\alpha,g}^{mix} = \begin{cases} \{\alpha\}, & \alpha \sim \pi_{\alpha|M^B,Y}, & \text{if } z_g = 1, \\ [l(\phi_g), u(\phi_g)], & \phi_g \sim \pi_{\phi|M^S,Y} & \text{if } z_g = 0. \end{cases} \quad (2.11)$$

3 Formal Analysis

This section formalizes the idea in a general setting and proves the analytical claims made in the previous section.

3.1 Notation and Definitions

Consider $J + K \geq 2$ candidate models, $J, K \geq 0$, that can differ in various aspects, including the identifying assumptions and the parametrization of the structural model. The class of J models consists of *single-posterior models*, whose prior input always (i.e., independent of the realization of the data) leads to a single posterior for the parameter of interest. Examples are models that impose dogmatic point-identifying assumptions with a single prior for the reduced-form parameter (such as model M^P in Scenario 1), or models that assume a single prior for the structural parameter in spite of it being set-identified (such as model M^B in Scenario 2). We denote the class of single-posterior models by \mathcal{M}^P .

The class of K models consists of *multiple-posterior models*, defined by the following features: (1) under the identifying assumptions the parameter of interest is set-identified, i.e., knowledge of the distribution of observables (value of the reduced-form parameter) does not pin down a unique value for the parameter of interest, and (2) they specify a single prior for the reduced-form parameter. The posterior information in a multiple-posterior model is characterized by the set of posteriors. We denote the class of multiple-posterior models by \mathcal{M}^S .

Let $\mathcal{M} \equiv \mathcal{M}^P \cup \mathcal{M}^S$. The vector of structural parameters in model $M \in \mathcal{M}$ is $\theta_M \in \Theta_M$, where Θ_M is the set of structural parameters that satisfy the identifying assumptions imposed in model M . We assume that the scalar parameter of interest $\alpha = \alpha_M(\theta_M) \in \mathbb{R}$ is well-defined as a function of θ_M and it carries a common (causal) interpretation in all models. The reduced-form parameter ϕ_M is a function of the structural parameter, $\phi_M = g_M(\theta_M) \in \mathbb{R}^{d_M}$, where $g_M(\cdot)$ maps a set of observationally equivalent structural parameters subject to the identifying assumptions in model M to a point in the reduced-form parameter space, defined as $\Phi_M = g_M(\Theta_M)$.⁶ As reflected in the notation, our most general set-up allows the parameter

⁶The likelihood $\tilde{p}(Y|\theta_M, M)$ in model M depends on θ_M only through the reduced-form parameter $g_M(\theta_M)$ for any realization of Y , i.e., there exists $p(Y|\cdot, M)$ such that $\tilde{p}(Y|\theta_M, M) = p(Y|g_M(\theta_M), M)$ holds for every Y . The statistics literature refers to the reduced-form parameter as the minimally sufficient parameter (see, e.g., Dawid (1979)).

space of both structural and reduced-form parameters to differ across models.⁷ We express the likelihood in model $M \in \mathcal{M}$ in terms of the reduced-form parameter by $p(Y|\phi_M, M)$. For a multiple-posterior model $M \in \mathcal{M}^s$, define the identified set of α by $IS_\alpha(\phi_M|M) = \{\alpha_M(\theta_M) : \theta_M \in \Theta_M \cap g_M^{-1}(\phi_M)\}$, which is a set-valued mapping from Φ_M to \mathbb{R} .

Note that, by construction, the parameter space of the reduced form parameter Φ_M incorporates the testable implications, if any, of the imposed identifying assumptions. For a set-identified model $M^s \in \mathcal{M}^s$, Φ_{M^s} is equivalent to the set of ϕ_M 's that yield a nonempty identified set, $\Phi_{M^s} = \{\phi_{M^s} \in \mathbb{R}^{d_{M^s}} : IS_\alpha(\phi_{M^s}|M^s) \neq \emptyset\}$.⁸

The next definition introduces the concept of identical reduced-forms among the candidate models. Our analytical claims on the posterior model probabilities shown below (Lemma 3.1 and Proposition 3.3) assumes some or all of the candidate models admit an identical reduced-form.

Definition 3.1 *Let \mathcal{M} be a collection of models. \mathcal{M} admits an identical reduced-form if the following conditions hold:*

- (a) Φ_M can be embedded into a common d -dimensional Euclidean space \mathbb{R}^d for all $M \in \mathcal{M}$ (hence ϕ_M can be denoted by $\phi \in \mathbb{R}^d$).
- (b) For every $M \in \mathcal{M}$, the reduced-form likelihood $p(Y|\phi_M = \phi, M)$ defines a probability distribution of Y on the extended domain $\Phi \equiv \cup_{M \in \mathcal{M}} \Phi_M$, and $p(Y|\phi_M = \phi, M) = p(Y|\phi)$ holds for all $\phi \in \Phi$, where $p(Y|\phi)$ is the likelihood common among $M \in \mathcal{M}$.

Definition 3.1 formalizes the situation where models imposing different identifying assumptions lead to the same parametric family of distributions for the observables (Condition (a)). Different identifying assumptions, nonetheless, can constrain the class of distributions of observables in the sense that the domain of reduced-form parameters Φ_M can differ among the models. The key condition in Definition 3.1 is (b), requiring that the distribution of the data Y in model M (indexed by ϕ) is well-defined over the extended domain $\Phi = \cup_{M \in \mathcal{M}} \Phi_M$ and the likelihood of ϕ is common among the models $M \in \mathcal{M}$. For instance, if \mathcal{M} consists of SVAR models with the same set of variables but subject to different identifying assumptions (including observationally restrictive ones such as sign restrictions), the conditions of Definition 3.1 are satisfied when the models have the same reduced-form VAR. See also the treatment effect models of Appendix A.2 as a microeconometrics example where all the candidate models

⁷For instance, in the model considered in Section 2, the reduced-form parameter space can differ depending on how many lagged endogenous variables and/or exogenous variables are included in each model.

⁸For instance, in a SVAR with observationally restrictive sign restrictions, Φ_M is the set of reduced-form parameters in the VAR yielding a nonempty impulse response identified set, which can be a proper subset of the reduced-form parameter space of the VAR.

admit an identical reduced-form. In what follows, whenever we assume \mathcal{M} admits an identical reduced-form, we notate the common reduced-form parameters by ϕ and the common reduced-form likelihood by $p(Y|\phi)$.

The next set of definitions defines the concepts of observational equivalence and distinguishability of the candidate models.

Definition 3.2 (i) *The models in \mathcal{M} are **observationally equivalent** at ϕ if \mathcal{M} admits an identical reduced-form and $\phi \in \cap_{M \in \mathcal{M}} \Phi_M$.*

(ii) *Two distinct models $M, M' \in \mathcal{M}$ that admit an identical reduced-form are **distinguishable** if $\Phi_M \neq \Phi_{M'}$.*

(iii) *The models in \mathcal{M} are **indistinguishable** if \mathcal{M} admits an identical reduced-form and $\Phi_M = \Phi$ for all $M \in \mathcal{M}$.*

Models that are observationally equivalent at ϕ (Definition 3.2 (i)) generate the same distribution of data (corresponding to ϕ), implying that knowledge of ϕ fails to uniquely identify what model generated the data. Note that our definition of observational equivalence is local to the given ϕ , and it does not constrain the relationship among their reduced-form parameter spaces except that they must have a non-empty intersection. In contrast, the concept of (in)distinguishability in Definition 3.2 (ii) and (iii) concerns relationships of the reduced-form parameter spaces across the models. If two models admitting an identical reduced-form are distinguishable, then there exists some reduced-form parameter value that allow us to falsify one model in support of the other. On the other hand, indistinguishability of Definition 3.1 (iii) can be interpreted as observational equivalence of the models in a global sense — if the models are indistinguishable, we are never able to support one model than the others based on the data independent of what knowledge we have for the distribution of observables.

3.2 Posterior Model Probabilities

This section shows when and how the data update the prior model probabilities when some or all of the candidate models admit an identical reduced form.

Let $(\pi_M : M \in \mathcal{M})$, $\sum_{M \in \mathcal{M}} \pi_M = 1$, be prior probabilities assigned over \mathcal{M} . By Bayes' rule, the posterior model probability for each model in the class is

$$\pi_{M|Y} = \frac{p(Y|M)\pi_M}{\sum_{M' \in \mathcal{M}} p(Y|M')\pi_{M'}}. \quad (3.1)$$

By the definition of reduced-form parameters, the value of the likelihood depends on θ_M only through ϕ_M , for which we assume a single prior. This implies that the marginal likelihood depends only on the ϕ_M -prior so it can be computed uniquely for all models including multiple-posterior models since every $M \in \mathcal{M}$ assumes a single prior for ϕ_M .

In situations where the models admit an identical reduced-form, we can simplify the expression of the posterior model probabilities, as shown in the next lemma.

Lemma 3.1 (i) Suppose that the multiple-posterior models $M^s \in \mathcal{M}^s$ admit an identical reduced-form with reduced-form parameters $\phi \in \Phi = \cup_{M^s \in \mathcal{M}^s} \Phi_{M^s} \subset \mathbb{R}^d$. Let $\tilde{\pi}_\phi$ be a proper prior on Φ and assume that $\tilde{\pi}_\phi(\Phi_{M^s}) = \tilde{\pi}_\phi(IS_\alpha(\phi|M^s) \neq \emptyset) > 0$ holds for all $M^s \in \mathcal{M}^s$. Let $\tilde{\pi}_{\phi|Y}$ be the posterior of ϕ obtained by updating $\tilde{\pi}_\phi$ with the likelihood $p(Y|\phi)$, which is common among all $M^s \in \mathcal{M}^s$. Suppose that the ϕ -prior in each model is specified according to

$$\pi_{\phi|M^s}(B) = \frac{\tilde{\pi}_\phi(B \cap \Phi_{M^s})}{\tilde{\pi}_\phi(\Phi_{M^s})}, \quad B \in \mathcal{B}(\Phi) \quad (3.2)$$

where $\mathcal{B}(\Phi)$ is the Borel σ -algebra of Φ , i.e., the ϕ -prior is constructed by trimming the support of $\tilde{\pi}_\phi$ to Φ_{M^s} . Then the posterior model probabilities are given by

$$\begin{cases} \pi_{M^p|Y} = \frac{p(Y|M^p)\pi_{M^p}}{\sum_{M^p \in \mathcal{M}^p} p(Y|M^p)\pi_{M^p} + \tilde{p}(Y)\sum_{M^s \in \mathcal{M}^s} O_{M^s}\pi_{M^s}}, & \text{for } M^p \in \mathcal{M}^p, \\ \pi_{M^s|Y} = \frac{\tilde{p}(Y)O_{M^s}\pi_{M^s}}{\sum_{M^p \in \mathcal{M}^p} p(Y|M^p)\pi_{M^p} + \tilde{p}(Y)\sum_{M^s \in \mathcal{M}^s} O_{M^s}\pi_{M^s}}, & \text{for } M^s \in \mathcal{M}^s, \end{cases} \quad (3.3)$$

where O_{M^s} is the posterior-prior plausibility ratio of the set-identifying assumptions of model $M^s \in \mathcal{M}^s$ and $\tilde{p}(Y)$ is the marginal likelihood with respect to $\tilde{\pi}_\phi$,

$$O_{M^s} \equiv \frac{\tilde{\pi}_{\phi|Y}(\Phi_{M^s})}{\tilde{\pi}_\phi(\Phi_{M^s})} = \frac{\tilde{\pi}_{\phi|Y}(IS_\alpha(\phi|M^s) \neq \emptyset)}{\tilde{\pi}_\phi(IS_\alpha(\phi|M^s) \neq \emptyset)}, \quad \tilde{p}(Y) = \int_{\Phi} p(Y|\phi) d\tilde{\pi}_\phi(\phi). \quad (3.4)$$

(ii) Suppose that, in addition to \mathcal{M}^s , all the single-posterior models M^p admit an identical reduced-form. Let $\tilde{\pi}_\phi$ be as defined in (i) of the current lemma and assume $\tilde{\pi}_\phi(\Phi_M) > 0$ holds for all $M \in \mathcal{M}$. If ϕ -prior satisfies (3.2) in every $M \in \mathcal{M}$, then the posterior model probabilities are further simplified to

$$\pi_{M|Y} = \frac{O_M \pi_M}{\sum_{M \in \mathcal{M}} O_M \pi_M} \quad \text{for } M \in \mathcal{M}, \quad (3.5)$$

where $O_M = \frac{\tilde{\pi}_{\phi|Y}(\Phi_M)}{\tilde{\pi}_\phi(\Phi_M)}$.

(iii) If all candidate models are indistinguishable and ϕ -prior is common among them, then the model probabilities are never updated, $\pi_{M|Y} = \pi_M$ for all $M \in \mathcal{M}$ and for any realization of Y .

Lemma 3.1 clarifies the sources of updating of the prior model probabilities. In the first claim, the specification of the ϕ -prior (3.2) simplifies the marginal likelihood of the set-identified model $M^s \in \mathcal{M}^s$ to $\tilde{p}(Y)O_{M^s}$. The computation of $\tilde{p}(Y)$ and O_{M^s} requires one set of Monte

Carlo draws of ϕ each from prior $\tilde{\pi}_\phi$ and from the posterior $\tilde{\pi}_{\phi|Y}$, and assessments of the validity of the identifying assumptions at drawn ϕ 's by checking emptiness of the corresponding identified set. Hence, computation can be saved by avoiding to run separate algorithms for each set-identified model. If all the candidate models admit an identical reduced-form (Lemma 3.1 (ii)), the posterior model probabilities only depend on $\{O_M : M \in \mathcal{M}\}$, so one does not even need to compute the marginal likelihoods. The claim in (iii) says that, if all the candidate models are indistinguishable and share unique ϕ -prior, the prior model probabilities can never be updated. This result is intuitive; assuming the same prior knowledge for ϕ for the indistinguishable model (i.e. the support of ϕ is common), all of them receive the same marginal likelihoods and they cancel out in (3.1).

Scenario 1 in Section 2 satisfies Lemma 3.1 (iii) and thus no update occurs for the model probabilities. Scenario 2 satisfies Lemma 3.1 (i) with $O_{M^s} = 1$, since the identified set in M^s is never empty. In the example of the treatment effect model in Appendix A.2, the point-identified and set-identified models are distinguishable since they have distinct testable implications. Hence, if the common kernel of the prior is maintained as in (3.2), Lemma 3.1 (ii) gives the formula of their posterior model probabilities.

3.3 Post-Averaging Ambiguous Belief and the Range of Posteriors

Estimation of the single-posterior models proceeds in the standard Bayesian way. We therefore take $\pi_{\alpha|M^p,Y}$, the posterior for α in each single-posterior model $M^p \in \mathcal{M}^p$, as given.

We perform posterior inference for model $M^s \in \mathcal{M}^s$ in the robust Bayesian way: we specify a single proper prior $\pi_{\phi_{M^s}|M^s}$ that is supported on Φ_{M^s} , and form the set of priors for θ_{M^s} as

$$\Pi_{\theta_{M^s}|M^s} \equiv \left\{ \pi_{\theta_{M^s}|M^s} : \pi_{\theta_{M^s}|M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(B)) = \pi_{\phi_{M^s}|M^s}(B), \forall B \in \mathcal{B}(\Phi_{M^s}) \right\}, \quad (3.6)$$

where $\mathcal{B}(\Phi_{M^s})$ is the Borel σ -algebra of Φ_{M^s} .⁹ In words, $\Pi_{\theta_{M^s}|M^s}$ collects priors for θ_{M^s} that satisfy the identifying assumptions with probability one (i.e., $\pi_{\theta_{M^s}|M^s}(\Theta_{M^s}) = 1$) and whose ϕ_{M^s} -marginals coincide with the specified ϕ_{M^s} -prior. Applying the Bayes rule to each θ_{M^s} -prior in $\Pi_{\theta_{M^s}|M^s}$ with the likelihood, $\tilde{p}(Y|\theta_{M^s}, M^s)$,¹⁰ and marginalizing the resulting posterior of

⁹By noting that the constraints in (3.6) are rewritten as $\int_B \pi_{\theta_{M^s}|M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi)) d\pi_{\phi_{M^s}|M^s}(\phi_{M^s}) = \pi_{\phi_{M^s}|M^s}(B)$ for all $B \in \mathcal{B}(\Phi_{M^s})$, the prior class (3.6) can be equivalently represented as

$$\Pi_{\theta_{M^s}|M^s} = \left\{ \int_{\Phi_{M^s}} \pi_{\theta_{M^s}|M^s} d\pi_{\phi_{M^s}|M^s} : \pi_{\theta_{M^s}|M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi_{M^s})) = 1, \pi_{\phi_{M^s}|M^s, Y}\text{-a.s.} \right\}.$$

This alternative expression is exploited in the illustrative example of Section 2.

¹⁰The likelihood of θ_M is linked to the likelihood of ϕ_M via $\tilde{p}(Y|\theta_{M^s}, M^s) = p(Y|g(\theta_{M^s}), M^s)$ by the definition of reduced-form parameters.

θ_M via $\alpha = \alpha_M(\theta_M)$, we obtain the following set of posteriors for α :¹¹

$$\begin{aligned} & \Pi_{\alpha|M^s, Y} \\ & \equiv \left\{ \pi_{\alpha|M^s, Y} = \int_{\Phi_M} \pi_{\alpha|M^s, \phi_{M^s}} d\pi_{\phi_{M^s}|M^s, Y} : \pi_{\alpha|M^s, \phi_{M^s}}(IS_{\alpha}(\phi_{M^s}|M^s)) = 1, \pi_{\phi_{M^s}|M^s}\text{-a.s.} \right\}. \end{aligned} \quad (3.7)$$

Given the posterior model probabilities, an posterior for α with the models averaged out is written as

$$\pi_{\alpha|Y} = \sum_{M^p \in \mathcal{M}^p} \pi_{\alpha|M^p, Y} \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\alpha|M^s, Y} \pi_{M^s|Y},$$

where the α -posterior for $M^p \in \mathcal{M}^p$ is unique, while there are multiple α -posteriors for $M^s \in \mathcal{M}^s$ as shown in (3.7). Since there is no restriction that constrains a choice of posterior across the set of posteriors, the set of averaged posteriors can be represented as

$$\Pi_{\alpha|Y} = \left\{ \sum_{M^p \in \mathcal{M}^p} \pi_{\alpha|M^p, Y} \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\alpha|M^s, Y} \pi_{M^s|Y} : \pi_{\alpha|M^s, Y} \in \Pi_{\alpha|M^s, Y} \forall M^s \in \mathcal{M}^s \right\}. \quad (3.8)$$

This is a representation of the post-averaging ambiguous belief that generalizes the two-model case shown in (2.7).

The next proposition provides a formal robust Bayes justification for our averaging formula (3.8) when the structural parameters are common across all models,¹² in which case (3.8) can be obtained by applying the Bayes rule to each prior in a certain well-defined class of priors.

Proposition 3.1 *Suppose that structural parameters are common in all models, $\theta_M = \theta \in \mathbb{R}^{d_{\theta}}$ for all $M \in \mathcal{M}$, and define $\Theta = \cup_{M \in \mathcal{M}} \Theta_M \subset \mathbb{R}^{d_{\theta}}$. Let prior model probabilities ($\pi_M : M \in \mathcal{M}$), prior $\pi_{\theta|M^p}$ for θ in $M^p \in \mathcal{M}^p$, and prior for the reduced-form parameters in $M^s \in \mathcal{M}^s$ be given. Define a set of priors for $(\theta, M) \in \Theta \times \mathcal{M}$:*

$$\Pi_{\theta, M} \equiv \left\{ \pi_{\theta, M} = \pi_{\theta|M} \pi_M : \pi_{\theta|M^s} \in \Pi_{\theta|M^s} \text{ for every } M^s \in \mathcal{M}^s \right\}, \quad (3.9)$$

where $\Pi_{\theta|M^s}$ is defined in (3.6). Then, the Bayes rule applied to each prior in $\Pi_{\theta, M}$ with likelihood $\tilde{p}(Y|\theta, M)$ and marginalization to α yield (3.8) as the class of posteriors for α .

The next proposition derives the range of posterior means, posterior quantiles, and the posterior probabilities when the posterior for α varies within $\Pi_{\alpha|Y}$.

¹¹Lemma A.1 in Appendix A shows a formal derivation of $\Pi_{\alpha|M^s, Y}$.

¹²The reason we assume the common structural parameter space is to ensure that we can construct a prior distribution on the product space of the structural parameter space and the model space.

Proposition 3.2 Let $[l(\phi_{M^s}|M^s), u(\phi_{M^s}|M^s)]$ be the convex hull of the identified set $IS_\alpha(\phi_{M^s}|M^s)$ in model $M^s \in \mathcal{M}^s$.

(i) The range of posterior means of $\Pi_{\alpha|Y}$ is the convex interval with lower and upper bounds:

$$\begin{aligned} \inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha) &= \sum_{M^p \in \mathcal{M}^p} E_{\alpha|M^p, Y}(\alpha) \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} E_{\phi_{M^s}|Y, M^s}[l(\phi_{M^s}|M^s)] \pi_{M^s|Y}, \\ \sup_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha) &= \sum_{M^p \in \mathcal{M}^p} E_{\alpha|M^p, Y}(\alpha) \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} E_{\phi_{M^s}|Y, M^s}[u(\phi_{M^s}|M^s)] \pi_{M^s|Y}, \end{aligned}$$

where $E_{\phi_{M^s}|Y, M^s}(\cdot)$ is the expectation with respect to the posterior of ϕ_{M^s} .

(ii) For any measurable subset H in \mathbb{R} , the lower bound of the posterior probabilities on $\{\alpha \in H\}$ in the class $\Pi_{\alpha|Y}$ (the lower posterior probability of $\Pi_{\alpha|Y}$) is

$$\inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} \pi_{\alpha|Y}(H) = \sum_{M^p \in \mathcal{M}^p} \pi_{\alpha|M^p, Y}(H) \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\phi_{M^s}|Y, M^s}(IS_\alpha(\phi_{M^s}|M^s) \subset H) \cdot \pi_{M^s|Y}.$$

(iii) The lower and upper bounds of the cumulative distribution function (cdf) of $\pi_{\alpha|Y} \in \Pi_{\alpha|Y}$ are

$$\begin{aligned} \underline{\pi}_{\alpha|Y}(a) &\equiv \inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} \pi_{\alpha|Y}([-\infty, a]) \\ &= \sum_{M^p \in \mathcal{M}^p} \pi_{\alpha|M^p, Y}([-\infty, a]) \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\phi_{M^s}|Y, M^s}(\{u(\phi_{M^s}|M^s) \leq a\}) \pi_{M^s|Y}, \\ \bar{\pi}_{\alpha|Y}(a) &\equiv \sup_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} \pi_{\alpha|Y}([-\infty, a]) \\ &= \sum_{M^p \in \mathcal{M}^p} \pi_{\alpha|M^p, Y}([-\infty, a]) \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\phi_{M^s}|Y, M^s}(\{l(\phi_{M^s}|M^s) \leq a\}) \pi_{M^s|Y}, \end{aligned}$$

and the range of posterior τ -th quantiles, $\tau \in (0, 1)$, is $\left[\inf\{a : \bar{\pi}_{\alpha|Y}(a) \geq \tau\}, \inf\{a : \underline{\pi}_{\alpha|Y}(a) \geq \tau\} \right]$.

If a set-identified model delivers $IS_\alpha(\phi_{M^s}|M^s)$ as a connected interval at every reduced-form parameter value, then we can view $[E_{\phi_{M^s}|Y, M^s}[l(\phi_{M^s}|M^s)], E_{\phi_{M^s}|Y, M^s}[u(\phi_{M^s}|M^s)]]$ as an estimator of the identified set in model M^s . We can therefore interpret the range of post-averaging posterior means as the weighted Minkowski sum of the Bayesian point estimators (posterior means) in the point-identified models and the identified set estimators in the set-identified models. The second claim of the proposition provides an analytical expression for the lower probability of $\Pi_{\alpha|Y}$. This lower probability is a mixture of the containment functionals of the random sets, which in turn can be viewed as the containment functional of the *mixture random sets* $\Pr(IS_\alpha^{mix} \subset A)$, where IS_α^{mix} is generated according to

$$\begin{aligned} M &\sim \text{Multinomial}(\{\pi_{M|Y}\}_{M \in \mathcal{M}}), \\ IS_\alpha^{mix} &= \begin{cases} \{\alpha\}, & \alpha|(M^p, Y) \sim \pi_{\alpha|M^p, Y} \text{ for } M^p \in \mathcal{M}^p, \\ IS_\alpha(\phi_{M^s}|M^s), & \phi_{M^s}|(M^s, Y) \sim \pi_{\phi_{M^s}|M^s, Y} \text{ for } M^s \in \mathcal{M}^s. \end{cases} \end{aligned} \tag{3.10}$$

This way of interpreting the lower probability of $\Pi_{\alpha|Y}$ simplifies its computation and justifies the algorithm presented in (2.9).

3.4 Computation

To report the range of posteriors based on the analytical expressions in Proposition 3.2, we need to compute (I) the posterior model probabilities (equivalently, the marginal likelihood in each $M \in \mathcal{M}$), (II) the posterior of α for each single-posterior model, and (III) the identified set $IS_{\alpha}(\phi_{M^s}|M^s)$ and the posterior of ϕ_{M^s} for each multiple-posterior model. Estimation of the single-posterior models in (II) is standard, and we assume some suitable posterior sampling algorithm is applicable to obtain Monte Carlo draws of $\alpha \sim \pi_{\alpha|M^p,Y}$. For (I), efficient and reliable algorithms to compute the marginal likelihood are available in the literature, e.g., see Chib and Jeliazkov (2001), Geweke (1999), and Sims, Waggoner, and Zha (2008). When Lemma 3.1 (i) or (ii) applies, such as in the empirical application of Section 5, the computation of the marginal likelihoods for multiple-posterior models can be reduced to the computation of the posterior-prior plausibility ratios O_M . Since O_M 's and the quantities in (III) are less standard, this section briefly discusses how to compute them under the setting of Lemma 3.1(i) or (ii), i.e., \mathcal{M}^s admits an identical reduced-form.

In each multiple-posterior model, if one can assess the non-emptiness of the identified set at each $\phi \in \Phi$, the posterior-prior plausibility ratio O_{M^s} can be computed simply by plugging in numerical approximations for the prior and posterior probabilities of non-emptiness of the identified set into (3.4). The denominator of O_{M^s} is computed by drawing many ϕ 's from the prior $\tilde{\pi}_{\phi}$ and computing the fraction of draws that yield nonempty identified sets. The numerator of O_{M^s} is computed similarly except that ϕ 's are drawn from the posterior $\tilde{\pi}_{\phi|Y}$.¹³

Monte Carlo draws of the lower and upper bounds of the identified set in model $M \in \mathcal{M}^s$ can be obtained by first drawing ϕ 's from the posterior $\tilde{\pi}_{\phi|Y}$, then retaining the draws of ϕ that yield nonempty $IS_{\alpha}(\phi|M^s)$, and computing the corresponding $l(\phi|M^s)$ and $u(\phi|M^s)$. Their sample averages approximate $E_{\phi|M^s,Y}(l(\phi|M^s))$ and $E_{\phi|M^s,Y}(u(\phi|M^s))$. Implementation of this procedure relies on computability of the lower and upper bounds of the identified set for each ϕ . Whether it is a simple task or not depends on the type of application. In the SVAR application of Section 5, we compute $l(\phi|M^s)$ and $u(\phi|M^s)$ by numerical optimization. Alternatively, adopting the criterion function approach of Chernozhukov, Hong, and Tamer (2007), the computation of the lower and upper bounds of the identified set can be facilitated by applying the slice sampling algorithm proposed by Kline and Tamer (2016).

¹³For instance, in the SVAR application considered in Section 5, we can assess non-emptiness of the identified set by drawing many non-identified parameters (rotation matrices) from the uniform distribution (Haar measure on the space of orthonormal matrices) using the sampling algorithm of Uhlig (2005), and then verifying if any of the draws satisfy the imposed sign restrictions. See also Algorithm 5.1 in Giacomini and Kitagawa (2015).

Utilizing the mixture random set representation shown in (3.10), we can use the following algorithm to approximate the lower posterior probability:

Algorithm 3.1

- Step 1: Draw a model $M \in \mathcal{M}$ from a multinomial distribution with parameters $(\pi_{M|Y} : M \in \mathcal{M})$.*
- Step 2: If the drawn M belongs to \mathcal{M}^p , then draw $\alpha \sim \pi_{\alpha|M,Y}$ and set $IS_{\alpha}^{mix} = \{\alpha\}$ (a singleton).
If the drawn M belongs to \mathcal{M}^s , draw $\phi_M \sim \pi_{\phi|M,Y}$ and set $IS_{\alpha}^{mix} = IS_{\alpha}(\phi_M|M)$.¹⁴*
- Step 3: Repeat Steps 1 and 2 many (G) times and obtain G draws of IS_{α}^{mix} : $IS_{\alpha,1}^{mix}, \dots, IS_{\alpha,G}^{mix}$.*
- Step 4: Let $[l_g^{mix}, u_g^{mix}]$ be the lower and upper bounds of $IS_{\alpha,g}^{mix}$, $g = 1, \dots, G$, where $l_g^{mix} = u_g^{mix}$ if $IS_{\alpha,g}^{mix}$ is a singleton (i.e., g -th draw of M belongs to \mathcal{M}^p). Approximate the mean bounds of the post-average posterior class by*

$$\inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha) = \frac{1}{G} \sum_{g=1}^G l_g^{mix}, \quad \sup_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha) = \frac{1}{G} \sum_{g=1}^G u_g^{mix}. \quad (3.11)$$

Approximate the lower probability of the post-averaging posterior class at $H \subset \mathbb{R}$ by

$$\inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} \pi_{\alpha|Y}(H) \approx \frac{1}{G} \sum_{g=1}^G 1\{IS_{\alpha,g}^{mix} \subset H\}. \quad (3.12)$$

The draws of IS_{α}^{mix} obtained in Steps 1-3 in Algorithm 3.1 are also useful for constructing the robust credible regions. The robust credible region with credibility $\gamma \in (0, 1)$ is defined as the shortest interval to which every posterior in the class assigns probability at least γ ;

$$C_{\gamma} \equiv \arg \min_{C \in \mathcal{C}} \text{length}(C), \quad \text{s.t.} \quad \inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} \pi_{\alpha|Y}(C) \geq \gamma, \quad (3.13)$$

where \mathcal{C} is the class of connected intervals in \mathbb{R} . Since the constraint in (3.13) can be interpreted equivalently as $\Pr(IS_{\alpha}^{mix} \subset C) \geq \gamma$, the computation of C_{γ} can be reduced to finding the shortest interval that contains the γ -proportion of the Monte Carlo draws of IS_{α}^{mix} . A simple computation algorithm for this optimization problem is shown in Proposition 5.1 of Kitagawa (2012) and it can be readily applied to the current context.

3.5 Asymptotic Properties

This section analyzes the asymptotic properties of our method. The procedure is finite-sample exact (up to Monte Carlo approximation errors) and does not rely on asymptotic approximations. The asymptotic analysis is nevertheless valuable, as it highlights what aspects of the

¹⁴Note that since $\pi_{\phi|M,Y}$ is supported only on the set of ϕ 's yielding nonempty identified set, $IS_{\alpha}(\phi|M)$ computed subsequently is nonempty.

prior input, if any, remain influential even in large sample. In this section, we make the sample size explicit in our notation by denoting a size n sample by Y^n .

We assume that \mathcal{M} admits an identical reduced-form (Definition 3.1) and that at least one model is correctly specified, so that the data-generating process is given by $p(Y^n|\phi_{true})$, where $\phi_{true} \in \Phi$ is the true reduced-form parameter value. We denote the unconstrained maximum likelihood estimator for ϕ by $\hat{\phi} \equiv \arg \max_{\phi \in \Phi} p(Y^n|\phi)$ and the true probability law of the sampling sequence $\{Y^n : n = 1, 2, \dots\}$ by $P_{Y^\infty|\phi_{true}}$.

For our asymptotic analysis, we impose the following regularity assumptions:

Assumption 3.2 (i) \mathcal{M} admits an identical reduced-form and every $M \in \mathcal{M}$ satisfies either one of the following conditions:

(A) Φ_M contains ϕ_{true} in its interior.

(B) Φ_M^c contains ϕ_{true} in its interior.

\mathcal{M}_A , denoting the set of models satisfying condition (A), is nonempty.

(ii) Let $l_n(\phi) \equiv n^{-1} \log p(Y^n|\phi)$. There exist an open neighborhood B of ϕ_{true} and $n_0 \geq 1$, such that for any $\{Y^n : n = n_0, n_0 + 1, \dots\}$, $l_n(\cdot)$ is third-time differentiable with the third-order derivatives bounded uniformly on B .

(iii) Let $H_n(\hat{\phi}) \equiv -\frac{\partial^2 l_n(\hat{\phi})}{\partial \phi' \partial \phi}$. $H_n(\hat{\phi})$ is a positive definite matrix and $\liminf_{n \rightarrow \infty} \det(H_n(\hat{\phi})) > 0$, with $P_{Y^\infty|\phi_{true}}$ -probability one.

(iv) For any open neighborhood B of ϕ_{true} ,

$$\limsup_{n \rightarrow \infty} \sup_{\phi \in \Phi \setminus B} \{l_n(\phi) - l_n(\phi_{true})\} < 0$$

holds with $P_{Y^\infty|\phi_{true}}$ -probability one.

(v) For every $M \in \mathcal{M}$, $\pi_{\phi|M}$ has probability density $f_{\phi|M}(\phi) \equiv \frac{d\pi_{\phi|M}}{d\phi}(\phi)$ with respect to the Lebesgue measure on Φ_M and $f_{\phi|M}(\phi)$ is continuously differentiable with a uniformly bounded derivative. For every $M \in \mathcal{M}_A$, $f_{\phi|M}(\phi_{true}) > 0$.

Assumption 3.2 (i) implies that none of the models has ϕ_{true} on the boundary of its reduced-form parameter space. \mathcal{M}_A defined in Assumption 3.2 (i) collects the models that are observationally equivalent at ϕ_{true} in the sense of Definition 3.2 (i). The requirement that ϕ_{true} be in the interior of Φ_M implies that Φ_M , $M \in \mathcal{M}_A$, has a nonempty interior in \mathbb{R}^d . For a set-identified model, condition (A) implies that $M^s \in \mathcal{M}_A$ has a nonempty identified set in an open neighborhood of ϕ_{true} , and condition (B) implies that $M^s \in \mathcal{M}^s \setminus \mathcal{M}_A$ has an empty

identified set in an open neighborhood of ϕ_{true} . Assumptions 3.2 (iii) and (iv) impose regularity conditions that imply almost sure consistency of $\hat{\phi}$. Assumptions 3.2 (ii) and (v), imposing smoothness of the log-likelihood and ϕ -prior, allow an application of the Laplace method to approximate the large sample marginal likelihood. Assumptions similar to Assumptions 3.2 (ii) - (v) appear in Kass, Tierney, and Kadane (1990) in their validation of the higher-order expansion of the marginal likelihood.

The next proposition, which is a large sample analogue of Lemma 3.1, derives the limits of the posterior model probabilities.

Proposition 3.3 (i) *Suppose Assumption 3.2 holds. Then*

$$\pi_{M|Y^\infty} \equiv \lim_{n \rightarrow \infty} \pi_{M|Y^n} = \begin{cases} \frac{f_{\phi|M}(\phi_{true}) \cdot \pi_M}{\sum_{M' \in \mathcal{M}_A} f_{\phi|M'}(\phi_{true}) \cdot \pi_{M'}}, & \text{for } M \in \mathcal{M}_A, \\ 0, & \text{for } M \notin \mathcal{M}_A. \end{cases} \quad (3.14)$$

with $P_{Y^\infty|\phi_{true}}$ -probability one.

(ii) *Suppose that Assumption 3.2 holds and a prior for ϕ given M is constructed according to (3.2) with a proper prior $\tilde{\pi}_\phi$. If $\tilde{\pi}_\phi(\Phi_M) > 0$ for all $M \in \mathcal{M}$,*

$$\pi_{M|Y^\infty} = \begin{cases} \frac{\tilde{\pi}_\phi(\Phi_M)^{-1} \cdot \pi_M}{\sum_{M' \in \mathcal{M}_A} \tilde{\pi}_\phi(\Phi_{M'})^{-1} \cdot \pi_{M'}}, & \text{for } M \in \mathcal{M}_A, \\ 0, & \text{for } M \notin \mathcal{M}_A. \end{cases} \quad (3.15)$$

with $P_{Y^\infty|\phi_{true}}$ -probability one.

(iii) *Under the assumptions of Lemma 3.1 (iii), $\pi_{M|Y^\infty} = \pi_M$ holds for every $M \in \mathcal{M}$ for any sampling sequence $\{Y^n : n = 1, 2, \dots\}$.*

The proposition clarifies the large sample behavior of the posterior model probabilities when the models admit an identical reduced-form. First, it shows that our procedure asymptotically screens out misspecified models $M \notin \mathcal{M}_A$, as their posterior probabilities converge to zero irrespective of the prior probabilities. If there is only one model consistent with the data generating process, asymptotically it has probability one. Second, if \mathcal{M}_A contains multiple models, their asymptotic probabilities are determined by the prior model probabilities and the densities of the ϕ -priors evaluated at ϕ_{true} . This implies that the sensitivity of the post-averaging posterior to the choices of ϕ -priors and prior model probabilities does not vanish asymptotically when multiple models are observationally equivalent at ϕ_{true} . Third, when the ϕ -priors share a common kernel, as assumed in Proposition 3.3 (ii), the asymptotic model probabilities are proportional to the reciprocal of the prior probability (in terms of $\tilde{\pi}_\phi$) that the distribution of data is consistent with the identifying assumptions. Hence, the asymptotic posterior model probabilities are higher for more observationally restrictive models, i.e., if $\Phi_{M_1} \subset \Phi_{M_2}$ for $M_1, M_2 \in \mathcal{M}_A$,

we have $\pi_{M_1|Y^\infty} \geq \pi_{M_2|Y^\infty}$. This result is in line with the principle of parsimony (Ockham's razor), which the standard Bayesian model selection/averaging is typically equipped with — we should prefer a more parsimonious model among those that explain the data equally well. Note that the notion of parsimony here refers to the size of the reduced-form parameter spaces, and has nothing to do with the strength of the identifying assumptions (often measured by the width of the identified set for α).¹⁵

A combination of the asymptotic posterior model probabilities obtained in Proposition 3.3 and the asymptotic behavior of $\pi_{\alpha|M,Y^n}$ for single-posterior models and of $\Pi_{\alpha|M,Y^n}$ for multiple-posterior models yields the asymptotic convergence properties of the range of post-averaging posteriors. To be specific, in addition to Assumption 3.2, we assume that (i) the posterior for ϕ is consistent for ϕ_{true} with $P_{Y^\infty|\phi_{true}}$ -probability one, (ii) for $M^p \in \mathcal{M}^p \cap \mathcal{M}_A$, $\alpha_{M^p}(\cdot)$ is continuous at ϕ_{true} and the posterior of $\alpha_{M^p}(\phi)$ is uniformly integrable with $P_{Y^\infty|\phi_{true}}$ -probability one, and (iii) for $M^s \in \mathcal{M}^s \cap \mathcal{M}_A$, $IS_\alpha(\phi|M^s)$ is a compact and continuous correspondence at ϕ_{true} and the posteriors of $l(\phi|M^s)$ and $u(\phi|M^s)$ are uniformly integrable with $P_{Y^\infty|\phi_{true}}$ -probability one. Then, the range of post-averaging posterior means considered in Proposition 3.2 (i) has the following limits:

$$\begin{aligned} & \lim_{n \rightarrow \infty} \left[\inf_{\pi_{\alpha|Y^n} \in \Pi_{\alpha|Y^n}} E_{\alpha|Y^n}(\alpha), \sup_{\pi_{\alpha|Y^n} \in \Pi_{\alpha|Y^n}} E_{\alpha|Y^n}(\alpha) \right] \\ &= \sum_{M^p \in \mathcal{M}^p \cap \mathcal{M}_A} \alpha_{M^p}(\phi_{true}) \pi_{M^p|Y^\infty} + \left[\sum_{M^s \in \mathcal{M}^s \cap \mathcal{M}_A} l(\phi_{true}|M^s) \pi_{M^s|Y^\infty}, \sum_{M^s \in \mathcal{M}^s \cap \mathcal{M}_A} u(\phi_{true}|M^s) \pi_{M^s|Y^\infty} \right]. \end{aligned}$$

4 Discussion

4.1 Relationship with ϵ -contaminated Class of Priors

The method proposed in this paper has a close link to performing robust Bayes analysis using an ϵ -contaminated class of priors (Huber (1973), Berger and Berliner (1986)). To clarify this, consider the simple case of one single posterior model and one multiple posterior model, $\mathcal{M} = \{M^p, M^s\}$. Further assume that the models share the same parametrization of the structural model and the likelihood for the common structural parameters θ does not depend on the model.

¹⁵For instance, in a SVAR, a model point-identified by a set of equality restrictions is not observationally restrictive, while a model set-identified by sign restrictions is observationally restrictive if the number of sign restrictions is larger than the number of variables in the SVAR system. If the ϕ -priors satisfy (3.2) and the two models are observationally equivalent at ϕ_{true} , then, relative to the prior model weights, the sign-restricted model receives a larger weight than the point-identified model in large sample.

Given (π_{M^p}, π_{M^s}) , $\pi_{\theta|M^p}$, and $\Pi_{\theta|M^s}$ in the form of (3.6), consider the set of priors for θ constructed by marginalizing $\Pi_{\theta,M}$ of Proposition 3.1 to θ ;

$$\Pi_{\theta} \equiv \left\{ \pi_{\theta} = \pi_{\theta|M^p} \pi_{M^p} + \pi_{\theta|M^s} \pi_{M^s} : \pi_{\theta|M^s} \in \Pi_{\theta|M^s} \right\}. \quad (4.1)$$

Similarly to Proposition 3.1, we obtain the post-averaging ambiguous belief $\Pi_{\alpha|Y}$ by updating Π_{θ} prior-by-prior with the common likelihood of θ and marginalizing to α .

A general formulation of an ϵ -contaminated class of priors is given by

$$\Pi_{\theta}^{\epsilon} \equiv \left\{ \pi_{\theta} = (1 - \epsilon) \pi_{\theta}^0 + \epsilon q_{\theta} : q_{\theta} \in \mathcal{Q} \right\}, \quad (4.2)$$

where $0 \leq \epsilon \leq 1$ is a prespecified constant, π_{θ}^0 is a *benchmark* prior for θ , and \mathcal{Q} is a set of priors of θ . Following Berger and Berliner (1986), a motivation for considering the ϵ -contaminated class of priors can be stated as follows. The researcher can express an initial believable prior for θ as π_{θ}^0 , but the elicitation process is subject to error by some amount specified by ϵ . q_{θ} captures in what way π_{θ}^0 differs from the most credible prior and \mathcal{Q} specifies the set of possible departures. Huber (1973) and Berger and Berliner (1986) show the ranges of posterior probabilities for various specifications of \mathcal{Q} when a prior varies over Π_{θ}^{ϵ} .

Despite the fact that the motivation for our averaging procedure differs from the original motivation of the ϵ -contaminated class of priors, the prior input of our averaging procedure specified in (4.1) has the same form as the ϵ -contaminated class of priors (4.2) — Π_{θ} is an ϵ -contaminated class of priors where the benchmark prior is the single-prior (point-identified) model $\pi_{\theta}^0 = \pi_{\theta|M^p}$, the amount of contamination is the prior model probability assigned to the set-identified model $\epsilon = \pi_{M^s}$, and the set of priors \mathcal{Q} corresponds to the multiple priors for the set-identified model $\Pi_{\theta|M^s}$. This clarifies a robust Bayes interpretation of our averaging method.¹⁶ If the single-posterior (point-identified) model plays the role of a sensible benchmark model subject to potential misspecification, averaging it with the set-identified model with weight π_{M^s} can be interpreted as performing sensitivity analysis by contaminating the prior of the point-identified model by an amount π_{M^s} in every possible direction subject to the set-identifying assumptions.

The robust Bayes literature on ϵ -contaminated priors has considered several specifications of \mathcal{Q} that lead to analytically tractable classes of posteriors (Berger and Berliner (1986)). To our knowledge, however, the class of priors in the form of $\Pi_{\theta|M^s}$ has not been investigated. Motivated by partial identification analysis, our analysis offers a new way to specify \mathcal{Q} without losing analytical and numerical tractability.

¹⁶As an alternative to the prior-by-prior updating, Berger and Berliner (1986) also considers the Type-II Maximum Likelihood updating rule (empirical Bayes updating rule) of Good (1965). This alternative approach resolves ambiguity by selecting from the class a prior that maximizes the marginal likelihood. Note that the Type-II Maximum Likelihood procedure fails to select a unique prior from Π_{θ} , because $\pi_{\theta|M^s} \in \Pi_{\theta|M^s}$ sharing a common prior for ϕ has the constant marginal likelihood.

4.2 Relationship with Hierarchical Bayesian Approach

Point-identifying assumptions or a prior for structural parameters sometimes come from a structural econometric model based on economic theory. A set-identified model, in contrast, may represent a “semi-structural” heuristic description of the underlying causal mechanisms with a flexible functional form. For instance, in empirical macroeconomic policy analysis, we can view a DSGE model as a single-posterior model and a sign restricted SVAR model as a set-identified model.

In such contexts, averaging models offers a way to combine the structural modelling approach and a more “reduced-form” approach.¹⁷ The macroeconometrics literature has proposed using hierarchical Bayesian methods to bridge the gap between structural and “reduced-form” approaches (Del Negro and Schorfheide (2004)), in which the structural parameters in the DSGE model act as hyperparameters of a prior for SVAR parameters.

The robust Bayes averaging approach, albeit similar in motivation in such contexts, differs from the hierarchical Bayesian approach in several ways. First, the hierarchical Bayesian approach always leads to a single posterior for the impulse responses, no matter whether they are identified or not in the SVAR model. If they are not, this means that the prior for the structural parameters in the DSGE model and the prior for the SVAR parameters (given the hyperparameters) have some part that is unrevisable by the data. Hence, if one cannot specify these priors with full confidence, posterior sensitivity may well become a concern. In contrast, our procedure classifies the DSGE model as a single posterior model and the set-identified SVAR as a multiple-posterior model. Limited credibility in the prior for the Bayesian DSGE model can be incorporated into the posterior inference by averaging it with the set-identified SVAR model with carefully specified π_{Ms} . Second, in the hierarchical Bayesian approach, tightness of the prior around the mean predicted by the DSGE model plays the role of prior confidence assigned to the structural model. In our procedure, the model probability assigned to the structural model governs the degree of confidence. It is however important to distinguish the notions of confidence between the two approaches, since the former is in the scale of Bayesian probabilistic uncertainty while the latter is in the scale of ambiguity (Knightian uncertainty).

4.3 Eliciting Prior Model Probabilities

The key prior input of our procedure is the prior model probability. A natural starting point is to assume a uniform distribution of prior probabilities, however our procedure can readily accommodate non-uniform probabilities. Discussions on how to determine prior probabilities in Bayesian averaging are in, e.g., George (1999) in the discussion of Clyde (1999), where, in order to prevent from overvaluing similar models, he suggests a “dilution” technique, i.e., if

¹⁷What we mean by “reduced-form” approach here differs from the technical terminology of the reduced-form model/parameters in our expositions.

some models are similar, the weight attached to the original model should be split between that model and its duplicates. Among others, Chipman (1996) attaches smaller prior probabilities to models that are unlikely, Hoeting, Madigan, Raftery, and Volinsky (1999) rely on variable selection in regression models to determine prior probabilities and Clyde and George (2004) propose a Bernoulli specification.

In our context, the robust Bayesian viewpoint based on the ϵ -contaminated class of priors can help clarify the interpretation of the prior model probabilities and facilitate their elicitation.

Suppose again that the set of candidate models consists of one point-identified model M^p and one set-identified model M^s . Assume in addition that M^p is nested in M^s , in the sense that the identifying assumptions in M^p include those in M^s . In this case, the prior model probability assigned to M^p can be interpreted as the *minimal* amount of credibility assigned to the identifying assumptions in model M^p , and the prior model probability assigned to the set-identified model can be interpreted as the *maximal* amount of contamination given to the point-identifying assumptions imposed in M^p but not in M^s . The reason that π_{M^p} is giving the credibility lower bound for model M^p is that, when model M^s nests model M^p , the set of priors specified in model M^s contains beliefs that assign full or partial credibility to the identifying assumptions in M^p . As a result, any prior probability between $[\pi_{M^p}, 1]$ can be attained for the credibility of the identifying assumptions in M^p .

The interpretation of the prior model probabilities differs when the identifying assumptions in models M^p and M^s are non-overlapping. In this case, the prior model probabilities are interpreted as the standard probabilistic belief assigned over mutually exclusive models.

When the identifying assumptions in models M^p and M^s are non-nested but overlapping (e.g., Scenario 1 in Section 2), interpreting the model probabilities may not appear as clear-cut as in the previous two cases. However, the lower credibility bound interpretation of π_{M^p} given in the nested case above remains valid. What differs from the nested case is that the maximal credibility that can be assigned to the identifying assumptions in M^p can be strictly less than one.

5 Empirical Application

We illustrate our method in the context of a conventional monetary SVAR for the federal funds rate i_t , real output growth Δy_t and inflation π_t , as in Aruoba and Schorfheide (2011), Moon, Schorfheide, and Granziera (2013) and Giacomini and Kitagawa (2015). Following Notation 3.1 in Giacomini and Kitagawa (2015), we order the variables so that we can easily verify the

conditions guaranteeing convexity of the identified set using their Lemmas 5.1 and 5.2.

$$A_0 \begin{pmatrix} i_t \\ \Delta y_t \\ \pi_t \end{pmatrix} = c + \sum_{j=1}^4 A_j \begin{pmatrix} i_{t-j} \\ \Delta y_{t-j} \\ \pi_{t-j} \end{pmatrix} + \begin{pmatrix} \epsilon_t^i \\ \epsilon_t^{\Delta y} \\ \epsilon_t^\pi \end{pmatrix} \text{ for } t = 1, \dots, T \quad (5.1)$$

where

$$A_0 = \begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \\ a_{31} & a_{32} & a_{33} \end{pmatrix}. \quad (5.2)$$

Assume $\epsilon_t = [\epsilon_t^i, \epsilon_t^{\Delta y}, \epsilon_t^\pi]'$ are i.i.d. normally distributed with mean zero and variance-covariance the identity matrix I_3 . The corresponding reduced-form VAR is:

$$y_t = b + \sum_{j=1}^4 B_j y_{t-j} + u_t, \quad (5.3)$$

where $b = A_0^{-1}c$, $B_j = A_0^{-1}A_j$, $u_t = A_0^{-1}\epsilon_t$, $\text{var}(u_t) = E(u_t u_t') = \Sigma = A_0^{-1}(A_0^{-1})'$. The reduced form parameter is $\phi = (b, B_1, \dots, B_4, \Sigma)$.

The first equation in (5.1) is interpreted as a monetary policy function: the Federal Reserve reacts to price and GDP, as well as lags of all variables. The second and third equations represent aggregate demand (AD) and aggregate supply (AS), respectively. The data are quarterly observations from 1965:1 to 2005:1 from the FRED2 database.

The prior for the reduced-form parameter is conjugate,¹⁸ relatively loose and belongs to the Normal Inverse-Wishart family:

$$\Sigma \sim \mathcal{IW}(\Psi, d), \quad \beta | \Sigma \sim \mathcal{N}(\bar{b}, \Sigma \otimes \Omega),$$

where $\beta \equiv \text{vec}([b, B_1, \dots, B_4]')$. Ψ is the location matrix of Σ , d is a scalar degrees of freedom hyperparameter and \bar{b} and Ω are the prior mean and variance-covariance matrix of β .

Following Christiano, Eichenbaum, and Evans (1999), we impose the sign normalization restrictions so that the diagonal elements of A_0 are nonnegative. As a result, we can interpret a unit positive change in a structural shock as a one standard-deviation positive shock to the corresponding variable.

5.1 Averaging Indistinguishable Models

Suppose we are interested in the output response to a unit positive shock in the federal funds rate ϵ_t^i at horizon h , $IR_{\Delta y i}^h$, and consider the following two sets of identifying assumptions.

¹⁸In order to reduce the computational burden, we use a conjugate prior as its posterior and marginal likelihood is analytically available.

- *Model 1 (M1, point-identified)*

Consider the standard recursive causal ordering restrictions (Bernanke (1986) and Sims (1980)), assuming that AD and AS do not react on impact to the interest rate shock:

This identification scheme restricts A_0 in (5.1) and (5.2) so that $a_{21} = a_{31} = a_{23} = 0$.

- *Model 2 (M2, set-identified through zero restrictions)*

The identification scheme in Model 1 is controversial. For example, assumption $a_{31} = 0$, implying that prices do not react contemporaneously to the interest rate shock, can be difficult to justify if the researcher relies on the stock price index rather than the GDP deflator.¹⁹ Thus, in Model 2 we leave AS unrestricted, i.e., AS can react to the interest rate within a quarter. By Lemma 5.1 in Giacomini and Kitagawa (2015), Model 2 delivers a convex identified set for $IR_{\Delta yi}^h$ for every value of the reduced form parameters.

Figure 1 focuses on the output response at horizon $h = 3$ implied by Model 1, Model 2 and their averages for different sets of prior probabilities. In the top panel, the vertical solid lines for Model 1 are the 90% credible region for the point-identified output response based on a single posterior for the impulse response; the vertical dashed lines for Model 2 are the posterior mean bounds (consistent estimator of the identified set) for the output response and the solid line represents credible regions for various levels of credibility. The bottom panels report the model average results when the prior weight assigned to model 1 is $w_1 = .5$ or $w_1 = .8$. The vertical dashed lines for the averaged model can be viewed as shrinking the identified set estimator from Model 2 towards the point estimator from Model 1. Figure 2 reports the output response credible sets for multiple horizons for the same models as in Figure 1.

Note that, as is common for standard recursive causal ordering restrictions in small-scale SVARs, the point-identified Model 1 shows a negative response of output in the short run, whereas the set-identified Model 2 supports both positive and negative effects. Averaging the models still does not support a negative output response, as the 90% robust credibility region always crosses the zero line. Note that in this case the models are indistinguishable and so the prior probabilities are not updated.

5.2 Averaging Distinguishable Models

Here we consider two distinguishable models that are widely used in empirical applications: a sign-restricted SVAR and a structural DSGE model.

- *Model 3 (M3, set-identified through sign restrictions)*

¹⁹See Kilian (2013) for details over the limitations of point-identifying assumptions.

We consider the following sign restrictions: the inflation response to a contractionary monetary policy shock is nonpositive for two quarters; the interest rate response is non-negative for two quarters. As in Uhlig (2005), the output response is unrestricted. By Lemma 5.2 in Giacomini and Kitagawa (2015), the identified set in Model 3 is convex.

Consider averaging Model 1 and Model 3 with equal prior probabilities. In contrast to the previous example, the prior probabilities can now be updated using equation (3.5) because the candidate models are distinguishable due to the imposition of observationally restrictive sign restrictions. Figures 3 and 4 report the results of averaging the two models: as in the case of Model 2, Model 3 does not support a negative output response (this is also the conclusion of Uhlig (2005), however based on a single-prior approach). Table 1 shows that the posterior model probabilities strongly favour Model 3 (with posterior probability 0.68), and the average of the two models does not support a negative output response.

- *Model 4 (M_4 , DSGE)*

We consider the Bayesian DSGE model in An and Schorfheide (2007), which is a simplified version of Smets and Wouters (2003) and Christiano, Eichenbaum, and Evans (2005). In order to estimate the model, we rely on the prior specification in An and Schorfheide (2007), Table 2 and use output, inflation and interest rate as observables. We use the Laplace approximation to compute the marginal likelihood.

Figures 5 and 6 show the results of averaging Models 3 and 4. Note that these models do not share a common reduced form, so the (equal) prior probabilities are updated according to equation (3.3). We see that Model 4 implies a negative output response, however its posterior model probability is only 0.27, and the averaged model does not support a negative response.

Finally, Figures 7 and 8 report the results of averaging all models (with equal prior weights). The posterior model probabilities (Table 1, sixth column) show strong evidence for the sign-restricted SVAR, while the support for the DSGE model is again weak. As in all previous cases, the averaged model does not support a negative output response.

5.3 Reverse-Engineering Prior Model Probabilities

Our method lends itself to a useful exercise of reverse engineering that helps shed light on the role of identifying assumptions in drawing inferences. Specifically, we compute the prior probability one needs to attach to a set of assumptions in order for the averaging to preserve certain model's conclusions. In our application, for example, we saw that the negative response of output to a contractionary monetary policy shock disappears once relaxing either standard causal ordering restrictions or the restrictions embedded in a DSGE model. We can thus

compute the prior weight one would assign to a given set of restrictions in order to preserve the negative output response.

First consider Model 1 (point-identification through causal ordering restrictions) and Model 2 (set-identification by relaxing one restriction from Model 1). Letting w be the prior probability of Model 1, the post-averaging interval of posterior means is

$$\left[\inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha), \sup_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha) \right] = \\ = \pi_{M^1|Y} E_{\alpha|M^1,Y}(\alpha) + \pi_{M^2|Y} [E_{\phi|Y,M^2}(l(\phi_{M^2}|M^2)), E_{\phi|Y,M^2}(u(\phi_{M^2}|M^2))]$$

and the posterior model probabilities are equal to the prior probabilities (since the models are indistinguishable).

We want to compute the prior model probability w such that the 90% robustified credible region of the weighted model lies below the zero line at $h = 4$:

$$C_\gamma \equiv \arg \min_{C \in \mathcal{C}} \text{length}(C), \text{ s.t. } \inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} \pi_{\alpha|Y}(C) \geq \gamma,$$

where $\gamma = 0.90$. This is equivalent to solving a system of inequalities where w is the unknown.

We find that one would need to attach a prior probability greater than 0.89 to the validity of the assumption that prices do not react contemporaneously to an interest rate shock in order to preserve a negative output response to a contractionary monetary policy shock.

We next consider Model 1 and Model 3 (set-identification through sign restrictions). The reverse-engineering exercise proceeds as before, with the only difference that now the posterior model probabilities are updated and are equal to

$$\pi_{M^1|Y} = \frac{O_1 \cdot w}{O_1 \cdot w + O_3 \cdot (1 - w)} \text{ and } \pi_{M^3|Y} = \frac{O_3 \cdot (1 - w)}{O_1 \cdot w + O_3 \cdot (1 - w)}.$$

We find that the 90% robustified credible region for the averaged model lies below the zero line only if $w > 0.98$. As expected, one would need to attach very high prior probability to the causal ordering restrictions to obtain a negative output response.

Similar reverse engineering exercises could usefully shed light on the role of identifying assumptions in generating so-called price and liquidity puzzles in monetary SVARs.²⁰

²⁰The price puzzle occurs when contractionary monetary policy shocks produce a positive response of the price level (Sims, 1992). The liquidity puzzle refers to shocks in monetary aggregates leading to an initial rising rather than falling of interest rates (Reichenstein, 1987).

6 Conclusion

We proposed a method to average point-identified models and set-identified models from the multiple prior (ambiguous belief) viewpoint. The method combines single priors in point-identified models with multiple priors in set-identified models, and delivers a set of posteriors. The post-averaging set of posteriors can be summarized by the range of posterior means and robust credible regions, which are easy to compute MCMC methods. Our averaging method can effectively reduce the amount of ambiguity (the size of the posterior class) relative to the analysis based on a set-identified model only, and hence offers a simple and flexible way to introduce additional identifying information into the set-identified model. In the opposite direction, when the set-identified model nests the point-identified model, our method can also offer a simple and flexible way to conduct sensitivity analysis for the point-identified model.

A Appendix

A.1 Omitted Proofs

Derivation of identified set (2.2). Following Uhlig (2005), we reparameterize A via the Cholesky matrix Σ_{tr} and a rotation matrix $Q = \begin{pmatrix} \cos \rho & -\sin \rho \\ \sin \rho & \cos \rho \end{pmatrix}$ with spherical coordinate $\rho \in [0, 2\pi]$. We can then express α as a function of ϕ and the non-identified parameter ρ indexing a rotation matrix;

$$A^{-1} = \Sigma_{tr} Q = \begin{pmatrix} \sigma_{11} \cos \rho & -\sigma_{11} \sin \rho \\ \sigma_{21} \cos \rho + \sigma_{22} \sin \rho & -\sigma_{21} \sin \rho + \sigma_{22} \cos \rho \end{pmatrix}$$

and the parameter of interest is $\alpha = \alpha(\rho, \phi) \equiv \sigma_{11} \cos \rho$. We impose the sign normalization restrictions throughout by constraining the diagonal elements of A to being nonnegative,

$$\sigma_{22} \cos \rho - \sigma_{21} \sin \rho \geq 0 \text{ and } \sigma_{11} \cos \rho \geq 0. \quad (\text{A.1})$$

The sign restrictions $a_{12} \geq 0$ and $a_{21} \leq 0$ are expressed as

$$\sigma_{11} \sin \rho \geq 0 \quad (\text{A.2})$$

$$-\sigma_{22} \sin \rho - \sigma_{21} \cos \rho \leq 0. \quad (\text{A.3})$$

Given ϕ , the identified set for $\alpha = \sigma_{11} \cos \rho$ is given by its range as ρ varies over the range characterized by the restrictions (A.1) - (A.3). Since the second constraint in (A.1) and (A.2) imply $\rho \in [0, \pi/2]$, we focus on how the other two restrictions (the first constraint in (A.1) and (A.3)) tighten up $\rho \in [0, \pi/2]$.

Assume $\sigma_{21} > 0$. Then, they imply $\arctan(-\sigma_{21}/\sigma_{22}) \leq \theta \leq \arctan(\sigma_{22}/\sigma_{21})$. Intersecting this interval with $\rho \in [0, \pi/2]$ leads to $[0, \arctan(\sigma_{22}/\sigma_{21})]$ as the identified set for ρ . Hence,

the identified set for α in the $\sigma_{21} > 0$ case follows. A similar argument leads to the α identified set for the $\sigma_{21} \leq 0$ case. ■

Proof of Lemma 3.1. (i) By the construction of ϕ -prior (3.2), the marginal likelihood for $M \in \mathcal{M}^s$ can be rewritten as

$$\begin{aligned} p(Y|M) &= \int_{\Phi} p(Y|\phi, M) d\pi_{\phi|M}(\phi) \\ &= \int_{\Phi} p(Y|\phi) \cdot \frac{1\{IS_{\alpha}(\phi|M) \neq \emptyset\}}{\tilde{\pi}_{\phi}(IS_{\alpha}(\phi|M) \neq \emptyset)} d\tilde{\pi}_{\phi}(\phi) \\ &= \tilde{p}(Y) \int_{\phi} \frac{1\{IS_{\alpha}(\phi|M) \neq \emptyset\}}{\tilde{\pi}_{\phi}(IS_{\alpha}(\phi|M) \neq \emptyset)} d\tilde{\pi}_{\phi|Y}(\phi) \\ &= \tilde{p}(Y) \frac{\tilde{\pi}_{\phi|Y}(IS_r(\phi|M) \neq \emptyset)}{\tilde{\pi}_{\phi}(IS_r(\phi|M) \neq \emptyset)} = \tilde{p}(Y) O_M, \end{aligned}$$

where the second line uses the assumption that the set-identified models admit an identical reduced-form and the third line follows from the Bayes theorem for the reduced-form parameters, $p(Y|\phi)\tilde{\pi}_{\phi}(\phi) = \tilde{p}(Y)\tilde{\pi}_{\phi|Y}(\phi)$. Plugging this expression of the marginal likelihood into (3.1) leads to the claim.

(ii) Under the additionally imposed assumptions, the marginal likelihood of model $M^p \in \mathcal{M}^p$ agrees with $\tilde{p}(Y)O_{M^p}$. Hence, (3.5) follows.

(iii) The claim follows immediately by setting O_M , $M \in \mathcal{M}$, to one in (3.5). ■

Derivation of $\Pi_{\alpha|M^s, Y}$ in equation (3.7). We derive $\Pi_{\alpha|M^s, Y}$ in the next lemma:

Lemma A.1 *For a set-identified model M^s with the structural parameters $\theta_{M^s} \in \Theta_{M^s}$ and reduced-form parameters $\phi_{M^s} = g_{M^s}(\theta_{M^s}) \in \Phi_{M^s} = g_{M^s}(\Theta_{M^s})$, let a prior for ϕ_{M^s} , $\pi_{\phi_{M^s}|M^s}$ be given. Define the class of priors of θ_{M^s} by*

$$\Pi_{\theta_{M^s}|M^s} \equiv \left\{ \pi_{\theta_{M^s}|M^s} : \pi_{\theta_{M^s}|M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(B)) = \pi_{\phi_{M^s}|M^s}(B), \forall B \in \mathcal{B}(\Phi_{M^s}) \right\}.$$

Updating $\Pi_{\theta_{M^s}|M^s}$ prior-by-prior with the likelihood $\tilde{p}(Y|\theta_{M^s}, M^s)$ and marginalizing the resulting posteriors via $\alpha = \alpha_{M^s}(\theta_{M^s})$ leads to the following set of posteriors for α :

$$\begin{aligned} &\Pi_{\alpha|M^s, Y} \\ &\equiv \left\{ \pi_{\alpha|M^s, Y} = \int_{\Phi_{M^s}} \pi_{\alpha|M^s, \phi_{M^s}} d\pi_{\phi_{M^s}|M^s, Y} : \pi_{\alpha|M^s, \phi_{M^s}}(IS_{\alpha}(\phi_{M^s}|M^s)) = 1, \pi_{\phi_{M^s}|M^s} \text{-a.s.} \right\}. \end{aligned} \tag{A.4}$$

Proof of Lemma A.1. The prior-by-prior updating rule updates $\Pi_{\theta_{M^s}|M^s}$ to

$$\Pi_{\theta_{M^s}|M^s, Y} \equiv \left\{ \pi_{\theta_{M^s}|M^s, Y} : \pi_{\theta_{M^s}|M^s, Y}(\Theta_{M^s} \cap g_{M^s}^{-1}(B)) = \pi_{\phi_{M^s}|M^s, Y}(B), \forall B \in \mathcal{B}(\Phi_{M^s}) \right\}.$$

Since $\pi_{\theta_{M^s}|M^s,Y}(\Theta_{M^s} \cap g_{M^s}^{-1}(B))$ can be written as

$$\pi_{\theta_{M^s}|M^s,Y}(\Theta_{M^s} \cap g_{M^s}^{-1}(B)) = \int_B \pi_{\theta_{M^s}|\phi_{M^s},M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi_{M^s})) d\pi_{\phi_{M^s}|M^s,Y}(\phi_{M^s}),$$

the ϕ_{M^s} -marginal constraints for $\pi_{\theta_{M^s}|M^s,Y}$ are equivalent to

$$\int_B \pi_{\theta_{M^s}|\phi_{M^s},M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi_{M^s})) d\pi_{\phi_{M^s}|M^s,Y}(\phi_{M^s}) = \pi_{\phi_{M^s}|M^s,Y}(B).$$

This equality holds for all $B \in \mathcal{B}(\Phi_{M^s})$ if and only if $\pi_{\theta_{M^s}|\phi_{M^s},M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi_{M^s})) = 1$, $\pi_{\phi_{M^s}|M^s,Y}$ -a.s. Accordingly, an equivalent representation of the class of posteriors for θ_{M^s} is

$$\Pi_{\theta_{M^s}|M^s,Y} = \left\{ \int_{\Phi_{M^s}} \pi_{\theta_{M^s}|\phi_{M^s},M^s} d\pi_{\Phi_{M^s}|Y} : \pi_{\theta_{M^s}|\phi_{M^s},M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi_{M^s})) = 1, \pi_{\phi_{M^s}|M^s,Y}\text{-a.s.} \right\}. \quad (\text{A.5})$$

Note that we have

$$\begin{aligned} \pi_{\alpha|\phi_{M^s},M^s}(IS_{\alpha}(\phi_{M^s}|M^s)) &= \pi_{\theta_{M^s}|\phi_{M^s},M^s}(\alpha_{M^s}^{-1}(IS_{\alpha}(\phi_{M^s}|M^s))) \\ &= \pi_{\theta_{M^s}|\phi_{M^s},M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi_{M^s})), \end{aligned}$$

where the second equality follows by the definition of the identified set of α . Hence, $\pi_{\theta_{M^s}|\phi_{M^s},M^s}(\Theta_{M^s} \cap g_{M^s}^{-1}(\phi_{M^s})) = 1$, $\pi_{\phi_{M^s}|M^s,Y}$ -a.s. holds if and only if $\pi_{\alpha|\phi_{M^s},M^s}(IS_{\alpha}(\phi_{M^s}|M^s)) = 1$, $\pi_{\phi_{M^s}|M^s,Y}$ -a.s. The class of marginalized posteriors for α (A.4) therefore follows. ■

Proof of Proposition 3.1. Let $\pi_{\theta,M} \in \Pi_{\theta,M}$ be a prior. The corresponding posterior for θ with M integrated out can be computed as follows: for any measurable subset $H \subset \Theta$,

$$\begin{aligned} \pi_{\theta|Y}(H) &= \frac{\sum_{M \in \mathcal{M}} \int_H \tilde{p}(Y|\theta, M) d\pi_{\theta|M}(\theta) \pi_M}{\sum_{M \in \mathcal{M}} \left[\int_{\Theta_M} \tilde{p}(Y|\theta, M) d\pi_{\theta|M}(\theta) \right] \pi_M} \\ &= \frac{\left(\sum_{M^p \in \mathcal{M}^p} \pi_{\theta|M^p,Y}(H) p(Y|M^p) \pi_{M^p} \right. \\ &\quad \left. + \sum_{M^s \in \mathcal{M}^s} \left[\int_{\Phi_{M^s}} \pi_{\theta|\phi_{M^s},M^s}(H) p(Y|\phi_{M^s}, M^s) d\pi_{\phi_{M^s}|M^s}(\phi_{M^s}) \right] \pi_{M^s} \right)}{\sum_{M^p \in \mathcal{M}^p} p(Y|M^p) \pi_{M^p} + \sum_{M^s \in \mathcal{M}^s} \left[\int_{\Phi_{M^s}} p(Y|\phi_{M^s}, M^s) d\pi_{\phi_{M^s}|M^s}(\phi_{M^s}) \right] \pi_{M^s}} \\ &= \sum_{M^p \in \mathcal{M}^p} \pi_{\theta|M^p}(H) \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \left[\int_{\Phi_{M^s}} \pi_{\theta|\phi_{M^s},M^s}(H) d\pi_{\phi_{M^s}|M^s,Y}(\phi_{M^s}) \right] \pi_{M^s|Y} \end{aligned}$$

where the second line uses

$$\begin{aligned} \int_H \tilde{p}(Y|\theta, M) d\pi_{\theta|M}(\theta) &= \int_{\Phi_M} \left[\int_{\Theta} 1\{\theta \in H\} \tilde{p}(Y|\theta, M) d\pi_{\theta|\phi_M,M}(\theta) \right] d\pi_{\phi_M|M}(\phi_M) \\ &= \int_{\Phi_M} \left[\int_{\Theta} 1\{\theta \in H\} d\pi_{\theta|\phi_M,M}(\theta) \right] p(Y|\phi_M, M) d\pi_{\phi_M|M}(\phi_M) \\ &= \int_{\Phi_M} \pi_{\theta|\phi_M,M}(H) p(Y|\phi_M, M) d\pi_{\phi_M|M}(\phi_M). \end{aligned}$$

The class of posteriors for θ can be therefore represented as

$$\Pi_{\theta|Y} \equiv \left\{ \sum_{M^p \in \mathcal{M}^p} \pi_{\theta|M^p,Y} \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\theta|M^s,Y} \pi_{M^s|Y} : \pi_{\theta|M^s,Y} \in \Pi_{\theta|M^s,Y}, \forall M^s \in \mathcal{M}^s \right\},$$

where $\Pi_{\theta|M^s,Y}$ is as defined in (A.5). As shown in the proof of Lemma A.1 above, marginalizing $\Pi_{\theta|M^s,Y}$ to α leads to $\Pi_{\alpha|M^s,Y}$ defined in (3.7). We therefore conclude that marginalizing $\Pi_{\theta|Y}$ to α results in $\Pi_{\alpha|Y}$ shown in (3.8). ■

Proof of Proposition 3.2. (i) Since there is no constraint across the posteriors belonging to different posterior classes, it holds

$$\inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} E_{\alpha|Y}(\alpha) = \sum_{M^p \in \mathcal{M}^p} E_{\alpha|M^p,Y}(\alpha) \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\alpha|M^s,Y} \inf_{\pi_{\alpha|M^s,Y} \in \Pi_{\alpha|M^s,Y}} \{E_{\alpha|M^s,Y}(\alpha)\} \cdot \pi_{M^s|Y}.$$

By the construction of $\Pi_{\alpha|M^s,Y}$, an application of Proposition 4.1 (ii) of Giacomini and Kitagawa (2015) shows $\inf_{\pi_{\alpha|M^s,Y} \in \Pi_{\alpha|M^s,Y}} \{E_{\alpha|M^s,Y}(\alpha)\} = E_{\phi|M^s,Y}(l(\phi|M^s))$. The claim of the mean lower bound therefore follows. The mean upper bound can be shown similarly.

(ii) Note that

$$\inf_{\pi_{\alpha|Y} \in \Pi_{\alpha|Y}} \pi_{\alpha|Y}(H) = \sum_{M^p \in \mathcal{M}^p} \pi_{\alpha|M^p,Y}(H) \cdot \pi_{M^p|Y} + \sum_{M^s \in \mathcal{M}^s} \pi_{\alpha|M^s,Y} \inf_{\pi_{\alpha|M^s,Y} \in \Pi_{\alpha|M^s,Y}} \{\pi_{\alpha|M^s,Y}(H)\} \cdot \pi_{M^s|Y}.$$

Proposition 3.1 of Kitagawa (2012) shows

$$\inf_{\pi_{\alpha|M^s,Y} \in \Pi_{\alpha|M^s,Y}} \{\pi_{\alpha|M^s,Y}(H)\} = \pi_{\phi_{M^s}|M^s,Y}(IS_{\alpha}(\phi_{M^s}|M^s) \subset H).$$

(iii) By setting H to $[-\infty, a]$, the lower probability obtained in part (ii) yields the lower bound of the cdfs, since the event $IS_{\alpha}(\phi_{M^s}|M^s) \subset [-\infty, a]$ is equivalent to $u(\phi_{M^s}|M^s) \leq a$. The upper bound follows by noting

$$\begin{aligned} \sup_{\pi_{\alpha|M^s,Y} \in \Pi_{\alpha|M^s,Y}} \pi_{\alpha|M^s,Y}([\infty, a]) &= \pi_{\phi_{M^s}|M^s,Y}(IS_{\alpha}(\phi_{M^s}|M^s) \cap [\infty, a] \neq \emptyset) \\ &= \pi_{\phi_{M^s}|M^s,Y}(l(\phi_{M^s}|M^s) \leq a). \end{aligned}$$

The range of quantiles then follows by inverting these cdf bounds. ■

Next, we show two lemmas used to prove Proposition 3.3. We denote the set of candidate models satisfying condition (A) of Assumption 3.2 (i) by \mathcal{M}_A and the set of those satisfying condition (B) by \mathcal{M}_B . Under Assumption 3.2 (i), $\mathcal{M} = \mathcal{M}_A \cup \mathcal{M}_B$ holds. Note that through these lemmas and the proof of Proposition 3.3, \mathcal{M} is assumed to admit an identical reduced-form with reduced-form parameter dimension $d \geq 1$.

Lemma A.2 Suppose Assumption 3.2 holds. For $M \in \mathcal{M}_A$,

$$\frac{n^{d/2} \det(H_n(\hat{\phi}))^{1/2} p(Y^n|M)}{(2\pi)^{d/2} p(Y^n|\hat{\phi})} - f_{\phi|M}(\hat{\phi}) = O(n^{-1/2}),$$

with $P_{Y^\infty|\phi_{true}}$ -probability one.

Proof of Lemma A.2. Denote the reduced-form parameter vector by $\phi = (\phi_1, \dots, \phi_d)$ and the third-derivative of $l_n(\cdot)$ by $h_{ijk}(\cdot) \equiv \frac{\partial^3}{\partial \phi_i \partial \phi_j \partial \phi_k} l_n(\cdot)$, $1 \leq i, j, k \leq d$. By Assumptions 3.2 (i), (ii) and (iv), there exists B^* an open neighborhood of ϕ_{true} such that $B^* \subset \Phi_M$ holds for all $M \in \mathcal{M}_A$, and

$$\sup_{\phi \in B^*} \max_{1 \leq i, j, k \leq d} |h_{ijk}(\phi)| < \infty, \quad (\text{A.6})$$

and

$$\limsup_{n \rightarrow \infty} \sup_{\phi \in \Phi \setminus B^*} \{l_n(\phi) - l_n(\phi_{true})\} < 0, \quad \text{with } P_{Y^\infty|\phi_{true}}\text{-probability one} \quad (\text{A.7})$$

hold. Since Assumptions 3.2 (iii) and (iv) imply the strong convergence of $\hat{\phi}$, for all sufficiently large n , $\hat{\phi} \in B^*$ holds. Given $\hat{\phi} \in B^*$, consider the third-order mean value expansions of $nl_n(\phi)$:

$$\begin{aligned} nl_n(\phi) &= nl_n(\hat{\phi}) - \frac{n}{2}(\phi - \hat{\phi})' H_n(\hat{\phi})(\phi - \hat{\phi}) + \frac{n}{6} \sum_{1 \leq i, j, k \leq d} h_{ijk}(\tilde{\phi})(\phi_i - \hat{\phi}_i)(\phi_j - \hat{\phi}_j)(\phi_k - \hat{\phi}_k) \\ &= nl_n(\hat{\phi}) - \frac{1}{2} u' H_n(\hat{\phi}) u + \frac{1}{\sqrt{n}} R_{1n}(u), \end{aligned}$$

where $\tilde{\phi}$ is a convex combination of ϕ and $\hat{\phi}$, $u \equiv \sqrt{n}(\phi - \hat{\phi})$, and $R_{1n}(u) = \frac{1}{6} \sum_{1 \leq i, j, k \leq d} h_{ijk}(\tilde{\phi}) u_i u_j u_k$, where u_i is the i -th entry of vector u . By the boundedness of h_{ijk} on B^* , $R_{1n}(u)$ can be bounded by a third-order polynomial of u with bounded coefficients on $\sqrt{n}(B^* - \hat{\phi})$, where $\sqrt{n}(B^* - \hat{\phi})$ is the subset in \mathbb{R}^d that translates B^* by $\hat{\phi}$ and scales up by \sqrt{n} . Plugging this expansion in $p(Y^n|\phi) = \exp(nl_n(\phi))$ and combining it with the first-order expansion of $f_{\phi|M}(\phi)$, we obtain on $\phi \in B^*$ (or equivalently on $u \in \sqrt{n}(B^* - \hat{\phi})$)

$$\begin{aligned} p(Y^n|\phi) f_{\phi|M}(\phi) &= \exp \left\{ nl_n(\hat{\phi}) - \frac{1}{2} u' H_n(\hat{\phi}) u \right\} \left\{ 1 + \frac{1}{\sqrt{n}} R_{1n}(u) + \frac{1}{2n} R_{1n}(u)^2 + \dots \right\} \\ &\quad \times \left\{ f_{\phi|M}(\hat{\phi}) + \frac{1}{\sqrt{n}} R_{2n}(u) \right\} \\ &= \exp \left\{ nl_n(\hat{\phi}) - \frac{1}{2} u' H_n(\hat{\phi}) u \right\} \left\{ f_{\phi|M}(\hat{\phi}) + \frac{1}{\sqrt{n}} R_{3n}(u) \right\}, \end{aligned} \quad (\text{A.8})$$

where the first equality invokes the expansion of $\exp(x) = 1 + x + \frac{1}{2}x^2 + \dots$, $R_{2n} = f'_{\phi|M}(\tilde{\phi})u$, and R_{3n} collects the residual terms that can be bounded uniformly on $\sqrt{n}(B^* - \hat{\phi})$ by a finite order polynomial of u with bounded coefficients.

Integration of $p(Y^n|\phi)f_{\phi|M}(\phi)$ over $\phi \in B^*$ is equivalent to integrating (A.8) in u over $\sqrt{n}(B^* - \hat{\phi})$:

$$\begin{aligned}
& \int_{B^*} p(Y^n|\phi)f_{\phi|M}(\phi)d\phi \\
&= n^{-d/2} \exp\{nl_n(\hat{\phi})\} \left(\int_{\sqrt{n}(B^* - \phi_{true})} \left(f_{\phi|M}(\hat{\phi}) + R_{3n}(u) \right) \exp \left\{ -\frac{1}{2} u' H_n(\hat{\phi}) u \right\} du \right) \\
&= (2\pi)^{d/2} p(Y^n|\hat{\phi}) n^{-d/2} \det(H_n(\hat{\phi}))^{1/2} \left(f_{\phi|M}(\hat{\phi}) E_{H_n}[1_{\sqrt{n}(B^* - \hat{\phi})}(u)] + n^{-1/2} E_{H_n}[R_{3n}(u) \cdot 1_{\sqrt{n}(B^* - \hat{\phi})}(u)] \right) \\
&= (2\pi)^{d/2} p(Y^n|\hat{\phi}) n^{-d/2} \det(H_n(\hat{\phi}))^{1/2} \left(f_{\phi|M}(\hat{\phi}) + O(n^{-1/2}) \right), \tag{A.9}
\end{aligned}$$

where $E_{H_n}(\cdot)$ is the expectation taken with respect to $u \sim \mathcal{N}(0, H_n(\hat{\phi})^{-1})$. Note that the third equality follows since the replacement of $\sqrt{n}(B^* - \hat{\phi})$ with \mathbb{R}^d incurs an error of exponentially decreasing order and $E_{H_n}(R_{3n}(u))$ is finite, i.e., the multivariate normal distribution has finite moments at any order.

Consider now integrating $p(Y^n|\phi)f_{\phi|M}(\phi)$ over $\Phi_M \setminus B^*$.

$$\begin{aligned}
& \int_{\Phi_M \setminus B^*} p(Y^n|\phi)f_{\phi|M}(\phi)d\phi \\
&= (2\pi)^{d/2} p(Y^n|\hat{\phi}) n^{-d/2} \det(H_n(\hat{\phi}))^{1/2} \\
& \quad \times \left((2\pi)^{-d/2} n^{d/2} \det(H_n(\hat{\phi}))^{-1/2} \int_{\Phi_M \setminus B^*} \exp\{n(l_n(\phi) - l_n(\hat{\phi}))\} f_{\phi|M}(\phi) d\phi \right) \\
&\leq (2\pi)^{d/2} p(Y^n|\hat{\phi}) n^{-d/2} \det(H_n(\hat{\phi}))^{1/2} \\
& \quad \times \left((2\pi)^{-d/2} n^{d/2} \det(H_n(\hat{\phi}))^{-1/2} \bar{f}_{\phi|M} \sup_{\phi \in \Phi \setminus B^*} \{ \exp\{n(l_n(\phi) - l_n(\phi_{true}))\} \} \right), \tag{A.10}
\end{aligned}$$

where by Assumption 3.2 (v), $\bar{f}_{\phi|M} \equiv \sup_{\phi \in \Phi} f_{\phi|M}(\phi) < \infty$. Assumptions 3.2 (iii) and (iv) imply that the term in the parentheses of (A.10) converges to zero faster than $n^{-1/2}$ -rate with $P_{Y^\infty|\phi_{true}}$ -probability one. Summing up (A.9) and (A.10) gives the following asymptotic approximation of the marginal likelihood in model $M \in \mathcal{M}_A$.

$$\begin{aligned}
p(Y^n|M) &= \int_{B^*} p(Y^n|\phi)f_{\phi|M}(\phi)d\phi + \int_{\Phi_M \setminus B^*} p(Y^n|\phi)f_{\phi|M}(\phi)d\phi \\
&= (2\pi)^{d/2} p(Y^n|\hat{\phi}) n^{-d/2} \det(H_n(\hat{\phi}))^{1/2} \left(f_{\phi|M}(\hat{\phi}) + O(n^{-1/2}) \right), \tag{A.11}
\end{aligned}$$

with $P_{Y^\infty|\phi_{true}}$ -probability one. Bringing the multiplicative terms in the right-hand side of (A.11) to the left-hand side completes the proof. ■

Lemma A.3 *Suppose Assumption 3.2 holds. For model $M \in \mathcal{M}_B$,*

$$\frac{n^{d/2} \det(H_n(\hat{\phi}))^{1/2} p(Y^n|M)}{(2\pi)^{d/2} p(Y^n|\hat{\phi})} = o(n^{-1/2}),$$

with $P_{Y^\infty|\phi_{true}}$ -probability one.

Proof of Lemma A.3. Let B^* be an open neighborhood of ϕ_{true} as defined in the proof of Lemma A.2. Consider the marginal likelihood of model $M \in \mathcal{M}_B$ divided by $(2\pi)^{d/2}p(Y^n|\hat{\phi})n^{-d/2}\det(H_n(\hat{\phi}))^{1/2}$

$$\begin{aligned} \frac{n^{d/2}\det(H_n(\hat{\phi}))^{1/2}p(Y^n|M)}{(2\pi)^{d/2}p(Y^n|\hat{\phi})} &= \frac{n^{d/2}\det(H_n(\hat{\phi}))^{1/2}}{(2\pi)^{d/2}} \int_{\Phi_M} \exp\{n(l_n(\phi) - l_n(\hat{\phi}))\} f_{\phi|M}(\phi) d\phi \\ &\leq \frac{n^{d/2}\det(H_n(\hat{\phi}))^{1/2}}{(2\pi)^{d/2}} \bar{f}_{\phi|M} \sup_{\phi \in \Phi_M} \exp\{n(l_n(\phi) - l_n(\hat{\phi}))\} \\ &\leq \frac{n^{d/2}\det(H_n(\hat{\phi}))^{1/2}}{(2\pi)^{d/2}} \bar{f}_{\phi|M} \sup_{\phi \in \Phi \setminus B^*} \exp\{n(l_n(\phi) - l_n(\phi_{true}))\}, \end{aligned} \tag{A.12}$$

where $\bar{f}_{\phi|M} = \sup_{\phi} f_{\phi|M}(\phi) < \infty$, and the third line follows since $B^* \subset \Phi_M^c$ implies $\Phi_M \subset \Phi \setminus B^*$. Note that by Assumption 3.2 (iv), the upper bound shown in (A.12) converges to zero faster than the polynomial rate of $n^{-1/2}$ with $P_{Y^\infty|\phi_{true}}$ -probability one. ■

Proof of Proposition 3.3. (i) Under Assumption 3.2 (i), the posterior model probability of model $M \in \mathcal{M}$ can be written as

$$\pi_{M|Y^n} = \frac{p(Y^n|M)\pi_M}{\sum_{M' \in \mathcal{M}_A} p(Y^n|M')\pi_{M'} + \sum_{M' \in \mathcal{M}_B} p(Y^n|M')\pi_{M'}}$$

By dividing both the numerator and denominator by $(2\pi)^{d/2}p(Y^n|\hat{\phi})n^{-d/2}\det(H_n(\hat{\phi}))^{1/2}$ and applying Lemmas A.2 and A.3, we have

$$\pi_{M|Y^n} = \begin{cases} \frac{f_{\phi|M}(\hat{\phi})\pi_M}{\sum_{M' \in \mathcal{M}_A} f_{\phi|M'}(\hat{\phi})\pi_{M'}} + O(n^{-1/2}), & \text{for } M \in \mathcal{M}_A, \\ o(n^{-1/2}), & \text{for } M \in \mathcal{M}_B, \end{cases}$$

with $P_{Y^\infty|\phi_{true}}$ -probability one.

Since $f_{\phi|M}(\cdot)$ is assumed to be continuous and Assumptions 3.2 (iii) and (iv) imply almost sure convergence of $\hat{\phi}$ to ϕ_{true} , $\pi_{M|Y^\infty}$ of the current proposition follows.

(ii) With the given specifications of the ϕ -prior, $f_{\phi|M}(\phi_{true})$ is proportional to $\tilde{\pi}(\Phi_M)^{-1}$ up to the model-independent constant (the Lebesgue density of $\tilde{\pi}_\phi$ evaluated at $\phi = \phi_{true}$). Hence,

(i) of the current proposition is reduced to the asymptotic model probabilities of (ii).

(iii) This trivially follows from Lemma 3.1 (iii). ■

A.2 Example 2: Treatment Effect Model with an Instrument

This appendix illustrate applicability of our averaging proposal to the treatment effect model with noncompliance and a binary instrumental variable $Z \in \{0, 1\}$ (Imbens and Angrist (1994)).

Assume that the treatment status and the outcome of interest are both binary. Let $(W_1, W_0) \in \{1, 0\}^2$ be the potential treatment status in response to the instrument, and $W = ZW_1 + (1 - Z)W_0$ be the observed treatment status. $(Y_1, Y_0) \in \{1, 0\}^2$ is a pair of treated and control outcomes and $Y = WY_1 + (1 - W)Y_0$ is the observed outcome. Following Imbens and Angrist (1994), consider partitioning the population into four subpopulations defined in terms of the potential treatment-selection responses:

$$T = \begin{cases} c & \text{if } W_1 = 1 \text{ and } W_0 = 0 & : \text{complier,} \\ at & \text{if } W_1 = W_0 = 1 & : \text{always-taker,} \\ nt & \text{if } W_1 = W_0 = 0 & : \text{never-taker,} \\ d & \text{if } W_1 = 0 \text{ and } W_0 = 1 & : \text{defier,} \end{cases}$$

where T is the indicator for the types of selection responses.

Assume that the instrument is randomized in the sense that $Z \perp (Y_1, Y_0, W_1, W_0)$.²¹ Then, the distribution of observables and the distribution of potential outcomes satisfy the following equalities for $y \in \{1, 0\}$:

$$\begin{aligned} \Pr(Y = y, W = 1 | Z = 1) &= \Pr(Y_1 = y, T = c) + \Pr(Y_1 = y, T = at), \\ \Pr(Y = y, W = 1 | Z = 0) &= \Pr(Y_1 = y, T = d) + \Pr(Y_1 = y, T = at), \\ \Pr(Y = y, W = 0 | Z = 1) &= \Pr(Y_0 = y, T = d) + \Pr(Y_1 = y, T = nt), \\ \Pr(Y = y, W = 0 | Z = 0) &= \Pr(Y_0 = y, T = c) + \Pr(Y_1 = y, T = nt). \end{aligned} \tag{A.13}$$

Ruling out the marginal distribution of Z , the structural parameters index a joint distribution of (Y_1, Y_0, T) :

$$\theta = (\Pr(Y_1 = y, Y_0 = y', T = t) : y = 1, 0, \quad y' = 1, 0, \quad t = c, nt, at, d) \in \Theta,$$

where Θ is the 16-dimensional probability simplex.

Let the average treatment effect (ATE) be the parameter of interest.

$$\begin{aligned} \alpha &\equiv E(Y_1 - Y_0) = \sum_{t=c, nt, at, d} [\Pr(Y_1 = 1, T = t) - \Pr(Y_0 = 1, T = t)] \\ &= \sum_{t=c, nt, at, d} \sum_{y=1, 0} [\Pr(Y_1 = 1, Y_0 = y, T = t) - \Pr(Y_1 = y, Y_0 = 1, T = t)]. \end{aligned}$$

The reduced-form parameter vector consists of the eight probability masses:

$$\phi = (\Pr(Y = y, W = w | Z = z) : y = 1, 0, \quad d = 1, 0, \quad z = 1, 0).$$

Consider the following two candidate models.

Candidate Models

²¹As reflected in the notation of the potential outcomes (Y_1, Y_0) , we assume the exclusion restriction of the instrument.

- *Model M^p (point-identified)*: In addition to the randomized instrument assumption $Z \perp (Y_1, Y_0, W_1, W_0)$, the *instrument monotonicity* (no-defier) assumption of Imbens and Angrist (1994) holds and the causal effects are *homogeneous* in the sense that $E(Y_1 - Y_0|T = c) = E(Y_1 - Y_0|T = at) = E(Y_1 - Y_0|T = nt) = E(Y_1 - Y_0)$.
- *Model M^s (set-identified)*: The randomized instrument assumption holds. Heterogeneity of the treatment effects is unrestricted.

In model M^p , the complier's average treatment effect is identified by the Wald estimand (Imbens and Angrist (1994)), and combined with the homogeneity of the causal effects, we achieve the point-identification of ATE,

$$\alpha_{M^p}(\phi) = \frac{\Pr(Y = 1|Z = 1) - \Pr(Y = 1|Z = 0)}{\Pr(W = 1|Z = 1) - \Pr(W = 1|Z = 0)}.$$

In model M^s , what the Wald estimand identifies is the complier's average treatment effect, while ATE becomes set-identified. See Balke and Pearl (1997) for the construction of the ATE identified set, $IS_\alpha(\phi|M^s)$.

The two models considered admit the identical reduced-form (the distribution of $(Y, W)|Z$), whereas these two models are distinguishable, since they have different testable implications. The testable implication for model M^p is given by the testable implication for the joint restriction of randomized instrument and instrument monotonicity shown by Balke and Pearl (1997):²²

$$\begin{aligned} \Pr(Y = 1, D = 1|Z = 1) &\geq \Pr(Y = 1, D = 1|Z = 0), \\ \Pr(Y = 0, D = 1|Z = 1) &\geq \Pr(Y = 0, D = 1|Z = 0), \\ \Pr(Y = 1, D = 0|Z = 1) &\leq \Pr(Y = 1, D = 0|Z = 0), \\ \Pr(Y = 0, D = 0|Z = 1) &\geq \Pr(Y = 0, D = 0|Z = 0). \end{aligned}$$

Accordingly, Φ_{M^p} is given by the set of ϕ 's that satisfy these four inequalities.

Kitagawa (2009) shows that the instrument inequality of Pearl (1995) gives the sharp testable implication for the randomized instrument assumption, i.e., $IS_\alpha(\phi|M^s)$ is empty if and only if

$$\max_w \sum_y \max_z \{\Pr(Y = y, W = w)|Z = z\} \leq 1. \quad (\text{A.14})$$

Hence, the reduced-form parameter space of model M^s , Φ_{M^s} , is obtained as the set of ϕ 's that fulfills (A.14).

²²Under the joint restriction of randomized instrument and instrument monotonicity, additionally imposing homogeneity of the treatment effects does not strengthen the testable implication of Balke and Pearl (1997).

Set prior model probabilities at $(\pi_{M^p}, \pi_{M^s}) = (w, 1 - w)$. Construct a prior for ϕ in each model as

$$\pi_{\phi|M^p}(B) = \frac{\tilde{\pi}_{\phi}(B \cap \Phi_{M^p})}{\tilde{\pi}_{\phi}(\Phi_{M^p})},$$

$$\pi_{\phi|M^s}(B) = \frac{\tilde{\pi}_{\phi}(B \cap \Phi_{M^s})}{\tilde{\pi}_{\phi}(\Phi_{M^s})}.$$

for any measurable subset B in the probability simplex that ϕ lies, where $\tilde{\pi}_{\phi}$ is a prior for ϕ such as a Dirichlet distribution.

The two models M^p and M^s are distinguishable since Φ_{M^p} is a proper subset of Φ_{M^s} . With the current construction of the priors for ϕ , Lemma 3.1 (ii) gives their posterior model probabilities,

$$\pi_{M^p|Y} = \frac{O_{M^p} \cdot w}{O_{M^p} \cdot w + O_{M^s} \cdot (1 - w)},$$

$$\pi_{M^s|Y} = \frac{O_{M^s} \cdot (1 - w)}{O_{M^p} \cdot w + O_{M^s} \cdot (1 - w)},$$

where O_{M^p} and O_{M^s} are the posterior-prior plausibility ratio as defined in Lemma 3.1.

With these posterior model probabilities, the robust Bayes averaging operates as presented in Scenario 1 of Example 1. The resulting range of posterior means shrinks the Balke and Pearl's ATE identified set toward the posterior mean of the Wald estimand that one would report in the point-identified model. Since the posterior model probabilities can differ from the prior ones, the degree of shrinkage can reflect how well the identifying assumptions fit the data. The current analysis offers one way to aggregate the Wald instrumental variable estimator and the ATE bounds with exploiting a partially credible assumption on homogeneity of the causal effects.

References

- ALTONJI, J. G., T. E. ELDER, AND C. R. TABER (2005): "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy*, 113(1), 151–184.
- AN, S., AND F. SCHORFHEIDE (2007): "Bayesian analysis of DSGE models," *Econometric reviews*, 26(2-4), 113–172.
- ARUOBA, B., AND F. SCHORFHEIDE (2011): "Sticky Prices versus Monetary Frictions: An Estimation of Policy Trade-offs," *American Economic Journal: Macroeconomics*, 3, 60–90.
- ASHENFELTER, O. (1978): "Estimating the Effect of Training Programs on Earnings," *Review of Economics and Statistics*, 60, 47–57.

- BAJARI, P., H. HONG, AND S. P. RYAN (2010): “Identification and Estimation of a Discrete Game of Complete Information,” *Econometrica*, 78(5), 1529–1568.
- BALKE, A., AND J. PEARL (1997): “Bounds on Treatment Effects from Studies with Imperfect Compliance,” *Journal of the American Statistical Association*, 92, 1171–1176.
- BATES, J., AND C. GRANGER (1969): “The Combination of Forecasts,” *Operational Research Quarterly*, 20, 451–468.
- BAUMEISTER, C., AND J. HAMILTON (2015): “Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information,” *Econometrica*, 83, 1963–1999.
- BERESTEANU, A., I. MOLCHANOV, AND F. MOLINARI (2011): “Sharp Identification Regions in Models with Convex Moment Predictions,” *Econometrica*, 79(6), 1785–1821.
- BERGER, J., AND L. BERLINER (1986): “Robust Bayes and Empirical Bayes Analysis with ϵ -contaminated Priors,” *The Annals of Statistics*, 14, 461–486.
- BERNANKE, B. (1986): “Alternative Explorations of the Money-Income Correlation,” *Carnegie-Rochester Conference Series on Public Policy*, 25, 49–99.
- BLANCHARD, O., AND D. QUAH (1993): “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *American Economic Review*, 83, 655–673.
- CANOVA, F., AND G. D. NICOLO (2002): “Monetary Disturbances Matter for Business Fluctuations in the G-7,” *Journal of Monetary Economics*, 49, 1121–1159.
- CHERNOZHUKOV, V., H. HONG, AND E. TAMER (2007): “Estimation and Confidence Regions for Parameter Sets in Econometric Models,” *Econometrica*, 75, 1243–1284.
- CHIB, S., AND I. JELIAZKOV (2001): “Marginal Likelihoods from the Metropolis Hastings Output,” *Journal of the American Statistical Association*, 96, 270–281.
- CHIPMAN, H. (1996): “Bayesian variable selection with related predictors,” *The Canadian Journal of Statistics/La Revue Canadienne de Statistique*, pp. 17–36.
- CHRISTIANO, L., M. EICHENBAUM, AND C. EVANS (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): “Monetary Policy Shock: What Have We Learned and to What End?,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, Part A, pp. 65–148. Elsevier.
- CILIBERTO, F., AND E. TAMER (2009): “Market Structure and Multiple Equilibria in Airline Markets,” *Econometrica*, 77(6), 1791–1828.

- CLAESKENS, G., AND N. L. HJORT (2008): *Model Selection and Model Averaging*. Cambridge University Press, Cambridge, UK.
- CLYDE, M. (1999): “Bayesian model averaging and model search strategies,” *Bayesian statistics*, 6, 157.
- CLYDE, M., AND E. I. GEORGE (2004): “Model uncertainty,” *Statistical science*, pp. 81–94.
- DAWID, A. (1979): “Conditional Independence in Statistical Theory,” *Journal of the Royal Statistical Society. Series B*, 41, 1–31.
- DEDOLA, L., AND S. NERI (2007): “What Does a Technology Shock Do? A VAR Analysis with Model-Based Sign Restrictions,” *Journal of Monetary Economics*, 54(2), 512–549.
- DEL NEGRO, M., AND F. SCHORFHEIDE (2004): “Priors from general equilibrium models for vars,” *International Economic Review*, 45(2), 643–673.
- FAUST, J. (1998): “The Robustness of Identified VAR Conclusions about Money,” *Carnegie-Rochester Conference Series on Public Policy*, 48, 207–244.
- FUJITA, S. (2011): “Dynamics of Worker Flows and Vacancies: Evidence from the Sign Restriction Approach,” *Journal of Applied Econometrics*, 26(1), 89–121.
- GAFAROV, B., M. MEIER, AND J.-L. MONTIEL-OLEA (2016a): “Delta-Method Inference for a Class of Set-identified SVARs,” *unpublished manuscript*.
- (2016b): “Projection Inference for Set-identified SVARs,” *unpublished manuscript*.
- GEORGE, E. (1999): “Discussion of Model averaging and model search strategies by M. Clyde,” in *Bayesian Statistics 6—Proceedings of the Sixth Valencia International Meeting*.
- GEWEKE, J. (1999): “Using Simulation Methods for Bayesian Econometric Models: Inference, Development, and Communication,” *Econometric Reviews*, 18, 1–126.
- GIACOMINI, R., AND T. KITAGAWA (2015): “Robust Inference about Partially-identified SVARs,” *cemmap working paper*, University College London.
- GOOD, I. J. (1965): *The Estimation of Probabilities: An Essay on Modern Bayesian Methods*. The MIT Press.
- HANSEN, B. E. (2007): “Least Squares Model Averaging,” *Econometrica*, 75(4), 1175–1189.
- (2014): “Model Averaging, Asymptotic Risk, and Regressor Groups,” *Quantitative Economics*, 5(3), 495–530.

- HANSEN, B. E., AND J. S. RACINE (2012): “Jackknife Model Averaging,” *Journal of Econometrics*, 167, 38–46.
- HJORT, N. L., AND G. CLAESKENS (2003): “Frequentist Model Average Estimators,” *Journal of the American Statistical Association*, 98, 879–899.
- HOETING, J. A., D. MADIGAN, A. E. RAFTERY, AND C. T. VOLINSKY (1999): “Bayesian Model Averaging: A Tutorial,” *Statistical Science*, 14(4), 382–417.
- HUBER, P. (1973): “The Use of Choquet Capacities in Statistics,” *Bulletin of the International Statistical Institute*, 45, 181–191.
- IMBENS, G., AND J. ANGRIST (1994): “Identification and Estimation of Local Average Treatment Effects,” *Econometrica*, 62, 467–475.
- KASS, R. E., L. TIERNEY, AND J. B. KADANE (1990): “The Validity of Posterior Expansions Based on Laplace Method,” in *Bayesian and Likelihood Methods in Statistics and Econometrics*, ed. by S. Geisser, J. Hodges, S. Press, and A. Zellner, pp. 473–488. Elsevier, North-Holland.
- KILIAN, L. (2013): “Structural Vector Autoregressions,” in *Handbook of Research Methods and Applications in Empirical Macroeconomics*, ed. by M. Hashimzade, N. & Thornton, pp. 515–554. Edward Elgar, Cheltenham, UK.
- KILIAN, L., AND D. MURPHY (2012): “Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models,” *Journal of the European Economic Association*, 10, 1166–1188.
- KITAGAWA, T. (2009): “Identification Region of the Potential Outcome Distributions under Instrument Independence,” *unpublished manuscript*.
- KITAGAWA, T. (2012): “Estimation and Inference for Set-identified Parameters Using Posterior Lower Probabilities,” *unpublished manuscript*.
- KITAGAWA, T., AND C. MURIS (2016): “Model Averaging in Semiparametric Estimation of Treatment Effects,” *Journal of Econometrics*, 193(1), 271–289.
- KLINE, B., AND E. TAMER (2016): “Bayesian Inference in a Class of Partially Identified Models,” *Quantitative Economics*, 7(2), 329–366.
- LEAMER, E. E. (1978): *Specification Searches*. Wiley, New York.
- LIAO, Y., AND A. SIMONI (2013): “Semi-parametric Bayesian Partially Identified Models based on Support Function,” *unpublished manuscript*.

- LIU, C.-A. (2015): “Distribution Theory of the Least Squares Model Averaging Estimator,” *Journal of Econometrics*, 186, 142–159.
- LIU, Q., AND R. OKUI (2013): “Heteroskedasticity-robust C_p Model Averaging,” *Econometrics Journal*, 16(3), 463–472.
- MAGNUS, J. R., O. POWELL, AND P. PRÜFER (2010): “A Comparison of Two Model Averaging Techniques with an Application to Growth Empirics,” *Journal of Econometrics*, 154, 139–153.
- MANSKI, C. F. (1989): “Anatomy of the Selection Problems,” *Journal of Human Resources*, 24, 343–360.
- MANSKI, C. F., AND J. V. PEPPER (2000): “Monotone Instrumental Variables: With an Application to the Returns to Schooling,” *Econometrica*, 68(4), 997–1010.
- MOON, H., AND F. SCHORFHEIDE (2011): “Bayesian and Frequentist Inference in Partially Identified Models,” *NBER working paper*.
- (2012): “Bayesian and Frequentist Inference in Partially Identified Models,” *Econometrica*, 80, 755–782.
- MOON, H., F. SCHORFHEIDE, AND E. GRANZIERA (2013): “Inference for VARs Identified with Sign Restrictions,” *unpublished manuscript*.
- MOUNTFORD, A. (2005): “Leaning into the Wind: A Structural VAR Investigation of UK Monetary Policy,” *Oxford Bulletin of Economics and Statistics*, 67(5), 597–621.
- MOUNTFORD, A., AND H. UHLIG (2009): “What Are the Effects of Fiscal Policy Shocks?,” *Journal of applied econometrics*, 24(6), 960–992.
- NORETS, A., AND X. TANG (2014): “Semiparametric Inference in Dynamic Binary Choice Models,” *Review of Economic Studies*, 81(3), 1229–1262.
- PEARL, J. (1995): “On the Testability of Causal Models with Latent and Instrumental Variables,” *Uncertainty in Artificial Intelligence*, 11, 435–443.
- PEERSMAN, G., AND R. STRAUB (2009): “Technology Shocks and Robust Sign Restrictions in a Euro Area SVAR,” *International Economic Review*, 50(3), 727–750.
- POIRIER, D. (1998): “Revising Beliefs in Nonidentified Models,” *Econometric Theory*, 14, 483–509.
- RAFIQ, S., AND S. MALICK (2008): “The Effect of Monetary Policy on Output in EMU3: A Sign Restriction Approach,” *Journal of Macroeconomics*, 30(4), 1756–1791.

- REICHENSTEIN, W. (1987): “The Impact of money on short-term interest rates,” *Economic Inquiry*, 25(1), 67–82.
- ROSENBAUM, P., AND D. B. RUBIN (1983): “The Central Role of the Propensity Score in Observational Studies,” *Biometrika*, 70, 41–55.
- RUBIN, D. B. (1987): *Multiple Imputation for Nonresponse in Surveys*. John Wiley & Sons, Hoboken, New Jersey.
- SCHOLL, A., AND H. UHLIG (2008): “New Evidence on the Puzzles: Results from Agnostic Identification on Monetary Policy and Exchange Rates,” *Journal of International Economics*, 76(1), 1–13.
- SIMS, C. (1980): “Macroeconomics and Reality,” *Econometrica*, 48, 1–48.
- (1992): “Interpreting the macroeconomic time series facts: The effects of monetary policy,” *European Economic Review*, 36(5), 975–1000.
- SIMS, C., D. WAGGONER, AND T. ZHA (2008): “Methods for Inference in Large Multiple-Equation Markov-Switching Models,” *Journal of Econometrics*, 146, 255–274.
- SMETS, F., AND R. WOUTERS (2003): “An estimated dynamic stochastic general equilibrium model of the euro area,” *Journal of the European economic association*, 1(5), 1123–1175.
- UHLIG, H. (2005): “What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure,” *Journal of Monetary Economics*, 52, 381–419.
- VARGAS-SILVA, C. (2008): “Monetary Policy and the US Housing Market: A VAR Analysis Imposing Sign Restrictions,” *Journal of Macroeconomics*, 30(3), 977–990.
- ZHANG, X., AND H. LIANG (2011): “Focused Information Criterion and Model Averaging for Generalized Additive Partial Linear Models,” *Annals of Statistics*, 39, 174–200.

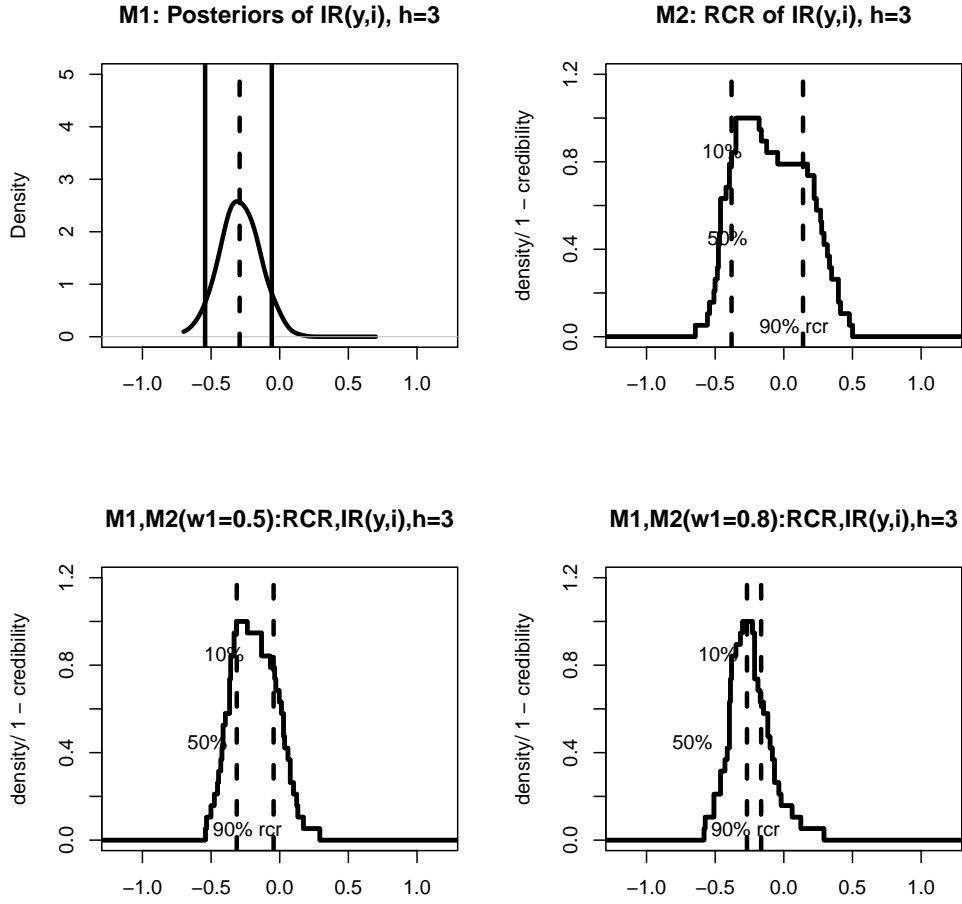


Figure 1: Density and Robust Credible Region of Output Impulse Responses

Note: Output Impulse Response at horizon $h = 3$. For set-identified models, step lines represent the Robust Credible Region (RCR) at different credibility levels (90%, 50%, 10% levels are explicitly indicated) as described in the last paragraph of Section 2.1 by modifying (Step 5) of Algorithm 4.1 in Giacomini and Kitagawa (2015). The vertical dashed lines represent the posterior mean bounds. For point-identified models (Model 1 and Model 4 in Figure 3), the vertical solid lines display the standard credible region. In such a case, we reported its posterior density.

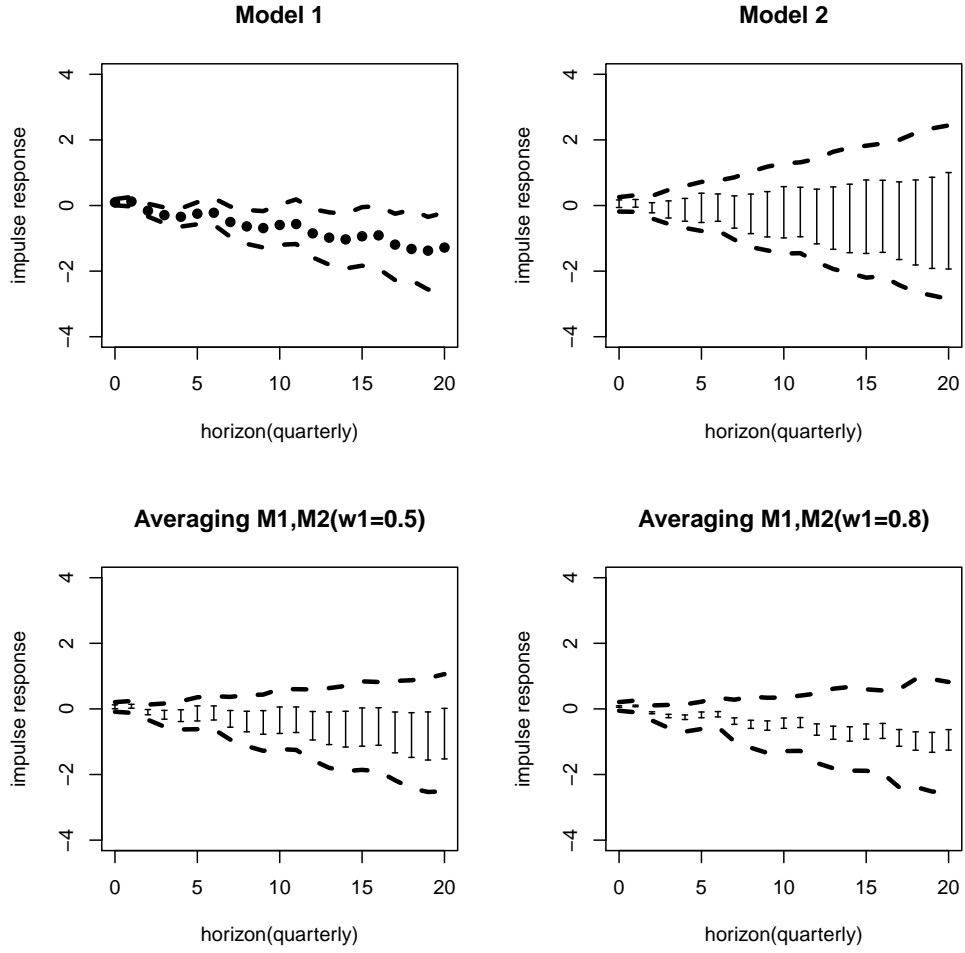


Figure 2: Plots of Output Impulse Responses

Note: for point-identified models, the points plot the (unique) posterior mean and the dashed curve represent the highest posterior density regions with credibility 90%. For set-identified models (Model 2, the averaged models and Model 3 in Figure 4), the vertical bars show the posterior mean bounds and the dashed curves connect the upper/lower bounds of posterior robust credible regions with credibility 90%.

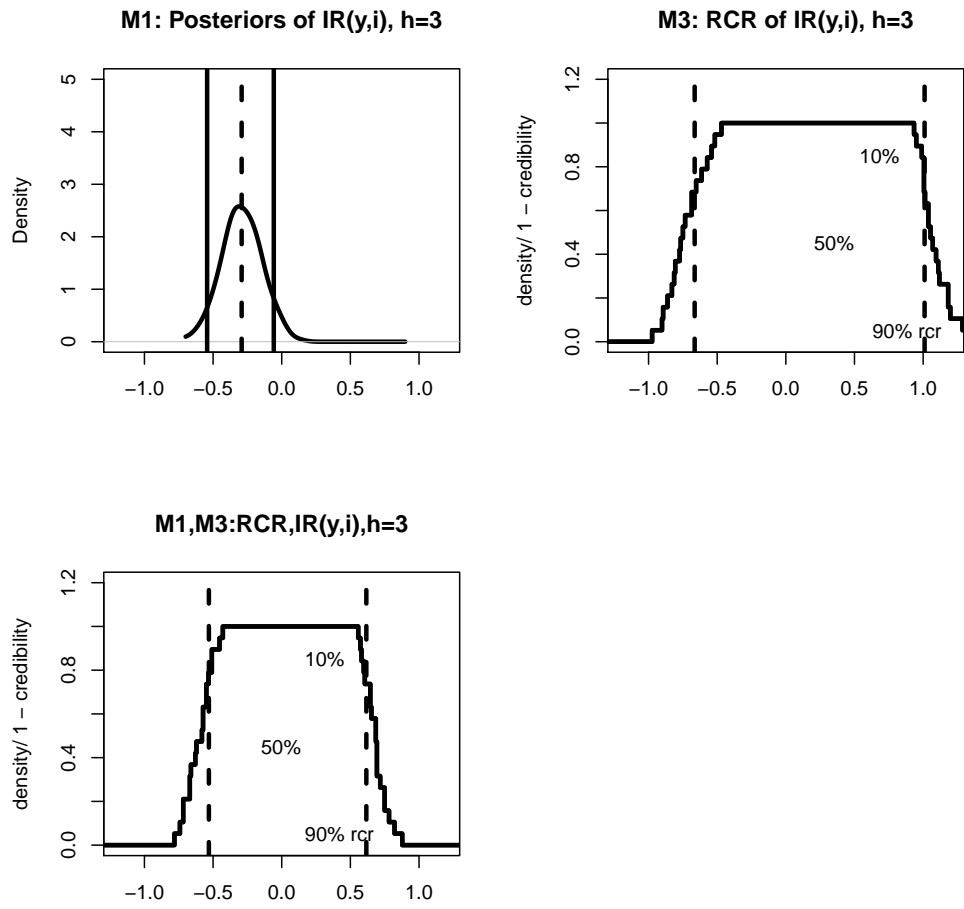


Figure 3: Density and Robust Credible Region of Output Impulse Responses

See the caption of Figure 1 for remarks.

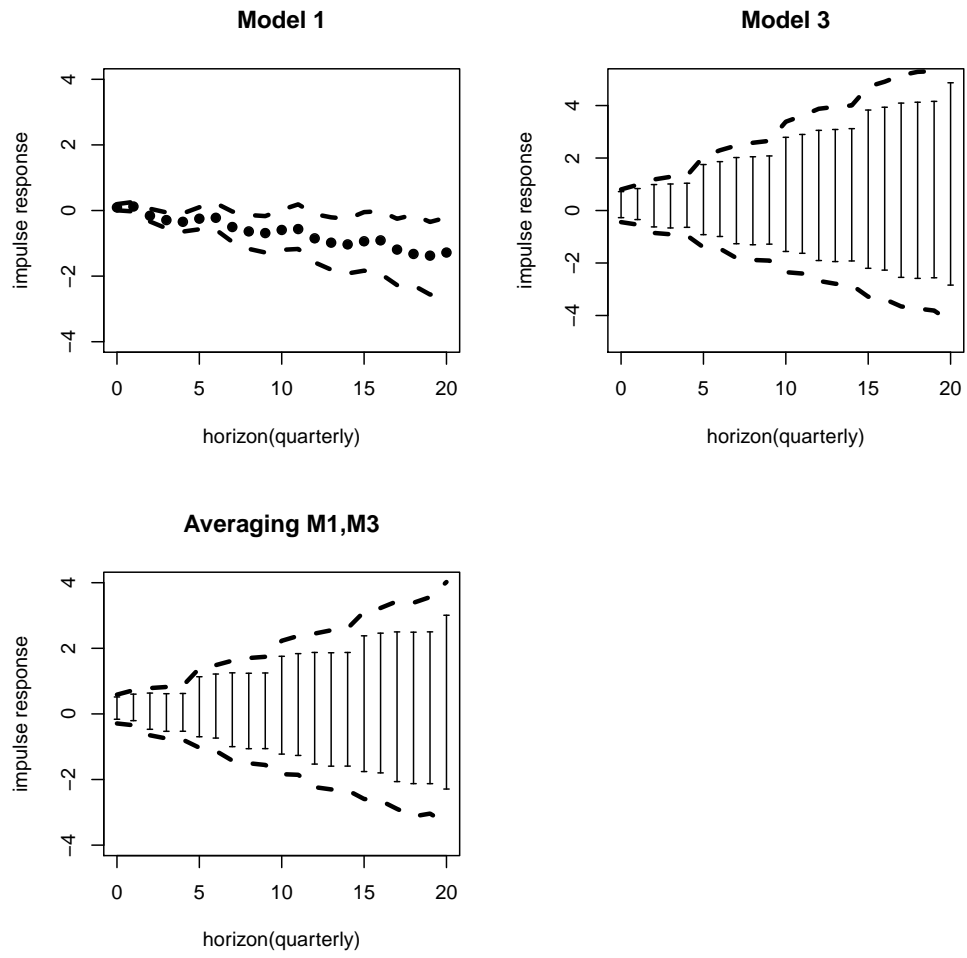


Figure 4: Plots of Output Impulse Responses

See the caption of Figure 2 for remarks.

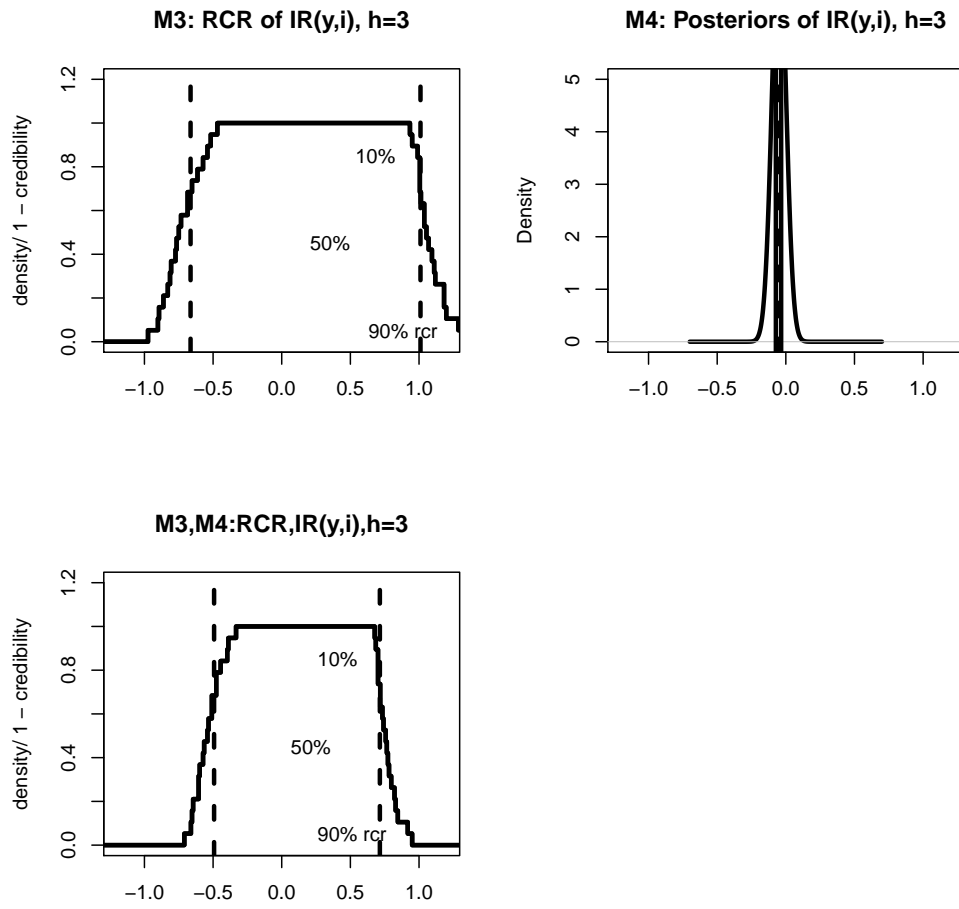


Figure 5: Density and Robust Credible Region of Output Impulse Responses

See the caption of Figure 1 for remarks.

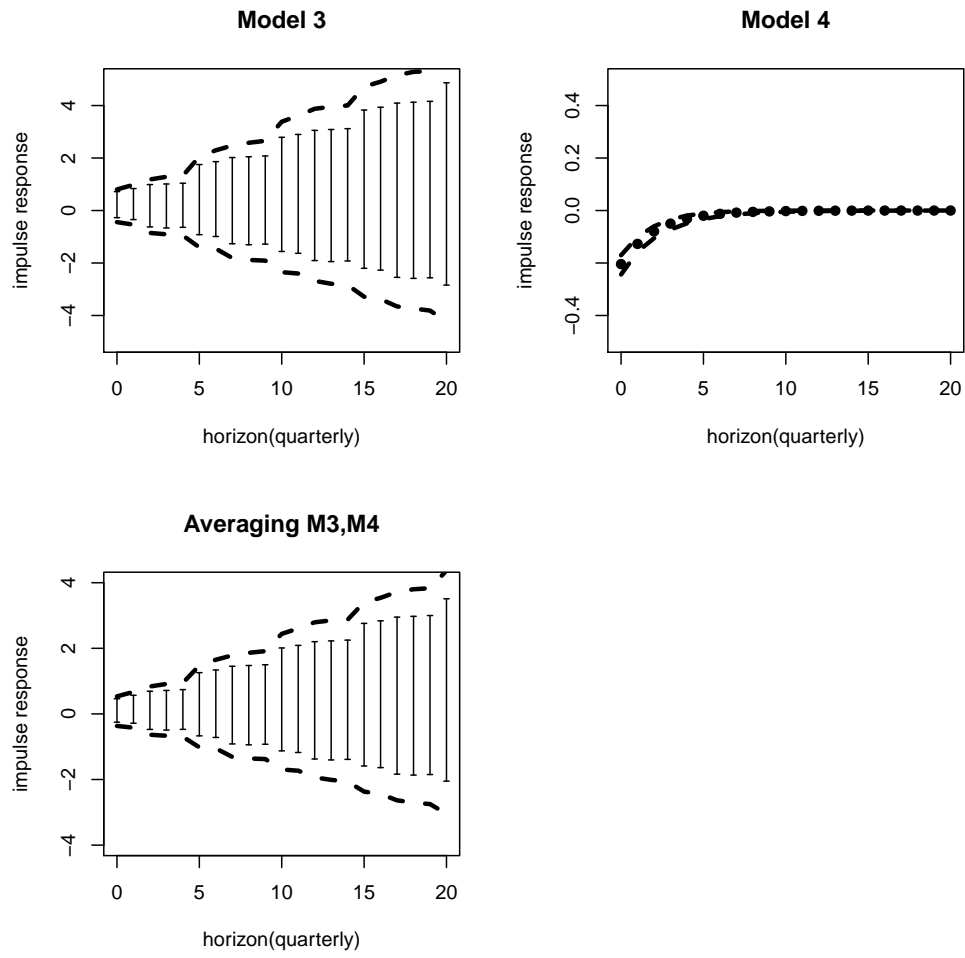


Figure 6: Plots of Output Impulse Responses

See the caption of Figure 2 for remarks.

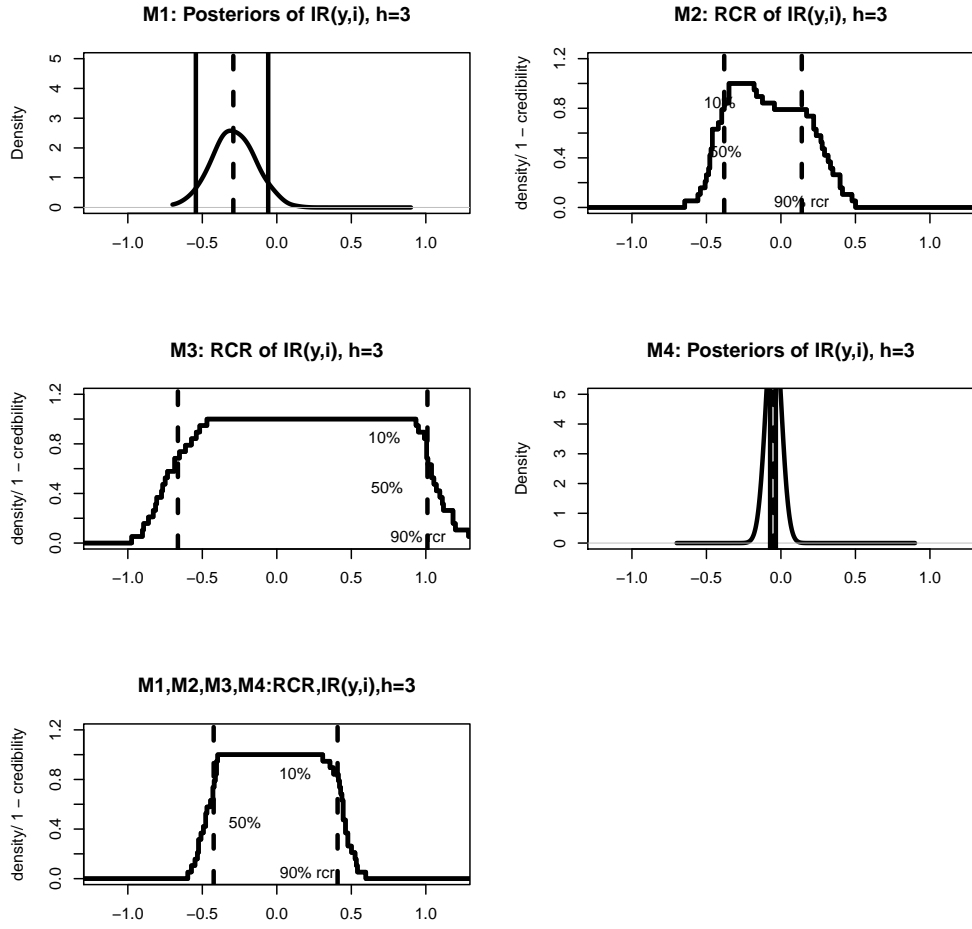


Figure 7: Density and Robust Credible Region of Output Impulse Responses

See the caption of Figure 1 for remarks.

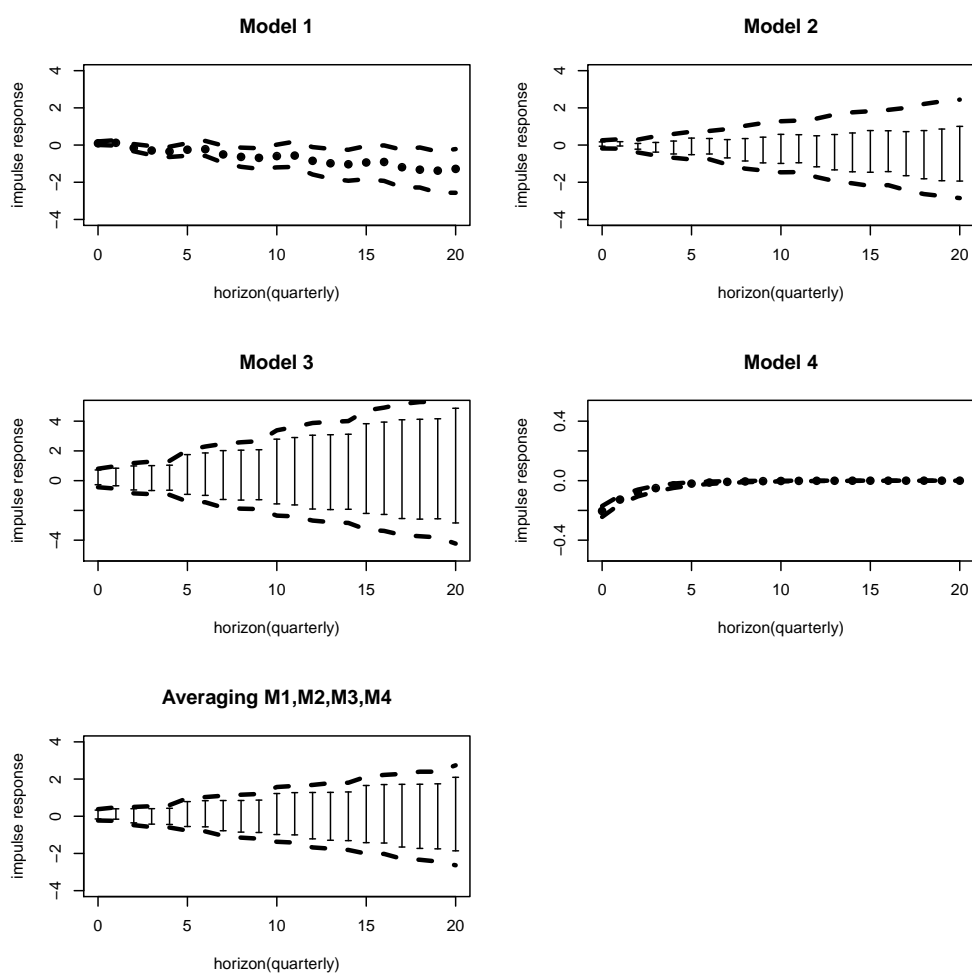


Figure 8: Plots of Output Impulse Responses

See the caption of Figure 2 for remarks.

	Averaging M1, M2	Averaging M1,M2	Averaging M1,M3	Averaging M3,M4	Averaging M1,M2,M3,M4
Prior w_1	0.50	0.80	0.50	/	0.25
Prior w_2	0.50	0.20	/	/	0.25
Prior w_3	/	/	0.50	0.50	0.25
Prior w_4	/	/	/	0.50	0.25
O_1	1	1	1	/	1
O_2	1	1	/	/	1
O_3	/	/	2.16	2.16	2.16
O_4	/	/	/	1	1
$\ln \tilde{p}(Y)$	-781.05	-781.05	-781.05	-781.05	-781.05
$\ln p(Y M^1)$	-781.05	-781.05	-781.05	/	-781.05
$\ln p(Y M^4)$	/	/	/	-781.29	-781.29
Posterior w_1^*	0.50	0.80	0.32	/	0.20
Posterior w_2^*	0.50	0.20	/	/	0.20
Posterior w_3^*	/	/	0.68	0.73	0.44
Posterior w_4^*	/	/	/	0.27	0.16

Table 1: Output Responses: Prior and Posterior Weights

Note: prior w_i , O_i and posterior w_i^* denote prior model probability, posterior-prior credibility ratio and posterior model probability for candidate Model i , respectively; $\ln \tilde{p}(Y)$, $\ln p(Y|M^1)$ and $\ln p(Y|M^4)$ represent log marginal likelihood for the common reduced form, for Model 1 and for Model 4, respectively.

	M1			M2		
	$h = 1$	$h = 10$	$h = 20$	$h = 1$	$h = 10$	$h = 20$
Post. Mean	.122	-.593	-1.281	/	/	/
90% CR	[-.029, .261]	[-1.200, .022]	[-2.566, -.215]	/	/	/
Post. Mean Bounds	/	/	/	[-.050, .183]	[-.987, .575]	[-1.936, 1.005]
90% robustified CR	/	/	/	[-.193, .313]	[-1.465, 1.281]	[-2.851, 2.443]
	M3			M4		
	$h = 1$	$h = 10$	$h = 20$	$h = 1$	$h = 10$	$h = 20$
Post. Mean	/	/	/	-.127	-.002	-.000
90% CR	/	/	/	[-.157, -.100]	[-.004, -.001]	[-.000, -.000]
Post. Mean Bounds	[-.344, .836]	[-1.561, 2.789]	[-2.843, 4.867]	/	/	/
90% robustified CR	[-.538, .979]	[-2.350, 3.384]	[-4.231, 6.051]	/	/	/
	Averaging M1,M2($w1 = 0.5$)			Averaging M1,M2($w1 = 0.8$)		
	$h = 1$	$h = 10$	$h = 20$	$h = 1$	$h = 10$	$h = 20$
Post. Mean	/	/	/	/	/	/
90% CR	/	/	/	/	/	/
Post. Mean Bounds	[.021, .141]	[-.748, .059]	[-1.526, .016]	[.067, .114]	[-.592, -.278]	[-1.257, -.631]
90% robustified CR	[-.121, .241]	[-1.224, .607]	[-2.512, 1.061]	[-.103, .255]	[-1.285, .347]	[-2.528, .821]
	Averaging M1,M3			Averaging M3,M4		
	$h = 1$	$h = 10$	$h = 20$	$h = 1$	$h = 10$	$h = 20$
Post. Mean	/	/	/	/	/	/
90% CR	/	/	/	/	/	/
Post. Mean Bounds	[-.205, .603]	[-1.225, 1.756]	[-2.291, 3.008]	[-.284, .568]	[-1.126, 2.010]	[-2.049, 3.509]
90% robustified CR	[-.346, .722]	[-1.837, 2.229]	[-3.339, 4.020]	[-.421, .669]	[-1.695, 2.440]	[-3.047, 4.369]
	Averaging M1,M2,M3,M4					
	$h = 1$	$h = 10$	$h = 20$			
Post. Mean	/	/	/			
90% CR	/	/	/			
Post. Mean Bounds	[-.158, .400]	[-.983, 1.222]	[-1.858, 2.095]			
90% robustified CR	[-.248, .481]	[-1.373, 1.573]	[-2.635, 2.739]			

Table 2: Output Responses: Estimation and Inference