Federal Reserve Bank of Dallas

When Is the Use of Gaussian-inverse Wishart-Haar Priors Appropriate?

Atsushi Inoue and Lutz Kilian

Working Paper 2404 **July 2024**

Research Department

https://doi.org/10.24149/wp2404

When Is the Use of Gaussian-inverse Wishart-Haar Priors Appropriate?*

Atsushi Inoue† and Lutz Kilian‡

July 7, 2024

Abstract

Several recent studies have expressed concern that the Haar prior typically employed in estimating sign-identified VAR models is driving the prior about the structural impulse responses and hence their posterior. In this paper, we provide evidence that the quantitative importance of the Haar prior for posterior inference has been overstated. How sensitive posterior inference is to the Haar prior depends on the width of the identified set of a given impulse response. We demonstrate that this width depends not only on how much the identified set is narrowed by the identifying restrictions imposed on the model, but also depends on the data through the reduced-form model parameters. Hence, the role of the Haar prior can only be assessed on a case-by-case basis. We show by example that, when the identification is sufficiently tight, posterior inference based on a Gaussian-inverse Wishart-Haar prior provides a reasonably accurate approximation.

JEL Codes: C22, C32, C52, E31

Keywords: Bayesian VAR, impulse response, sign restrictions, set identification, Haar prior

---

*The views expressed in the paper are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. We have benefited from comments from the editor, three anonymous referees, Juan Antolin-Diaz, Pooyan Amir Ahmadi, Thorsten Drautzburg, Juan Rubio-Ramírez, Neil Shephard, and Benjamin Wong. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

†Atsushi Inoue, Vanderbilt University, Department of Economics, Nashville, TN 37235-1819, E-mail: atsushi.inoue@vanderbilt.edu.

‡Corresponding author: Lutz Kilian, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA and CEPR, E-mail: lkilian2019@gmail.com.

1 Introduction

The conventional Bayesian approach to estimating VAR models subject to sign restrictions on the impulse responses involves specifying a conjugate Gaussian-inverse Wishart-Haar prior for the reduced-form slope parameters, A, the reduced-form error covariance matrix, Σ, and the orthogonal rotation matrix Q.[1] The prior for the scalar impulse response parameter $\theta = g(A, \Sigma, Q)$, is defined implicitly by the nonlinear function $g(\cdot)$.[2] In the asymptotic limit, the posterior distribution of θ over the identified set depends only on the prior for Q, since the likelihood is flat over the identified set. A number of recent studies have questioned the extent to which the impulse response posterior in these models is driven by the choice of a Haar prior for Q (e.g., Baumeister and Hamilton 2015, Giacomini and Kitagawa 2021, Giacomini, Kitagawa and Read 2023; Kilian 2013; Plagborg-Møller 2019; Watson 2020). Baumeister and Hamilton (2015) argue against the use of the conventional approach to posterior inference about sign-identified impulse responses on the grounds that it amounts to "pretending that [this] prior information ... has no effect on the reported conclusions (p. 1993)." In this paper, we provide evidence that the quantitative importance of the Haar prior for posterior inference has been overstated.

Our analysis clarifies the conditions under which the conventional approach to estimating sign-identified VAR models remains a viable option in applied work. We demonstrate that there is no presumption that the substantive conclusions of sign-identified models estimated using the conventional approach are necessarily spurious. The key question we address is how much the choice of the prior over the set-identified impulse response parameters matters, given the reduced-form posterior and the identifying sign restrictions. Our analysis recognizes that the identifying sign

[1] Prominent early applications of sign-identified VAR models include Faust (1998), Uhlig (1998), and Canova and de Nicolo (2002). The use of the Gaussian-inverse Wishart prior was proposed by Uhlig (2005). The Haar prior was introduced by Rubio-Ramírez, Waggoner and Zha (2010). For a review of the literature see Kilian and Lütkepohl (2017) and Uhlig (2017).

[2] More formal explanations of this notation can be found in the online appendix.

restrictions are part of the impulse response prior. The more informative these restrictions are, given the data, the smaller the identified set. We show that the smaller the identified set, the less misleading are the impulse response estimates and credible sets obtained using the conventional approach. In practice, the width of the identified set for a given impulse response depends not only on the identifying restrictions imposed on this response, but also on the identifying restrictions imposed on the other structural responses in the VAR model. In addition, the width of the identified set depends on the data through Σ and A, given the identifying restrictions. Hence, the role of the Haar prior can only be assessed on a case-by-case basis.

Our point is not that the prior for Q does not influence the prior for θ. Nor do we dispute that the posterior of θ is affected by the Haar prior. Rather our point is that this influence tends to be negligible in structural VAR models that are tightly identified by sign restrictions (possibly in conjunction with other restrictions such as elasticity bounds and narrative, shape or exclusion restrictions). As long as the identified set is narrow enough, the posterior distribution over the identified set, as implied by the Haar prior, makes no material difference for the impulse response estimates from an economic point of view.

We provide diagnostic tools for sign-identified VAR models that help applied users assess how much of a concern the use of the Haar prior is in a given application. Our approach builds on the estimate of the identified set of the impulse response proposed in Giacomini and Kitagawa (2021) and Giacomini, Kitagawa and Read (2023) and the corresponding robust credible interval. When the identification is sufficiently informative, given the data, the estimate of the identified set will be narrow enough for conventional inference to be insensitive to the Haar prior. In this case, the Bayes estimate of the impulse response based on the Haar prior will be close to the bounds of the identified set, as measured by the Hausdorff distance. Moreover, by the same metric, the endpoints of conventional credible intervals based on the Haar prior for Q will come close to the endpoints of Q

prior-robust credible intervals. Provided the Hausdorff distance between these sets is economically negligible, the influence of the Haar prior may be ignored. A similar approach may also be applied to the identified set of posterior probabilities. Compared to alternative metrics, one advantage of the Hausdorff distance is that it is measured in the same units as the impulse response, facilitating the interpretation of the results. The other advantage is that it accounts for both the width and location of the identified set in relation to posterior estimates based on the Haar prior.

These diagnostics allow us to quantify the influence of the Haar prior on posterior inference in several empirical examples drawn from the literature, starting with sign-identified VAR models of monetary policy shocks. We show that the conventional approach is indeed problematic in agnostic models such as Uhlig (2005) because the identified set is very wide, but these concerns greatly diminish when imposing additional narrative restrictions as in Antolin-Diaz and Rubio-Ramírez (2018). We also revisit the stylized model of monetary policy examined in Baumeister and Hamilton (2018). We find that adding sign and narrative restrictions reverses their substantive conclusions about the conventional approach being misleading. Finally, we use a recent study of the impact of gasoline price shocks on inflation and inflation expectations to demonstrate that even in the absence of narrative restrictions the identified set may be narrow enough for the conventional approach to be insensitive to the Haar prior. These examples illustrate the importance of assessing the influence of the Haar prior on a case-by-case basis.

The remainder of the paper is organized as follows. In Section 2, we study a stylized statistical model with one set-identified parameter proposed by Watson (2020) under a range of priors. We show that how sensitive posterior inference in this model is to the prior depends on how tightly identified the parameter of interest is, which explains why sometimes the prior unduly influences the posterior, as suggested by Baumeister and Hamilton (2015), and sometimes it does not. How tight the identification is may be judged by the width of the identified set of the responses or, alternatively, by the bounds on the posterior probability that the parameter of interest is positive or negative, respectively. While the posterior probability bounds in the Watson (2020) model approach $[0, 1]$ as $n \to \infty$, we find that for finite n these bounds are informative, provided the width of identified set is sufficiently narrow. Building on Uhlig (2017), Section 3 shows that how tightly identified the impulse responses in sign-identified VAR models are, depends not only on the identifying restrictions, but also on the data. In Section 4, we outline how applied users can assess the sensitivity of conventional estimates of sign-identified VAR models to the Haar prior, building on Giacomini and Kitagawa (2021) and Giacomini et al. (2023). Section 5 contains several empirical illustrations that illustrate that the conventional approach to posterior inference about impulse responses based on the Haar prior need not be driven by the Haar prior as well as other empirical examples that show that conventional inference can be misleading. The concluding remarks are in Section 6.

2 A stylized example of the role of the prior for set-identified parameters

To better understand the concern about the Haar prior, consider a stylized model with a generic set-identified scalar parameter, θ, building on Watson (2020). The observed data are independent normal random variables $Y \sim N(\mu, I_2/n)$, where $Y = (Y_1, Y_2)'$ and $\mu = (\mu_1, \mu_2)'$. There are three model parameters, (μ_1, μ_2, θ), where $\{(\mu_1, \mu_2) \in \mathbb{R}^2 : \mu_1 < \mu_2\}$. The parameter of interest, θ, satisfies the sign restrictions $\mu_1 \leq \theta \leq \mu_2$. We assume that the prior for μ_1 and μ_2 is a constant over its support. As shown in the online appendix, because the likelihood does not depend on θ, the posterior for θ is:

$$p(\theta|Y) = \int_{\mu_1<\mu_2} p(\theta|\mu_1, \mu_2, Y_1, Y_2) p(\mu_1, \mu_2|Y_1, Y_2) d(\mu_1, \mu_2) = \int_{\mu_1<\mu_2} p(\theta|\mu_1, \mu_2) p(\mu_1, \mu_2|Y_1, Y_2) d(\mu_1, \mu_2).$$

The question of interest is the effect of the prior $p(\theta|\mu)$ on the posterior $p(\theta|Y)$. For expository purposes, consider four alternative priors for θ conditional on μ_1 and μ_2:

Prior 1: $\theta \sim U(\mu_1, \mu_2)$

Prior 2: $\theta \sim N(0,1)$ truncated at μ_1 and μ_2 and rescaled to integrate to 1

Prior 3: $\theta = \mu_1$ with probability 0.99 and $\theta = \mu_2$ with probability 0.01.

Prior 4: $\theta = \mu_1$ with probability 0.01 and $\theta = \mu_2$ with probability 0.99.

Whereas prior 2 assigns most of the probability mass to the center of the identified set, priors 3 and 4 put most of the probability mass on the endpoints μ_1 and μ_2, respectively.[3]

If the posterior $p(\mu|Y)$ assigns most of its probability mass to values of (μ_1, μ_2) within a narrow identified set, the data suggest that θ lies in this narrow range of values. In the words of Watson (2020, p. 186), if the data indicate that μ_1 is close to μ_2, then the restriction $\mu_1 \leq \theta \leq \mu_2$ says a lot about the value of θ. From an economic point of view, the relative probability of values of θ in this narrow range, which is governed by the prior, might be unimportant because different values are substantively identical. In practice, the more important conclusion might be that θ falls into the specified range. In contrast, if $p(\mu|Y)$ assigns substantial probability mass to values of (μ_1, μ_2) within a wide identified set, the prior $p(\theta|\mu)$ is likely to have an important effect on the posterior $p(\mu|Y)$. In that case, $p(\theta|Y)$ may be concentrated in a narrow range simply because of the choice of prior $p(\theta|\mu)$ and caution is called for in interpreting the posterior.

This point is illustrated in Figure 1, which shows the posterior distribution of θ for each of the four priors above, for $n = 100$ and $n = \infty$.[4] By construction, the posterior variance of $\mu_i, i = 1, 2$, declines with rising n. Thus, as $n \to \infty$, the posterior distribution of μ_1 and μ_2 collapses on Y_1 and Y_2. For illustrative purposes, we evaluate the density at two data points: $Y_2 = -Y_1 = 2.5$ in

[3] We thank one of the referees for suggesting priors 3 and 4.

[4] The number of posterior draws of μ_1 and μ_2 is 10 million before imposing the restriction $\mu_1 < \mu_2$.

panel (a) and $Y_2 = -Y_1 = 0.05$ in panel (b). Figure 1 provides two key insights. First, the extent to which the prior for θ matters for posterior inference depends on the width of the identified set. When μ_1 and μ_2 are far apart, the posterior over the identified set will asymptotically resemble the prior over the identified set (see $n = \infty$ in panel (a)). This is the situation that Baumeister and Hamilton (2015) had in mind. When μ_1 and μ_2 are sufficiently close, in contrast, the fact that the posterior over the identified set asymptotically resembles the prior becomes irrelevant to the extent that the posterior draws contained in the identified set are sufficiently similar (see $n = \infty$ in panel (b)).[5] In the latter case, reporting the median or mean of the posterior as a measure of its central tendency and the corresponding credible intervals will yield similar estimates, regardless of the prior for θ. Even when $n = 100$ in panel (b) the prior has much less influence on the posterior than in panel (a). How narrow the identified set of θ has to be for the estimates to be sufficiently similar depends on the economic application. In some cases, numerically small differences may be economically significant, whereas the same difference in other cases may be immaterial. We will return to this point later.[6]

For illustrative purposes, Table 1 focuses on the posterior mean of θ. When $(Y_1, Y_2) = (-2.5, 2.5)$, the estimated identified set is wide, and the posterior means are sensitive to the prior. Table 1(a) shows that even asymptotically in this case the posterior mean of θ differs from the bounds μ_1 and μ_2 of the identified set by as much as 5.0. We call this the *maximum error* committed when reporting the posterior mean. This maximum error is the Hausdorff distance between the posterior estimate of the identified set and the posterior mean.[7] When $(Y_1, Y_2) = (-0.05, 0.05)$, in contrast,

[5] This argument does not depend on the posterior densities being locally similar within the identified set for $n = \infty$. Indeed, in our illustrative example the posterior under priors 3 and 4 is dramatically different from the posterior under prior 1 or 2 in the relevant range. All that matters is that the width of the identified set is narrow.

[6] Additional results for the posterior distribution are available in Table A.1 in the online appendix which reports the posterior estimate of the identified set and the mean width of this set, $\mu_2 - \mu_1$, for $n \in \{100, \infty\}$ as well as the posterior mean and the 90% credible interval for θ for each of the four priors. The evidence in Table A.1 confirms that, notwithstanding the readily apparent differences across credible intervals and posterior means constructed under different priors, these differences tend to be numerically much smaller when the identified set is tight, as in the lower panel of Figure 1, than when it is wide as in the upper panel.

[7] The Hausdorff distance measures how far two subsets of a metric space are from each other. Denote the endpoints

the situation changes. The maximum error is substantially reduced. While each prior puts different weights on the values in the identified set, the range of these values is so narrow that a user arguably may not be concerned about the difference between them. In the limit, for $n = \infty$, the identified set, $\theta \in [-0.05, 0.05]$ is known. In that case, under prior 1, the posterior mean in Table 1(b) differs from the bounds of identified set at most by 0.05 and, similarly, under prior 2. Under priors 3 and 4, that discrepancy rises to 0.1. This maximum error is smaller than in the upper panel of Figure 1 by a factor of 50. In practice, the user must decide whether the Hausdorff distance is negligible from an economic point of view in a given application. If it is, the posterior mean provides an adequate summary of what we learn from the data. Otherwise, it does not. The empirical examples in Section 5 illustrate that the answer to this question more often than not is readily apparent.

In Section 4, we discuss how this approach can be implemented in sign-identified VAR models, building on the work of Giacomini amd Kitagawa (2021) and Giacomini et al. (2023). A similar approach may also be applied to credible intervals for the parameter of interest. In the context of structural VAR analysis, one computes the maximum error of the endpoints of the conventional credible interval based on the Haar prior relative to the endpoints of a robust credible interval that does not take a stand on the Q prior.

It should be noted that the width of the identified set of θ is not the only way to assess the sensitivity of the posterior to the choice of the prior. Another approach is to focus on posterior probabilities such as the posterior probability that θ is positive or negative, respectively (see, e.g., Giacomini and Kitagawa 2021). Suppose that one is interested in $P(\theta \geq 0|Y)$. There are three situations of interest. If both μ_1 and μ_2 are on the same side of zero, then asymptotically this probability is 0 or 1, regardless of the prior for θ. If $\mu_1 \leq 0 \leq \mu_2$, as in the Watson (2020) example,

of these sets by (α, β) and (γ, δ), respectively. Then the Hausdorff distance is $\max([|\alpha - \gamma|, |\beta - \delta|])$.

much depends on n. As $n \to \infty$, the support of the probability may be as large as $[0,1]$, depending on the prior $p(\theta|\mu)$. However, for a given value of n, that probability has a smaller range defined by the endpoints:

$$\left[\int_0^\infty \frac{\sqrt{n}}{c}\phi\left(\sqrt{n}(\mu_1 - Y_1)\right)\Phi\left(-\sqrt{n}(\mu_1 - Y_2)\right)d\mu_1, \int_0^\infty \frac{\sqrt{n}}{c}\phi\left(\sqrt{n}(\mu_2 - Y_2)\right)\Phi\left(\sqrt{n}(\mu_2 - Y_1)\right)d\mu_2\right],$$

as shown in the online appendix, where ϕ denotes the standard normal pdf, Φ the standard normal cdf, and c is a normalizing constant defined in the online appendix. If $Y_2 - Y_1$ is small relative to $1/n$, $P(\theta \geq 0|Y)$ will be insensitive to the prior, making the model informative about $P(\theta \geq 0|Y)$. If it is large, the support of $P(\theta \geq 0|Y)$ widens, and we do not learn much from the data. These points are illustrated in Table 2 for the choices of Y_1 and Y_2 considered in Figure 1. Let $n \in \{10, 100, \infty\}$. The table shows that, when the identified set is narrow as in Figure 1(b), for $n = 10$, $P(\theta \geq 0|Y)$ can be bounded by $[0.23, 0.77]$. For $n = 100$, the bounds widen to $[0.19, 0.81]$ and for $n = \infty$ the bounds reach $[0,1]$. When the identified set is wide, in contrast, as in Figure 1(a), there is essentially no information about $P(\theta \geq 0|Y)$ for any n. Obviously, these numerical results are only suggestive. How narrow the identified set is, for given n, depends on the application. This analysis suggests that exploring the width of the support of $P(\theta \geq 0|Y)$ provides another metric for assessing how robust the posterior probabilities are to relaxing the assumption of the Haar prior in sign-identified VAR models. We will return to this point in Section 4.

3 What determines the width of the identified set?

The stylized example in Section 2 suggests that the empirical relevance of the critique of the Gaussian-inverse Wishart-Haar prior for sign-identified VAR models depends on the strength of the identifying restrictions. What determines how tightly identified the impulse responses are in practice? Even when there is no inequality restriction directly imposed on the structural impulse response of interest, sign restrictions imposed on other structural impulse responses will tend to reduce the width of the identified set and will affect the shape of the posterior over the identified set. In models with many economically motivated static and/or dynamic sign restrictions, one would expect the identified set to be smaller than in deliberately agnostic models involving only a handful of sign restrictions. This is not merely a question of the number of these restrictions, of course, but of how binding these inequality constraints jointly are.

How tight the identification is depends not only on the sign restrictions imposed on the impulse responses, however, but may also be affected by the use of other identifying restrictions. For example, some studies further restrict the values of selected impulse responses. Notably, Arias, Rubio-Ramírez and Waggoner (2018) consider models that combine traditional short-run and long-run exclusion restrictions with sign restrictions. Kilian and Murphy (2012, 2014) discuss the imposition of bounds on price elasticities of demand and supply, expressed as nonlinear inequality restrictions on impulse responses, based on extraneous evidence. Amir-Ahmadi and Drautzburg (2021) restrict the shape of selected response functions by means of slope restrictions. Another class of restrictions sometimes used involves imposing inequality restrictions on functions of the model parameters and the data. For example, Antolin-Diaz and Rubio-Ramírez (2018) discuss the imposition of narrative restrictions on the signs and relative magnitudes of structural shocks in selected periods or on the cumulative effect of structural shocks over selected subperiods, while Doh and Smith (2022) shrink the VAR model-based expectations toward extraneous measures of expectations, resulting in more informative posterior inference. Because such restrictions are imposed on functions of the data, they restrict the likelihood.

In addition, how much a given set of sign restrictions constrains the identified set also depends on the covariance structure of the reduced-form errors (see Uhlig 2017). To illustrate this point, consider the bivariate VAR(0) model of demand and supply shocks analyzed in Uhlig (2017). Price (P_t) and quantity (Q_t) are functions of mutually uncorrelated supply and demand shocks, w_t^S and w_t^D with mean zero.

$$\begin{bmatrix} P_t \\ Q_t \end{bmatrix} = B_0^{-1} \begin{bmatrix} w_t^S \\ w_t^D \end{bmatrix},$$

where the structural impact multiplier matrix is

$$B_0^{-1} = \begin{bmatrix} \cos(\nu) & \sin(\nu) \\ \sin(\varphi+\nu) & -\cos(\varphi+\nu) \end{bmatrix},$$

$\varphi \in \left(-\frac{\pi}{2}, \frac{\pi}{2}\right)$ controls the correlation between P_t and Q_t, and $\nu \in [0, 2\pi]$. Supply and demand shocks are identified by restrictions on the structural impact multiplier matrix: $b_0^{11} \leq 0$, $b_0^{21} \geq 0$, $b_0^{12} \geq 0$, $b_0^{22} \geq 0$, where b_0^{ij} is the (i,j)th element of B_0^{-1}.

The first column of Table 3 reports the identified set for the quantity response to a demand shock as a function of the correlation of the reduced-form VAR shocks, given by $\rho = \sin(\varphi)$. For expository purposes, we restrict attention to negative correlations and set the variances of P_t and Q_t to 1 such that the diagonal elements of the reduced-form error covariance matrix Σ are 1. Whereas for $\rho = 0$, the width of the identified set is 1, lowering ρ to -0.5 reduces that width to 0.87 and further lowering ρ to -0.999 reduces the width to 0.045. Likewise, raising ρ from zero toward 1 reduces the width of the identified set. Thus, changes in the reduced-form error correlation may have dramatic effects on the width of the identified set, for a given set of sign restrictions. Because Σ depends on the data, this evidence implies that the width and location of the identified set are data-dependent. These results are not specific to this example. Qualitatively similar results can be obtained with other models. Moreover, it can be shown that substantial reductions in width do not require correlations approaching 1 or -1 when considering larger-dimensional VAR models, and that not all pairwise correlations need to differ from zero.

The remaining columns of Table 3 illustrate how the reduced-form error correlation affects the posterior median and 90% credible interval for θ, given the same sign restrictions.[8] We consider three priors for ν: $\nu \sim U(0, 2\pi)$, $\nu \sim 2\pi\ Beta(2.5, 5)$, and $\nu \sim 2\pi\ Beta(5, 2.5)$ and compute the asymptotic posterior. For example, under the last prior, the posterior median of the quantity response to the demand shock ranges from 0.022 to 0.752, depending on ρ. The corresponding 90% credible interval ranges from $[0.002, 0.043]$ to $[0.523, 0.991]$, indicating substantial sensitivity to ρ. Qualitatively similar patterns hold for the other two priors.

As shown in the online appendix, this example can be generalized to show that the width of the identified set more generally depends not only on the reduced-form error covariance matrix Σ, but also on the reduced-form slope parameters in A. An immediate implication of this point and of the dependence of the identified set on the identifying restrictions is that the relevance of the critique of the Gaussian-inverse Wishart-Haar prior must be assessed on a case-by-case basis. This raises the question of how to assess the sensitivity of posterior inference to the Haar prior in practice.

4 How to assess the sensitivity of the posterior to the Haar prior

The stylized examples in Sections 2 and 3 suggest that a useful diagnostic in sign-identified Bayesian VAR models is the width of the identified set. If the identified set is narrow, the prior over the set-identified impulse response is unlikely to have a substantive impact on the posterior. If the identified set is wide, the prior will matter. While the identified set is not observable, it can be consistently estimated under assumptions covering a wide range of sign-identified models without taking a stand on the prior for Q (see Giacomini and Kitagawa 2021, Giacomini, Kitagawa and Read 2023). This set is also referred to as the set of posterior means. We follow the algorithm

[8] The number of draws is set to 10 million.

described in these studies with the specifics adapted to each VAR model.[9]

The question in practice is how narrow the width of this identified set has to be for the conventional Bayes estimate based on the Gaussian-inverse Wishart-Haar prior to produce a reasonably good approximation. One way of establishing the adequacy of the impulse response estimate is to show that the maximum error a user of this estimate commits relative to the bounds of the estimated identified set is small from an economic point of view. The maximum error corresponds to the Hausdorff distance introduced in Section 2. Whether the conventional approach is justified in a given application is not a question of how numerically small this distance is, but of how economically meaningful this distance is. In practice, it usually is reasonably clear whether the differences are large enough to matter from an economic point of view, as illustrated in Section 5. Compared to alternative metrics, one advantage of the Hausdorff distance is that it is measured in the same units as the impulse response, facilitating the interpretation of the results. The other advantage is that it accounts for both the width and location of the identified set in relation to the Bayes estimate.

Likewise, for the bounds of the conventional credible interval to be insensitive to the Haar prior, their limits should be economically indistinguishable from the bounds of the Q prior-robust credible interval proposed by Giacomini and Kitagawa (2021) and Giacomini, Kitagawa and Read (2023). In this case the Hausdorff distance may be computed as the larger of (a) the distance between the upper endpoint of the conventional credible interval and the upper endpoint of the robust credible interval, and (b) the corresponding difference between the respective lower interval endpoints.

An alternative approach to evaluating the information content of the Haar prior is to evaluate the bounds of the probability of a response being negative (or, alternatively, positive). The construction

[9] The basic idea is to generate many sets of impulse response posterior estimates that satisfy the identifying restrictions, with each draw based on the Haar prior, to compute the maximum and minimum impulse response in each set, and to average the maxima and minima across the sets. Since the support of any possible Q prior is the same as the support of the Haar prior, this approach allows the estimation of the identified set without taking a stand on what the appropriate Q prior is.

of these probability bounds requires a different approach than the construction of the identified sets for impulse responses. In the empirical section, we follow the approach in Theorem 1 of Giacomini and Kitagawa (2021) for constructing these bounds. By analogy to the earlier discussion, one could evaluate the Hausdorff distance between the posterior probability obtained under the Haar prior and the identified set for this probability. We do not pursue this question because it is less clear how to judge the economic importance of these distances than in the case of the identified set for the impulse response.

5 Empirical illustrations

In this section, we discuss three models of sign-identified VAR models drawn from the recent literature, in which posterior inference about the impulse responses depends primarily on the data and not on the prior. We contrast these examples with two alternative models in which the choice of the Haar prior substantially influences the posterior, and we show that in each case the differences are driven by the identifying assumptions.

5.1 Example 1

One of the leading examples in the literature in support of the view that the impulse response posterior is largely driven by the Haar prior is an example provided in Baumeister and Hamilton (2018, Figure 1, panels B and D). This study considered a stylized quarterly VAR model of U.S. monetary policy including the output gap, inflation and the federal funds rate. The analysis focused on the impact response of the output gap to an unexpected 25 bp tightening of monetary policy. No sign restrictions are imposed. Rather than comparing the posterior distribution of this response to its distribution conditional on the MLE of the reduced-form parameters, as in the original study, here we focus on the estimate of the identified set, following Giacomini, Kitagawa and Read (2023), and the conventional summary statistics of the posterior distribution. The first row of Table 4 shows that the identified set for this response is quite wide with 0.964 percent. The posterior mean of -0.005 percent is far from the bounds of the identified set, with a maximum error of 0.487 percent. This situation mirrors the example in Figure 1(a) shown earlier and argues against applying the conventional approach. This is also, for all practical purposes, the conclusion reached by different means in the original study.

This does not establish that the conventional approach to estimating sign-identified models must be abandoned more generally, however. While there is no question that conventional posterior inference would be questionable under this model specification, it is fair to say that no applied user would have estimated such a model without incorporating sign restrictions into the impulse response prior. In fact, it is not possible to identify the effects of monetary policy based on a model without identifying restrictions. Next, we therefore analyze the same model, but impose impact sign restrictions on the impact responses, as used in Peersman (2005) or Ouliaris and Pagan (2016) for a variation of this model:

$$\begin{pmatrix} u_t^{gap} \\ u_t^{\pi} \\ u_t^{i} \end{pmatrix} = \begin{bmatrix} + & + & - \\ - & + & - \\ - & + & + \end{bmatrix} \begin{pmatrix} w_t^{\text{cost}} \\ w_t^{\text{demand}} \\ w_t^{\text{monetary policy}} \end{pmatrix}. \tag{1}$$

We also impose narrative restrictions on the sign of the monetary policy shock in 1988.IV, 1994.I, 1990.IV, 1998.IV, 2001.II, and 2002.IV, motivated by the analysis in Antolin-Diaz and Rubio-Ramírez (2018). A natural question is how credible these identifying restrictions are (see Uhlig 2005). While our restrictions are likely to be less controversial than commonly used exclusion restrictions on the effects of monetary policy, we do not take a stand on whether these specific restrictions are necessarily correct. Rather our objective is to illustrate what difference the use of such a rich set of inequality restrictions makes for the support of the identified set of the structural response of interest and for posterior inference.[10]

The second row of Table 4 shows that merely adding the impact sign restrictions reduces the width of the identified set by more than half and reduces the maximum error of the posterior mean from 0.487 to 0.283 percent. While this is a substantial improvement, it takes the additional narrative sign restrictions to make conventional inference trustworthy. Obviously, the model remains set identified even with all additional restrictions imposed. The third row of Table 4 shows that when all restrictions are imposed, the width of the identified set shrinks to 0.089 (about 9% of the original width) and the maximum error of the posterior mean drops to 0.059 percent. Such an error is too small to make a difference for policy analysis and hence, for all practical purposes, can be ignored. The output gap in this model fluctuates between +3.5 and −3.5 percent. Thus, whether the output gap drops by 0.03 percent or 0.09 percent on impact is largely immaterial from a policy point of view. This situation mirrors the example in Figure 1(b) shown earlier. It illustrates that the critique of the conventional approach becomes largely irrelevant when the identified set is sufficiently narrow. Finally, as shown in Table 5, the corresponding maximum error of the 68% credible interval for the impact response of the output gap drops from 0.202 without any restrictions to 0.076 with all restrictions imposed. The latter Hausdorff distance is as economically negligible as the maximum error in the posterior mean.

[10] It may seem at first sight that imposing sign restrictions would not make a material difference in this model because imposing a sign restriction on the response of the output gap to a monetary policy tightening would simply truncate the posterior distribution at zero without affecting the shape of the left tail of the distribution. However, in practice, sign restrictions on the other impact responses further constrain the identified set of the response of the output gap, as do narrative restrictions, rendering the identified set much tighter than the sign restriction on b_0^{13} alone would.

5.2 Example 2

We now turn to larger-dimensional and more realistic models of monetary policy shocks. One of the classical examples of sign-identified VAR models is the model of U.S. monetary policy shocks proposed in Uhlig (2005), which postulates that a contractionary monetary policy shock increases the federal funds rate and reduces the GDP deflator, the commodity price index and non-borrowed reserves for the first six months. The estimation period is January 1965 through November 2007 to facilitate comparisons with the model in Antolin-Diaz and Rubio-Ramírez (2018) below. It is well known that the output responses in this model can be estimated only very imprecisely because the identifying restrictions are only weak. As a result, one would expect posterior inference based on the Gaussian-inverse Wishart-Haar prior to be questionable in this model. Table 6 formalizes this point. We focus on the response of real GDP to a 25 bp monetary policy tightening at horizons of 12, 24, 36, 48 and 60 months. The width of the identified sets ranges from 0.597 to 0.660 percent, indicating substantial identification uncertainty, similar to the situation in Figure 1(a). The maximum error committed by relying on the posterior mean is between 0.338 and 0.363 percent, depending on the horizon. Just how weakly identified this model is, is also illustrated by the left panel in Table 7 which shows that this model puts essentially no restrictions on the posterior probability of a recessionary response. It allows this probability to range from 0 to 100 percent at the one-year and two-year horizon. The probability range is not much narrower at longer horizons. Thus, this model is an example validating concerns in the literature over the use of the Haar prior raised in the literature.

More recently, Antolin-Diaz and Rubio-Ramírez (2018) proposed augmenting the Uhlig (2005) model with narrative restrictions. In our analysis, we employ their Narrative Sign Restriction 6 and 7 that restrict the sign of the policy shock and its relative contribution to unexpected movements in the federal fund rate in selected months. This tightens the identification considerably, resulting in a situation similar to Figure 1(b). Table 8 shows that the width of the identified set drops to between 0.210 and 0.263 percent. The maximum error implied by the posterior mean drops to between 0.106 and 0.156, making conventional posterior inference fairly accurate relative to the identified set, albeit not as accurate as in Example 1. A maximum error of about 0.1 percent in the impulse response at horizons up to three years, while not negligible, is on the small side relative to the historical variation in U.S. real GDP growth. This evidence suggests that the conclusions in Antolin-Diaz and Rubio-Ramírez (2018) are not an artifact of the use of the Haar prior and that conventional inference paints a fairly accurate picture of the responses to a monetary policy shock. Table A.2 in the online appendix shows similarly small maximum errors for the corresponding 68% credible sets. The fact that the narrative restrictions render the model much more tightly identified is also reflected in much stronger evidence of a recessionary effect in the right panel of Table 7. For example, at a horizon of two years, the probability of a recessionary response can be bounded between 69 and 100 percent compared to between 0 and 100 percent in the Uhlig (2005) model. This means that, independently of the Q prior, the probability of a recessionary effect is higher than 50 percent in the Antolin-Diaz and Rubio-Ramírez (2018) model.[11]

5.3 Example 3

Examples 1 and 2 relied on a combination of sign and narrative restrictions. The next example illustrates that identified sets may be tight even in the absence of narrative restrictions. We examine the structural VAR model of the impact of nominal gasoline price shocks on monthly consumer price inflation and one-year household inflation expectations in Kilian and Zhou (2022). The estimation period is July 1987 through April 2020. The model is identified by seven sign and two exclusion restrictions on the structural impact multiplier matrix. We focus on horizons of 0 to 5 months.

[11] A 50 percent probabilty of a recession is the threshold typically used for declaring a recession in regime-switching models (e.g., Chauvet and Hamilton 2006).

Table 9 illustrates that the responses to a nominal gasoline price shock are very tightly identified. For example, the range of the identified set for the impact response of the real price of gasoline is $[3.874, 4.396]$ percent, making the 0.522 width of the identified set small compared to the magnitude of the response estimate. Likewise the maximum error from relying on the posterior mean is only 0.357 which is negligible compared to the posterior mean response of 4.231 percent. Given a gasoline price of $3 per gallon, this translates to the difference between predicting an increase in the price of gasoline to $3.13 or to $3.12 per gallon in response to the shock, which is clearly economically immaterial. Similar results hold at the other horizons.

Similarly, the responses of headline inflation and inflation expectations are tightly identified. For example, the identified set of the impact response of headline inflation is $[0.174, 0.213]$ percent. Reporting the posterior mean of 0.197 implies a maximum error of only 0.023 percent in the inflation rate response. For the impact response of inflation expectations the identified set is $[0.160, 0.182]$ percent, implying a maximum error of 0.014 when reporting the posterior mean of 0.174. Again, similar results hold at longer horizons. These Hausdorff distances are negligible from a policy point of view. Table 9 demonstrates that the central conclusion in Kilian and Zhou (2022) that a positive gasoline price shock raises the real price of gasoline, headline inflation and inflation expectations is not an artifact of the prior and that the posterior mean implied by the conventional approach based on the Haar prior mainly reflects the information in the data. This conclusion is also consistent with the economically insignificant maximum errors for the credible intervals reported in Table A3 in the online appendix.

6 Concluding remarks

Several recent studies have voiced concerns that the Haar prior for the rotation matrix typically used in estimating sign-identified VAR models may contaminate posterior inference about impulse responses in ways applied researchers are not aware of. Our analysis in this paper calls into question the view that much of the posterior uncertainty about the impulse response in applied work must necessarily be attributed to the Haar prior. We showed that the importance of the Haar prior for impulse response inference depends on the data through the reduced-form VAR parameters and depends on how informative the identifying restrictions are, making it impossible to derive general results about the role of the Haar prior. As a result, applied users need to examine on a case-by-case basis whether their substantive conclusions are an artifact of the prior. We developed diagnostic tools that may be used for this purpose.

Using these tools, we demonstrated that in models with sufficiently informative sign (and possibly other identifying) restrictions the identified set of the impulse responses may be narrow enough for posterior inference based on a Gaussian-inverse Wishart-Haar prior to be insensitive to the Haar prior. We showed by example that this situation is not just a theoretical curiosity. Not only are there economic applications that lend themselves to richer sets of economically plausible sign restrictions than typical of early applications of sign restrictions such as Uhlig (2005), but there has been much effort in recent years to bring additional identifying information (such as narrative restrictions, shape restrictions, elasticity bounds, or selective exclusion restrictions) to bear that complements conventional sign restrictions and narrows the identified set. We presented several empirical examples drawn from the recent literature that suggest that the posterior distribution of the impulse responses sometimes depends primarily on the data rather than the Haar prior. We presented other empirical examples that support the concerns about the use of the Haar prior.

References

1. Amir-Ahmadi, P., Drautzburg, T., 2021. Identification and inference with ranking restrictions. Quantitative Economics, 12, 1-39.

2. Antolin-Diaz, J., Rubio-Ramírez, J.F., 2018. Narrative sign restrictions for SVARs. American Economic Review 108, 2802-2839.

3. Arias, J., Rubio-Ramírez, J.F., Waggoner, D.F., 2018. Inference based on SVARs identified with sign and zero restrictions: Theory and applications. Econometrica 86, 685-720.

4. Baumeister, C., Hamilton, J.D., 2015. Sign restrictions, vector autoregressions, and useful prior information. Econometrica 83, 1963-1999.

5. Baumeister, C., Hamilton, J.D., 2018. Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations. Journal of Monetary Economics 100, 48-65.

6. Canova, F., De Nicolo, G., 2002. Monetary disturbances matter for business fluctuations in the G-7. Journal of Monetary Economics 49, 1121-1159.

7. Chauvet, M., Hamilton, J.D., 2006. Dating business cycle turning points in real time. In Milas, C., Rothman, P., Van Dijk, D. (eds.), Nonlinear time series analysis of business cycles. Contributions to Economic Analysis 276, 1-54.

8. Doh, T., Smith, A.L. 2022. A new approach to integrating expectations into VAR models. Journal of Monetary Economics 132, 24-43.

9. Faust, J., 1998. The robustness of identified VAR conclusions about money. Carnegie-Rochester Conference Series on Public Policy 48, 207-244.

10. Giacomini, R., Kitagawa, T. 2021. Robust Bayesian inference for set-identified models. Econometrica 89, 1519-1556.

11. Giacomini, R., Kitagawa, T., Read, M., 2023. Identification and inference under narrative restrictions. Manuscript, University College London.

12. Kilian, L., 2013. Structural vector autoregressions, in: N. Hashimzade and M.A. Thornton (eds.), Handbook of Research Methods and Applications in Empirical Macroeconomics, Cheltenham, UK: Edward Elgar, 2013, 515-554.

13. Kilian, L., Lütkepohl, H., 2017. Structural vector autoregressive analysis. Cambridge University Press, New York.

14. Kilian, L., Murphy, D.P. 2012. Why agnostic sign restrictions are not enough: Understanding the dynamics of oil market VAR models. Journal of the European Economic Association, 10, 1166-1188.

15. Kilian, L., Murphy, D.P. 2014. The role of inventories and speculative trading in the global market for crude oil. Journal of Applied Econometrics 29, 454-478.

16. Kilian, L., Zhou, X. 2022. Oil prices, gasoline prices and inflation expectations. Journal of Applied Econometrics, 37, 867-881.

17. Moon, H.R., Schorfheide, F., 2012. Bayesian and frequentist inference in partially identified models. Econometrica 80, 755-782.

18. Ouliaris, S., Pagan, A., 2016. A method for working with sign restrictions in structural equation modelling. Oxford Bulletin of Economics and Statistics, 78, 605-622.

19. Peersman, G., 2005. What caused the early millennium slowdown? Evidence based on vector autoregressions. Journal of Applied Econometrics 20, 185-207.

20. Plagborg-Møller, M., 2019. Bayesian inference on structural impulse response functions. Quantitative Economics 10, 145-184.

21. Rubio-Ramírez, J.F., Waggoner, D., Zha, Z., 2010. Structural vector autoregressions: Theory of identification and algorithms for inference. Review of Economic Studies 77, 665-696.

22. Uhlig, H., 1998. "The robustness of identified VAR conclusions about money: A comment," Carnegie Rochester Conference Series 49, 245-263.

23. 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.

24. Uhlig, H., 2017. Shocks, sign restrictions, and identification. Advances in Economics and Econometrics: Eleventh World Congress, edited by Honoré, B., Pakes, A., Piazzesi, M., Samuelson, L., Cambridge University Press, pp. 95-127.

25. Watson, M.W., 2020. Comment on 'On the empirical (ir)relevance of the zero lower bound' by D. Debortoli, J. Gali, and L. Gambetti. NBER Macroeconomics Annual 2019 34, 182-193.

Table 1: Posterior Median of θ in Watson (2020) Model

	(a) $Y_1 = -2.5$ and $Y_2 = 2.5$				(b) $Y_1 = -0.05$ and $Y_2 = 0.05$			
	$n = 100$		$n = \infty$		$n = 100$		$n = \infty$	
	Posterior mean	Maximum error	Posterior mean	Maximum error	Posterior Mean	Maximum error	Posterior mean	Maximum error
Prior 1	-0.00	2.50	0.00	2.50	0.00	0.08	0.00	0.05
Prior 2	0.00	2.50	0.00	2.50	-0.00	0.08	-0.00	0.05
Prior 3	-2.45	4.95	-2.45	4.95	-0.08	0.16	-0.05	0.10
Prior 4	2.45	4.95	2.45	4.95	0.08	0.16	0.05	0.10

NOTES: The Watson (2020) model is a stylized model with a generic set-identified parameter. The structure of the model and priors is discussed in the main text. The maximum error is the Hausdorff distance between the posterior mean and the posterior estimate of the identified set shown in Table A1.

Table 2: Identified Set of $P(\theta \geq 0 \mid Y)$ in Watson (2020) Model

	Identified Set of $P(\theta \geq 0 \mid Y)$					
Y_1, Y_2	$n = 10$		$n = 100$		$n = \infty$	
±2.5	0.00	1.00	0.00	1.00	0.00	1.00
±0.05	0.23	0.77	0.19	0.81	0.00	1.00

NOTES: Computed based on the expression for the interval endpoints derived in the online appendix.

Table 3: Posterior Response of Quantity to Demand Shock in Uhlig (2017) Model for $n = \infty$

			$U(0, 2\pi)$			$2\pi Beta(2.5, 5)$			$2\pi Beta(5, 2.5)$		
ρ	Identified set		Median	90% credible interval		Median	90% credible interval		Median	90% credible interval	
-0.999	0.000	0.045	0.022	0.002	0.042	0.022	0.002	0.042	0.022	0.002	0.043
-0.990	0.000	0.141	0.071	0.007	0.134	0.069	0.007	0.134	0.073	0.007	0.134
-0.900	0.000	0.436	0.224	0.023	0.416	0.207	0.020	0.412	0.246	0.028	0.419
-0.500	0.000	0.866	0.500	0.052	0.839	0.440	0.043	0.828	0.617	0.099	0.849
0.000	0.000	1.000	0.866	0.545	0.999	0.896	0.559	0.999	0.752	0.523	0.991

NOTES: Bivariate model of prices and quantities with sign restrictions on the response to demand and supply shocks. Because $n = \infty$, the posterior over the identified set equals the prior.

Table 4: Posterior Mean of Impact Response of Output Gap to Monetary Policy Tightening

	Estimate of identified set			Conventional Bayes estimate	
Model	Lower bound	Upper bound	Width	Posterior Mean	Maximum error
No sign restrictions	-0.482	0.482	0.964	-0.005	0.487
Impact sign restrictions	-0.394	0.000	0.394	-0.111	0.283
Impact + shock sign restrictions	-0.089	0.000	0.089	-0.030	0.059

NOTES: Reduced-form VAR model as in Baumeister and Hamilton (2018). Identified set computed following Giacomini, Kitagawa and Read (2023). The shock is a 25 bp increase in the policy rate.

Table 5: 68% Credible Interval for Impact Response of Output Gap to Monetary Policy Tightening

Model	Robust Credible Interval			Conventional Credible Interval			
	Lower bound	Upper bound	Width	Lower bound	Upper bound	Width	Maximum Error
No sign restrictions	-0.500	0.500	1.000	-0.298	0.346	0.643	0.202
Impact sign restrictions	-0.409	0.000	0.409	-0.394	0.000	0.394	0.015
Impact+shock sign restrictions	-0.109	0.000	0.109	-0.033	0.000	0.033	0.076

NOTES: See Table 4.

Table 6: Posterior Mean of GDP Response to Monetary Policy Tightening in Uhlig (2005) Model

Horizon (months)	Estimate of identified set			Conventional Bayes estimate	
	Lower bound	Upper bound	Width	Posterior Mean	Maximum error
12	-0.232	0.365	0.597	0.116	0.347
24	-0.304	0.315	0.619	0.039	0.343
36	-0.323	0.337	0.660	0.041	0.363
48	-0.315	0.343	0.657	0.043	0.357
60	-0.304	0.328	0.631	0.035	0.338

NOTES: The static and dynamic sign restrictions are described in Uhlig (2005). The shock is a 25 bp increase in the policy rate. The responses are in percent.

Table 7: Range of Posterior Probabilities of a Recessionary Effect (Percent)

Horizon (months)	Uhlig (2005) Model		Antolin-Diaz and Rubio-Ramirez (2018) Model	
	Min	Max	Min	Max
12	0.0	100.0	1.0	97.9
24	0.0	100.0	68.7	100.0
36	0.7	100.0	61.8	99.9
48	2.7	99.8	58.5	99.7
60	6.0	99.4	56.7	98.9

NOTES: The static and dynamic sign restrictions are described in Uhlig (2005). The Antolin-Diaz and Rubio-Ramírez (2018) model includes additional narrative restrictions on the shock sign and on the relative importance of policy shocks for explaining the federal funds rate in selected periods as in Table 8.

Table 8: Posterior Mean GDP Response to Monetary Policy Tightening in Uhlig (2005) Model with Added Narrative Restrictions 6 and 7 as in Antolin-Diaz and Rubio-Ramírez (2018)

Horizon (months)	Estimate of identified set			Conventional Bayes estimate	
	Lower bound	Upper bound	Width	Posterior Mean	Maximum error
12	-0.128	-0.083	0.210	-0.021	0.106
24	-0.309	-0.071	0.239	-0.203	0.133
36	-0.332	-0.060	0.272	-0.214	0.154
48	-0.317	-0.042	0.275	-0.197	0.156
60	-0.298	-0.035	0.263	-0.184	0.149

NOTES: The two sets of additional narrative restrictions include shock sign restrictions and restrictions on the relative importance of policy shocks for explaining the federal funds rate in selected period, as described in Antolin-Diaz and Rubio-Ramírez (2018). The responses are in percent.

Table 9: Responses to Nominal Gasoline Price Shock in Kilian and Zhou (2022) Model

		Estimate of identified set			Conventional Bayes estimate	
	Horizon (months)	Lower bound	Upper bound	Width	Posterior Mean	Maximum error
Real gasoline price	0	3.874	4.396	0.522	4.231	0.357
	1	5.796	6.477	0.681	6.276	0.479
	2	5.866	6.493	0.627	6.312	0.446
	3	5.453	5.885	0.432	5.770	0.317
	4	5.218	5.552	0.334	5.466	0.248
	5	4.732	5.160	0.428	5.038	0.306
Headline inflation	0	0.174	0.213	0.038	0.197	0.023
	1	0.092	0.105	0.013	0.100	0.008
	2	0.006	0.014	0.009	0.010	0.004
	3	-0.012	-0.004	0.008	-0.008	0.004
	4	0.004	0.008	0.005	0.006	0.002
	5	-0.006	0.011	0.017	0.003	0.009
Expected inflation	0	0.160	0.182	0.022	0.174	0.014
	1	0.226	0.246	0.020	0.240	0.014
	2	0.217	0.239	0.022	0.232	0.014
	3	0.163	0.197	0.035	0.183	0.021
	4	0.152	0.195	0.043	0.177	0.025
	5	0.115	0.138	0.023	0.129	0.013

NOTES: The model is identified based on sign and exclusion restrictions on the structural impact multiplier matrix, as described in Kilian and Zhou (2022). The responses are in percent.

Figure 1: Posterior Distribution of θ in the Watson (2020) Model under Alternative Priors Given Wide and Narrow Identified Sets

(a) $Y_1 = -2.5$ and $Y_2 = 2.5$

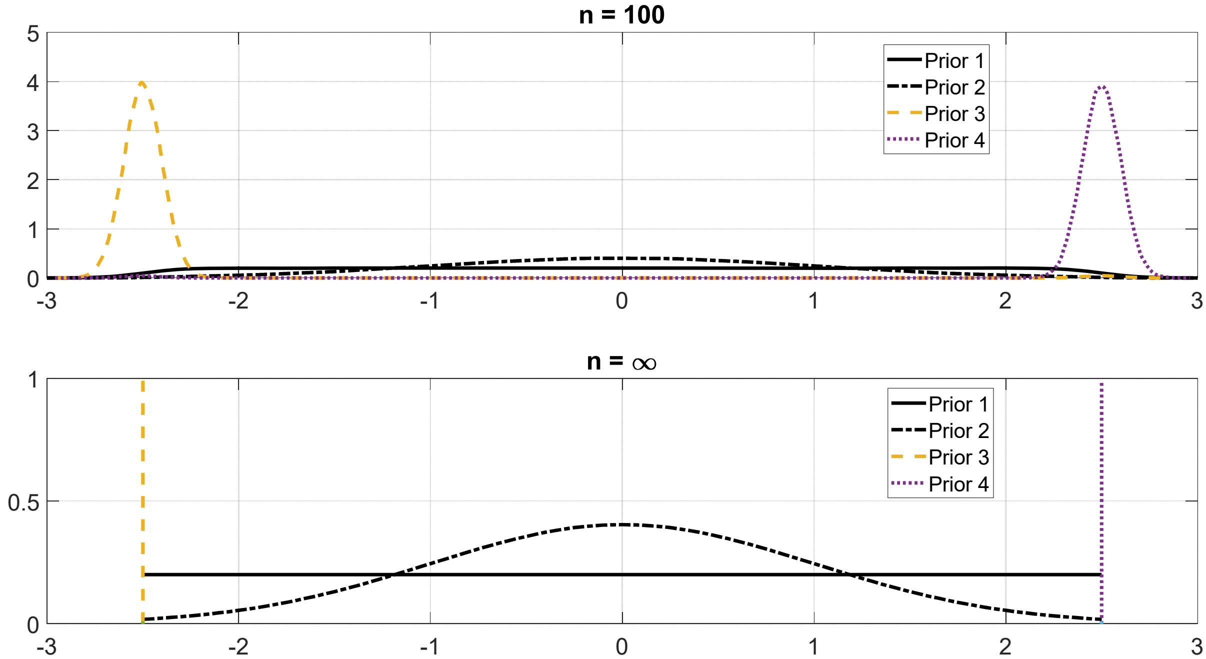

(b) $Y_1 = -0.05$ and $Y_2 = 0.05$

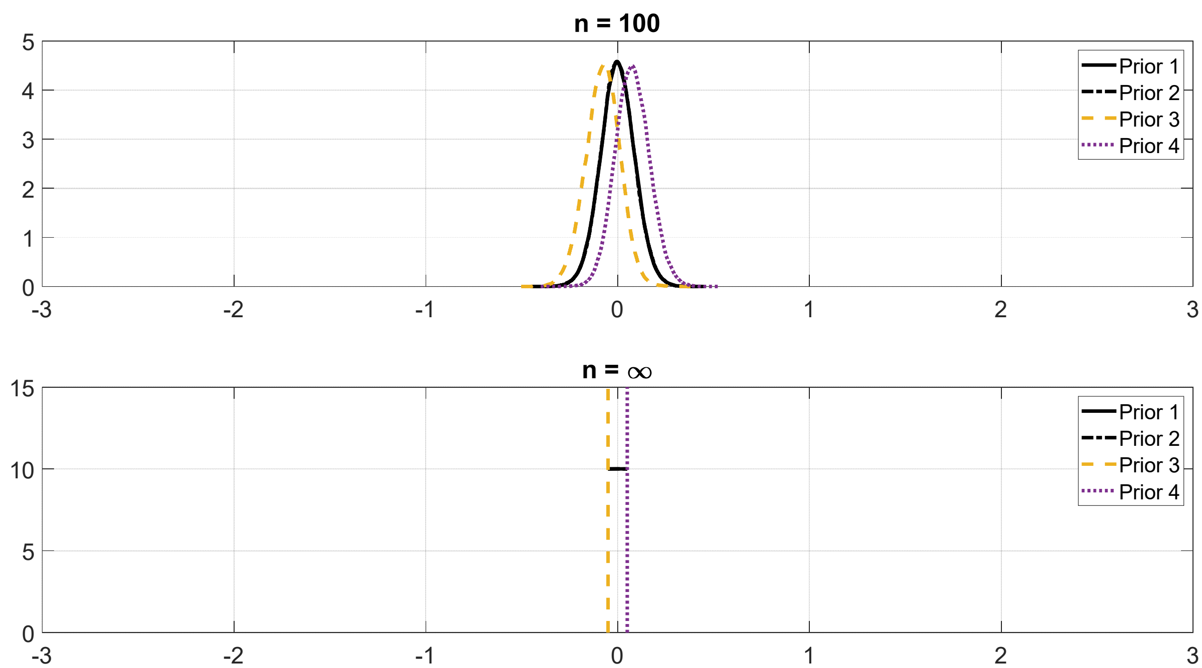

NOTES: For $n = 100$, we estimate the posterior densities using a kernel density smoother. The densities for $n = \infty$ are derived analytically. The posterior under priors 3 and 4 assigns 0.99 probability mass to one end point of the identified set and 0.01 to the other. For simplicity, we represent this result as a vertical line at the endpoint with 0.99 probability mass.

When Is the Use of Gaussian-inverse Wishart-Haar Priors Appropriate? Online Appendix

Atsushi Inoue*
Vanderbilt University

Lutz Kilian†
Federal Reserve Bank of Dallas
CEPR

July 7, 2024

*Vanderbilt University, Department of Economics, Nashville, TN 37235-1819. E-mail: atsushi.inoue@vanderbilt.edu.

†Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA. E-mail: lkilian2019@gmail.com (corresponding author).

1 Notation

- Let y_t be a $K \times 1$ vector of variables generated by the structural VAR(p) model $B_0 y_t = Bx_t + w_t$, with reduced-form representation $y_t = Ax_t + u_t$, for $t = 1, ..., n$, where $x_t = (y'_{t-1}, ..., y'_{t-p})$, B_0 is invertible, the vector of structural shocks is $w_t \overset{iid}{\sim} N(0, I_K)$, the vector of reduced-form shocks is $u_t = B_0^{-1} w_t$ such that $u_t \overset{iid}{\sim} N(0, \Sigma)$, $\Sigma = B_0^{-1} B_0^{-1\prime}$, $B = (B_1, ..., B_p)$, $A = (A_1, ..., A_p)$, and deterministic regressors have been suppressed.

- Let Q be a $K \times K$ matrix such that $QQ' = Q'Q = I_K$ and let P denote the lower triangular Cholesky decomposition of Σ with the diagonal elements normalized to be positive. Then $PP' = PQ(PQ)' = \Sigma$.

- We assume that the VAR(p) model can be inverted to obtain the VMA(∞) representation, $y_t = \sum_{h=0}^{\infty} C_h u_{t-h} = \sum_{h=0}^{\infty} C_h PQw_{t-h}$, $t = 1, ..., n$, where C_h is the hth term in $(I_K - \sum_{l=1}^{p} A_l L^l)^{-1}$ and L is the lag operator. The(i, j)th element of the structural impulse response matrix $C_h PQ$ is the horizon h response of variable i to shock j. We denote a generic example of this response for some combination of (i, j, h) by $\theta = g(A, \Sigma, Q)$, emphasizing that this response has been obtained by a nonlinear transformation $g(\cdot)$ of the model parameters A, Σ, and Q.

2 Derivation of the Bounds on P($\theta \geq 0|Y$) in the Watson (2020) Model

Let Y_1 and Y_2 be independent normal random variables with means μ_1 and μ_2, respectively, and variances $\sigma^2 = \frac{1}{n}$, where $\mu_1 < \mu_2$. Given a flat prior for μ_1 and μ_2, their posterior density is given by

$$p(\mu_1, \mu_2|Y_1, Y_2) = \frac{1}{\sigma^2 c} \phi\left(\frac{\mu_1 - Y_1}{\sigma}\right) \phi\left(\frac{\mu_2 - Y_2}{\sigma}\right) I(\mu_1 < \mu_2), \tag{1}$$

where

$$c \;=\; 1-\frac{1}{\sigma}\int_{-\infty}^{\infty}\phi\left(\frac{\mu_1-Y_1}{\sigma}\right)\Phi\left(\frac{\mu_1-Y_2}{\sigma}\right)d\mu_1 \;=\; \frac{1}{\sigma}\int_{-\infty}^{\infty}\phi\left(\frac{\mu_2-Y_2}{\sigma}\right)\Phi\left(\frac{\mu_2-Y_1}{\sigma}\right)d\mu_2 \quad (2)$$

By Theorem 1 of Giacomini and Kitagawa (2021), the lower bound on the probability is given by

$$\begin{aligned} P(\mu_1 \;\geq\; 0|Y_1,\; Y_2) \;&=\; \int_0^{\infty}\int_{\mu_1}^{\infty} p(\mu_1,\; \mu_2|Y_1,\; Y_2)d\mu_2 d\mu_1 \\ &=\; \int_0^{\infty}\frac{1}{\sigma c}\phi\left(\frac{\mu_1-Y_1}{\sigma}\right)\Phi\left(-\frac{\mu_1-Y_2}{\sigma}\right)d\mu_1, \end{aligned} \quad (3)$$

when n is finite. Similarly, the upper bound is

$$\begin{aligned} P(\mu_2 \;\geq\; 0|Y_1,\; Y_2) \;&=\; \int_0^{\infty}\int_{-\infty}^{\mu_2} p(\mu_1,\; \mu_2|Y_1,\; Y_2)d\mu_1 d\mu_2 \\ &=\; \int_0^{\infty}\frac{1}{\sigma c}\phi\left(\frac{\mu_2-Y_2}{\sigma}\right)\Phi\left(\frac{\mu_2-Y_1}{\sigma}\right)d\mu_2, \end{aligned} \quad (4)$$

when n is finite. In Table 2, we report Monte Carlo simulation estimates of $P(\mu_1 \geq 0|Y_1, Y_2)$ and $P(\mu_2 \geq 0|Y_1, Y_2)$ when n is finite. When $n = \infty$, the probability bounds are given by

$$P(\theta > 0|Y_1, Y_2) \;\in\; \begin{cases} \{0\} & \text{if } Y_2 < 0 \\ [0,1] & \text{if } Y_1 \leq 0 \leq Y_2 \\ \{1\} & \text{if } Y_1 > 0 \end{cases} .$$

3 Generalization of the Uhlig (2017) Example

The Uhlig (2017) example may be generalized to show that the width of the identified set depends not only on Σ, but also on A. Consider a dynamic version of Uhlig's supply and demand model.

The identifying restrictions are unchanged, but we allow for one autoregressive lag:

$$\begin{bmatrix} P_t \\ Q_t \end{bmatrix} = \begin{bmatrix} a_{11,1} & a_{12,1} \\ a_{21,1} & a_{22,1} \end{bmatrix} \begin{bmatrix} P_{t-1} \\ Q_{t-1} \end{bmatrix} + \begin{bmatrix} \cos(\nu) & \sin(\nu) \\ \sin(\varphi+\nu) & -\cos(\varphi+\nu) \end{bmatrix} \begin{bmatrix} w_t^S \\ w_t^D \end{bmatrix}. \tag{5}$$

The one-step-ahead response matrix is

$$\begin{bmatrix} a_{11,1}\cos(\nu) + a_{12,1}\sin(\varphi+\nu) & a_{11,1}\sin(\nu) - a_{12,1}\cos(\varphi+\nu) \\ a_{21,1}\cos(\nu) + a_{22,1}\sin(\varphi+\nu) & a_{21,1}\sin(\nu) - a_{22,1}\cos(\varphi+\nu) \end{bmatrix}. \tag{6}$$

The width of the identified set of the quantity response to a one-standard-deviation demand shock, $(a_{21,1}\sin(\nu) - a_{22,1}\cos(\varphi+\nu))$, for example, then becomes:

$$\max_{\frac{\pi}{2}-\varphi\leq\nu\leq\pi}(a_{21,1}\sin(\nu) - a_{22,1}\cos(\varphi+\nu)) - \min_{\frac{\pi}{2}-\varphi\leq\nu\leq\pi}(a_{21,1}\sin(\nu) - a_{22,1}\cos(\varphi+\nu)) \quad \text{if} \quad \varphi < 0,$$
$$\max_{\frac{\pi}{2}\leq\nu\leq\pi-\varphi}(a_{21,1}\sin(\nu) - a_{22,1}\cos(\varphi+\nu)) - \min_{\frac{\pi}{2}\leq\nu\leq\pi-\varphi}(a_{21,1}\sin(\nu) - a_{22,1}\cos(\varphi+\nu)) \quad \text{if} \quad \varphi > 0.$$

It is readily apparent that this width depends not only on the reduced-form error correlations through φ, as in the static model, but also on the slope parameters in A. For example, let $a_{11,1} = a_{22,1} = 0.9$, $a_{12,1} = 0.5$, and $\varphi = \arcsin(0.5)$, as we vary the value of $a_{21,1}$. Table A.4 shows that the width of the identified set for the one-step-ahead quantity response to a supply shock changes from 0.95 to 0.246 as the value of $a_{21,1}$ changes, all else equal, from -1.0 to 1.0. The location of the identified set also changes with the value of $a_{21,1}$.

4 Additional empirical results

Tables A.1, A.2, A.3 and A.4 contain additional empirical results that were omitted to conserve space, but are discussed in the text.

Disclaimer

The views expressed in this appendix are those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

Table A1: Posterior Distribution of θ in Watson (2020) Model

(a) $Y_1 = -2.5$ and $Y_2 = 2.5$

	$n = 100$			$n = \infty$		
Posterior identified set	Mean width	Range		Mean width	Range	
	5.00	-2.50	2.50	5.00	-2.50	2.50
	Mean	90% Credible Interval		Mean	90% Credible Interval	
Posterior under prior 1	-0.00	-2.25	2.25	0.00	-2.25	2.25
Posterior under prior 2	0.00	-1.59	1.59	0.00	-1.59	1.59
Posterior under prior 3	-2.45	-2.66	-2.33	-2.45	-2.50	-2.50
Posterior under prior 4	2.45	2.33	2.66	2.45	2.50	2.50

(b) $Y_1 = -0.05$ and $Y_2 = 0.05$

	$n = 100$			$n = \infty$		
Posterior identified set	Mean width	Range		Mean width	Range	
	0.16	-0.08	0.08	0.10	-0.05	0.05
	Mean	90% Credible Interval		Mean	90% Credible Interval	
Posterior under prior 1	0.00	-0.15	0.15	0.00	-0.04	0.04
Posterior under prior 2	-0.00	-0.15	0.15	-0.00	-0.05	0.04
Posterior under prior 3	-0.08	-0.23	0.07	-0.05	-0.05	-0.05
Posterior under prior 4	0.08	-0.07	0.23	0.05	0.05	0.05

NOTES: The Watson (2020) model is a stylized model with a generic set-identified parameter. The structure of the model and priors is discussed in the main text.

Table A.2: 68% Credible Intervals for GDP Response to Monetary Policy Tightening in Uhlig (2005) Model and in Antolin-Diaz and Rubio-Ramírez (2018) Model

Uhlig (2005) Model							
	Robust Credible Interval			Conventional Credible Interval			
Horizon (months)	Lower bound	Upper bound	Width	Lower bound	Upper bound	Width	Maximum Error
12	-0.300	0.404	0.703	-0.007	0.259	0.267	0.292
24	-0.363	0.374	0.737	-0.092	0.191	0.283	0.271
36	-0.418	0.417	0.835	-0.077	0.235	0.312	0.341
48	-0.444	0.429	0.873	-0.137	0.201	0.338	0.307
60	-0.425	0.451	0.876	-0.105	0.247	0.352	0.320
Antolin-Diaz and Rubio-Ramírez (2018)							
	Robust Credible Interval			Conventional Credible Interval			
Horizon (years)	Lower bound	Upper bound	Width	Lower bound	Upper bound	Width	Maximum Error
12	-0.195	0.171	0.366	-0.098	0.063	0.161	0.108
24	-0.377	0.010	0.388	-0.288	-0.114	0.174	0.124
36	-0.435	0.034	0.469	-0.310	-0.089	0.221	0.124
48	-0.423	0.089	0.512	-0.307	-0.054	0.253	0.142
60	-0.418	0.112	0.530	-0.312	-0.036	0.276	0.148

NOTES: See Tables 6 and 8. All responses are in percent.

Table A.3: 68% Credible Intervals for Responses to Nominal Gasoline Price Shock in Kilian and Zhou (2022) Model

	Horizon (months)	Robust Credible Interval Lower bound	Upper bound	Width	Conventional Credible Interval Lower bound	Upper Bound	Width	Maximum Error
Real gasoline price	0	3.659	4.576	0.917	4.060	4.490	0.430	0.401
	1	5.477	6.859	1.382	5.971	6.715	0.745	0.494
	2	5.345	6.909	1.564	5.893	6.858	0.965	0.548
	3	4.826	6.340	1.514	5.201	6.284	1.083	0.376
	4	4.578	6.125	1.547	4.743	5.965	1.222	0.164
	5	3.966	5.733	1.767	4.382	5.731	1.350	0.416
Headline Inflation	0	0.166	0.221	0.055	0.186	0.215	0.030	0.020
	1	0.080	0.115	0.034	0.090	0.113	0.023	0.009
	2	-0.006	0.025	0.031	-0.002	0.022	0.023	0.004
	3	-0.024	0.008	0.032	-0.019	0.006	0.025	0.006
	4	-0.009	0.019	0.028	-0.008	0.017	0.024	0.003
	5	-0.017	0.023	0.041	-0.009	0.017	0.026	0.008
Expected inflation	0	0.012	0.017	0.005	0.013	0.016	0.003	0.001
	1	0.017	0.023	0.006	0.018	0.022	0.004	0.001
	2	0.016	0.022	0.006	0.017	0.022	0.005	0.001
	3	0.011	0.018	0.008	0.013	0.018	0.006	0.002
	4	0.010	0.019	0.009	0.012	0.018	0.006	0.002
	5	0.007	0.014	0.007	0.008	0.013	0.005	0.001

NOTES: See Table 9. All responses are in percent.

Table A.4: Posterior of one-step ahead response of quantity to demand shock for $n = \infty$

	$a_{21,1}$								
Identified set	-1.000	-0.750	-0.500	-0.250	0.000	0.250	0.500	0.750	1.000
Lower bound	-0.550	-0.300	-0.050	0.200	0.450	0.700	0.950	1.200	1.400
Upper bound	0.400	0.525	0.650	0.775	0.900	1.048	1.229	1.431	1.646
Width	0.950	0.825	0.700	0.575	0.450	0.348	0.279	0.231	0.246

NOTES: Based on the bivariate VAR(1) extension of the Uhlig (2017) model with the coefficient settings described in the online appendix.