The Conventional Impulse Response Prior in VAR Models with Sign Restrictions

Atsushi Inoue*
Vanderbilt University

Lutz Kilian†
Federal Reserve Bank of Dallas
CEPR

March 18, 2025

Abstract: Some studies have expressed concern that the Gaussian-inverse Wishart-Haar prior typically employed in estimating sign-identified VAR models may be unintentionally informative about the implied prior for the structural impulse responses. We discuss how this prior may be reported and make explicit what impulse response priors a number of recently published studies specified, allowing the readers to decide whether they are comfortable with this prior. We discuss what features to look for in this prior in the absence of specific prior information about the responses and in the presence of such information. Our empirical examples illustrate that the Gaussian-inverse Wishart-Haar prior need not be unintentionally informative about the impulse responses. Moreover, even when it is, there are empirically verifiable conditions under which this fact becomes immaterial for the substantive conclusions.

JEL code: C22, C32, C52, E31, Q43

Key words: Gaussian-inverse Wishart prior, Haar prior, impulse response, set identification, sign restrictions.

*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 Introduction

Structural VAR models identified based on sign restrictions on structural impulse responses play an important role in empirical research.[1] The conventional Bayesian approach to estimating such models involves specifying a conjugate Gaussian-inverse Wishart prior for the parameters A and Σ in the reduced-form VAR model, where A denotes the slope parameters and Σ is the error covariance matrix, and a Haar prior for the orthogonal rotation matrix Q. The prior for the vector of structural impulse responses $\theta = g(A, \Sigma, Q)$ is defined implicitly by the nonlinear function $g(\cdot)$.[2]

Several studies have expressed concern that this impulse response prior may not reflect the beliefs of researchers (e.g., Baumeister and Hamilton 2015, Watson 2019). If so, posterior inference about the impulse responses may be driven by features of the prior that were never intended by the user of the conventional approach. Whether the implicit prior is unintentionally informative is impossible to know without formally evaluating it, yet this question has not been commonly addressed in empirical work. In this paper, we propose evaluating this prior by simulation. We provide tools for reporting this prior and make explicit the impulse response priors in a number of published studies.[3]

Reporting the implicit impulse response prior is important because there is nothing wrong with posterior inference about the impulse responses if this prior is proper and reflects the researcher's beliefs (see, e.g., Poirier 1998; Gustafson 2009).[4] As long as users of sign-identified structural VAR

[1] Earlier prominent applications of sign-identified VAR models include Faust (1998), Canova and de Nicolo (2002), and Uhlig (2005). For a review of the literature on sign-identified structural VAR models see Kilian and Lütkepohl (2017).

[2] The Haar prior was introduced in Rubio-Ramirez et al. (2010) as a computationally efficient device to simulate draws from sign-identified structural VAR models, building on and generalizing the methodology in Uhlig (2005) (see also Arias et al. (2018) and Antolin-Diaz and Rubio-Ramirez (2018)).

[3] In related work, Arias, Rubio-Ramírez and Waggoner (2025) show that a minor modification of the conventional algorithm used to estimate sign-identified structural VAR models results in a jointly uniform prior on the vector of impulse responses. In this paper, in contrast, we examine the conventional approach as traditionally employed in applied work.

[4] Some empirical studies employ the improper reduced-form parameter prior popularized by Uhlig (2005) (see, e.g., Kilian and Murphy 2014). By construction, the implied impulse response prior is not proper in that case, which is why we do not consider that prior. Often the impulse response posterior is not very sensitive to replacing an improper reduced-form prior by an alternative proper prior.

models agree that the implicit impulse response prior reflects their beliefs, posterior inference based on conventional priors for the model parameters is justified. Some researchers may prefer to remain weakly informative a priori about the features of the impulse responses and may be content with an impulse response prior that does not take a stand on the sign of the impulse responses, but allows for approximately symmetric deviations from the zero line with lower probability mass assigned to more distant draws. Other researchers may prefer a prior that is strongly informative about the impulse responses in line with the implications of an economic model. Either way, reporting and defending that prior is essential, so authors as well as readers can decide for themselves whether they are comfortable with this prior.

In this paper, we show by example that in many models identified by static sign restrictions the implicit priors on the impact responses are only weakly informative. Moreover, the impulse response priors tend to be neutral with respect to the sign of the responses at longer horizons, provided the reduced-form prior for the slope parameters is centered on zero and the prior variance is chosen based on approximating AR(1) models for each variable. In contrast, models with dynamic sign are informative about the sign of these responses. We show by example that, even in the presence of static and dynamic sign restrictions (or for that matter narrative restrictions), the implied impulse response prior need not be at odds with economic reasoning and may correspond to a researcher's prior beliefs. The examples from applied work we consider provide no evidence of impulse response priors that are patently unreasonable, say, because their probability mass is concentrated in the tails or because their shape is arbitrary.

Of course, in general, the impulse response posterior depends on the choice of the impulse prior. Different prior views may result in different posteriors. This raises the question of how to proceed when the prior does not correspond to one's beliefs. Recent analysis in Inoue and Kilian (2025) shows that the use of the Gaussian-inverse Wishart-Haar prior remains justified even in that case,

provided the identified set is sufficiently tight. Under this condition posterior inference will be insensitive to the choice of the prior and notably to the use of a Haar prior for the rotation matrix, even asymptotically. Only when the identified set is wide and the implicit impulse response prior is at odds with one's prior views, caution is called for in using the conventional approach to estimating sign-identified VAR models.

The remainder of the paper is organized as follows. In Section 2, we stress that, in practice, the joint prior distribution and the corresponding marginal prior distributions of the individual elements of θ may be evaluated by simulation. We stress that the impulse response prior distribution depends not only on the prior for the rotation matrix Q, but also on the prior for the reduced-form parameters A and Σ. Our analysis highlights that this prior differs from the distribution of θ conditional on specific values of A and Σ, as discussed in Baumeister and Hamilton (2015) and Watson (2019). In Section 3, we discuss how to differentiate between weakly informative and strongly informative impulse response priors. In Section 4, we consider a stylized model of monetary policy shocks and demonstrate by example what marginal impulse response prior the conventional approach implies for the impact period. We then show that the prior is only weakly informative not only about the impact responses, but also about the sign of the responses at longer horizons when the prior for A is centered on zero. We also examine how this conclusion changes when centering the prior for A on an independent random walk. In Section 5, we examine a sign-identified model of the transmission of gasoline price shocks to inflation and inflation expectations and in Section 6 a sign-identified model of the global oil market. In Section 7, we extend the analysis to larger sign-identified models of monetary policy shocks employed in the literature. In Section 8, we discuss how to proceed when the impulse response prior is unacceptable to the user, building on recent insights in Inoue and Kilian (2025). The concluding remarks are in Section 9.

2 The impulse response prior

The primary object of interest in structural VAR analysis is the structural impulse response vector θ obtained by stacking the structural responses of interest, which captures the shape and co-movement of response functions that are of interest in applied work (see, e.g., Sims and Zha 1999; Inoue and Kilian 2022, Arias et al. 2025). Consider the n-dimensional structural VAR(p) model

$$B_0 y_t = B_1 y_{t-1} + ... + B_p y_{t-p} + w_t,$$

where the intercept has been suppressed for expositional purposes, B_i, $i = 0, ..., p$, are $n \times n$ coefficient matrices, and the structural error, w_t, is mean zero mutually uncorrelated Gaussian white noise. Without loss of generality, we impose the normalization that w_t has the covariance matrix I_n. The corresponding reduced-form VAR(p) representation is

$$y_t = A_1 y_{t-1} + ... + A_p y_{t-p} + u_t,$$

where $A_i = B_0^{-1} B_i$, $i = 1, ..., p$, and $u_t = B_0^{-1} w_t$. The variance-covariance matrix of the reduced-form error, u_t, is $\Sigma = B_0^{-1} B_0^{-1\prime}$. Identification is typically achieved by imposing restrictions on B_0^{-1}. In models identified based on exclusion restrictions, the vector of structural impulse responses can be written as $\theta = g(A, \Sigma)$, where $A = [A_1, ..., A_p]$ and $g(\cdot)$ denotes a nonlinear function. In the case of sign-identified models, $B_0^{-1} = PQ$, where P is the lower triangular Cholesky decomposition of Σ with positive elements on the diagonal and Q denotes an $n \times n$ matrix such that $QQ' = Q'Q = I_n$. Hence, $\theta = g(A, \Sigma, Q)$, where $g(\cdot)$ denotes a nonlinear function. For details of the specification, estimation and identification of this class of models the reader is referred to Kilian and Lütkepohl (2017).

2.1 The conventional Gaussian-inverse Wishart-Haar prior

The conventional approach to estimating sign-identified VAR models postulates an inverse Wishart prior for Σ, a Gaussian prior for A conditional on Σ, and a Haar prior for Q, which may be viewed as a uniform prior in the space of possible rotation matrices. Some applications of sign-identified structural VAR models rely on a diffuse reduced-form prior popularized by Uhlig (2005). Although such diffuse priors imply a Gaussian-inverse Wishart-Haar posterior, they are not proper priors. This makes it impossible to evaluate whether the implied impulse response prior is unintentionally informative.

For expository purposes, we therefore work with a version of the popular conjugate reduced-form Gaussian-inverse Wishart prior discussed in Karlsson (2013) (see also the appendix in Uhlig (1994)). This prior involves a pre-specified rate of decay at higher lags, sets the relative tightness of other variables in a given equation to 1, and sets the prior standard deviation of the first own lag in each equation, λ, to 0.2. The prior variance also depends on estimates of the innovation variances obtained from fitting univariate AR(1) models to each variable. As the baseline, we set the prior mean to zero, but we also report selected results for a prior centered on a random walk mean for each variable. Given the Haar prior for Q, these assumptions also imply a Gaussian-inverse Wishart-Haar prior for structural VAR models identified by sign restrictions. In many applications, replacing the diffuse reduced-form parameter prior by this alternative prior does not materially change the posterior distribution of θ.

2.2 The implied prior for θ

The conventional approach is for applied researchers to specify a joint prior for the parameters Q, A, and Σ, rather than directly specifying a prior for the vector of the structural impulse responses, $\theta = g(A, \Sigma, Q)$. In doing so, the researcher implicitly also defines the prior for θ. Since θ depends on

Σ at all horizons and, in addition, on A at all horizons but the impact period, it follows immediately from an application of the change-of-variable theorem that the implicit impulse response prior $f(\theta)$ depends both on the Haar prior and on the prior uncertainty about the reduced-form parameters:

$$f(\theta) = \int_{(A,\Sigma,Q):\ \theta=g(A,\Sigma,Q) \text{ satisfies the identifying restrictions}} \phi(A|\Sigma)\nu(\Sigma)\ d\Sigma\ d\Sigma\ dQ, \tag{1}$$

where $\nu(\cdot)$ denotes the inverse Wishart density and $\phi(\cdot)$ the conditionally Gaussian density (see also Arias et al. 2018).

One could also evaluate the density of θ conditional on the value of A and Σ:

$$h(\theta|A,\Sigma) = \frac{\int_{Q:\ \theta=g(A,\Sigma,Q) \text{ satisfies the identifying restrictions}} \phi(A|\Sigma)\nu(\Sigma)\ d\Sigma\ dQ}{\phi(A|\Sigma)\nu(\Sigma)}. \tag{2}$$

Unless the user's prior uncertainty about A and Σ is zero, however, this conditional distribution will differ from the prior distribution $f(\theta)$ embodied in the conventional approach. The existing debate about the merits of the Gaussian-inverse Wishart-Haar prior has focused on $h(\theta|A,\Sigma)$ rather than $f(\theta)$. This distinction is important because, in general, $f(\theta)$ cannot inferred from $h(\theta|A,\Sigma)$, whether we are interested in the joint or in the marginal distribution of the impulse responses. For example, the widely cited stylized example in Baumeister and Hamilton (2015) and the evidence in Watson (2019) that $h(\theta|A,\Sigma)$ may be unintentionally informative when employing the Haar prior does not speak to the question of whether the impulse response priors actually used by researchers in applied work are unintentionally informative. While this point may seem obvious upon reflection, it has not been widely recognized in the literature.

This fact that little is known about the nature of $f(\theta)$ in applied work motivates us in this paper to develop easy-to-use tools to describe the marginal impulse response priors, allowing researchers to examine the implicit impulse response prior. We argue that these methods are a natural starting

point when assessing the information content of the implicit impulse response prior and should be routinely employed by users of the Gaussian-inverse Wishart-Haar prior. Extensions to joint impulse response priors are straightforward. In practice, the central question is whether the implicit impulse response prior $f(\theta)$ is consistent with the researchers's prior beliefs about the impulse responses.

Articulating these beliefs does not require knowledge of the identified set in sign-identified models nor does it require making a distinction between values inside that set and outside the identified set. While the likelihood is flat over the identified set in sign-identified models, there is no reason for the prior of the researcher to be flat. In fact, there is no reason for researcher's prior beliefs about the impulse responses to be different in sign-identified models than in point-identified models. After all these beliefs are formed about the same object. For example, a prior belief that monetary policy shocks have little effect on impact on real output or a prior belief that such shocks have hump-shaped effects on real output do not depend on the identifying assumptions in the econometric model. Thus, one would expect researchers to have prior views about how the probability mass of the impulse responses is distributed in sign-identified models.

2.3 Approximating the prior distribution by simulation

In practice, $f(\theta)$ and the corresponding marginal prior distributions for each element of θ may be approximated by simulation, given a sufficiently large number of admissible draws from the prior distributions of (A, Σ, Q) that satisfy the identifying restrictions. In this paper, we focus on the marginal prior distributions of the impulse responses. We not only plot the marginal priors for the impact responses of interest, but also report the central tendency of the marginal impulse priors at all horizons. Specifically, we follow the convention in much of applied work of postulating an additively separable absolute loss function and report the vector of prior medians of the impulse

responses and selected percentiles of the prior distribution. A complementary approach would have been to report the central tendency of the joint prior distribution and to approximate the corresponding lowest joint prior risk regions, as proposed in Inoue and Kilian (2022). This would have allowed us to also examine the shapes and co-movement of the response functions. In the interest of space, in this paper we follow much of the earlier literature and focus on pointwise inference.

Our approach is largely analogous to the simulation of the marginal impulse response posterior and can be summarized as follows:

Algorithm (Simulating the impulse response prior implied by the Gaussian-inverse Wishart-Haar prior and the identifying restrictions)

1. Generate M draws of A and Σ from the normal-inverse Wishart prior, as described in Section 3.2.1. of Karlsson (2013).

2. For each draw of (A, Σ) generate R draws of Q from the uniform distribution over $O(n)$, as discussed in Rubio-Ramirez et al. (2010).

3. Construct $\theta = g(A, \Sigma, Q)$ for each of the $M \times R$ draws of (A, Σ, Q) and verify that the identifying restrictions hold. Retain the draw if it satisfies the identifying restrictions. Otherwise discard the draw as inadmissible.

4. Repeat steps 1-3 until a sufficiently large number of admissible θ draws has been generated. Build up the prior distribution for each element of θ of interest. Characterize this distribution based on its moments or quantiles and report the histogram of the draws.

We have found the upper and lower 68% and 90% percentile intervals along with the median to be useful in summarizing the information contained in the prior. How many draws of θ are required

depends on the identifying restrictions. Models with tighter identification require more draws of (A, Σ, Q) to generate enough admissible draws of θ for a good approximation of the prior. This algorithm can be adapted to allow for combinations of sign and zero restrictions, as in Arias et al. (2018), in which case the θ draws in step 3 must be re-weighted using the importance sampler they describe. The algorithm can also be adapted to incorporate narrative inequality restrictions as in Antolin-Diaz and Rubio-Ramirez (2018). In the latter case, the narrative restrictions are evaluated in step 3 at the same time as the sign restrictions on θ. Evaluating narrative restrictions requires knowledge of the data, but there is no need for an importance sampler to correct the likelihood as required for posterior inference.

3 Characterizing the impulse response prior

As is widely recognized in the literature, there is no conceptual difficulty in conducting Bayesian inference based on a Gaussian-inverse Wishart-Haar prior, as long as the implied impulse response prior is proper and not unintentionally informative (e.g., Poirier 1998; Gustafson 2009; Gelman et al. 2013). One can imagine several situations that satisfy this criterion. One involves a subjective prior based on the best scientific knowledge available. In some circumstances, it may be possible to defend *strongly informative* impulse response priors based on economic reasoning and extraneous evidence. In others, extraneous evidence may be unreliable or insufficient to parameterize the prior or there is no scientific consensus. In the latter situation, one may wish to work with less informative priors. Gelman et al. (2013) defines a *weakly informative* prior distribution as a prior containing some information but without attempting to fully capture one's scientific knowledge about the underlying parameter. Such priors are motivated by the observation that it is not always possible or desirable to encode all of one's prior beliefs about a subject into a set of probability distributions. Gelman et al. offer two principles for setting up weakly informative priors: Either

start with some version of a noninformative prior distribution and then add enough information so that inferences are constrained to be reasonable or start with a strong, highly informative prior and broaden it to account for uncertainty in one's prior beliefs and in the applicability of any historically based prior distribution to new data. Below we discuss how these ideas may be operationalized in the impulse response context.

3.1 Weakly informative impulse response priors

We follow Gelman et al. (2013) and characterize a prior distribution as weakly informative if it is proper but is set up so that the information it contains is intentionally weaker than whatever actual prior knowledge is available. Suppose that the researcher prefers to limit the use of extraneous information about the impulse response to a minimum. For starters, let us consider only the impact responses. In the absence of sign restrictions a natural starting point is to require that the impact response prior does not favor responses of positive sign or of negative sign. Instead, it should assign equal weight to positive and negative responses, making the prior sign-neutral. For example, under absolute loss, the impact response prior would be sign-neutral if the prior median is zero because in that case exactly 50% of the prior probability mass is in the positive region and 50% in the negative region. This postulate may be viewed as an application of the symmetry principle in Gelman et al. (2013) that the prior distribution should not pull inferences in any predetermined direction.

An additional plausible feature of the impulse response prior is that - regardless of one's other prior views about the response - values of the response far from zero are less plausible than values closer to zero. For example, most macroeconomists would not feel comfortable assigning as high a probability to an impact response of real GDP in the range between -2% and -4% than to a response in the range between 0% and -2%, given a 25 bp monetary tightening. This reasoning suggests that – in the absence of further information about the impulse response – the prior density

should be centered on zero and its value should monotonically decline with the distance from zero. We refer to such a prior as weakly informative in that it only reflects minimal extraneous information.

Clearly, such weakly informative priors are not unique. They may differ in their dispersion and in how quickly the value of the prior density declines away from zero. For example, both a Gaussian and a Student-t prior centered on zero can be viewed as a weakly informative prior about the impulse response in the sense defined above, but the latter assigns more probability mass to the tail, controlling for the variance. Thus, these priors differ in how confident the researcher is a priori that the response is close to the central tendency of the prior. The concentration of the prior distribution about its median can be captured by computing percentile intervals of the marginal prior distribution. The closer the upper and lower percentiles are, the more confident is the researcher a priori that the response is close to the median of the prior. The more the percentiles differ, the more diffuse is the prior. Whether a given degree of concentration of the prior is acceptable to a researcher must be decided on a case by case basis and may depend on extraneous evidence. Finally, percentile intervals being tilted in one direction relative to the prior median rather than being symmetric about the prior median provides an indication of the skewness of the prior impulse response distribution. For example, if the upper endpoint of the percentile interval is further away from the prior median than the lower endpoint, the researcher believes that with high probability the response will be close to the prior median, if it is below the median, whereas the probability mass is more spread out above the prior median, indicating greater uncertainty about the value of the response.

The impact response being sign-restricted, as is required in many cases for identification, causes an otherwise weakly informative prior to be truncated at zero. The resulting prior will still monotonically decline with the distance from zero and may still be viewed as weakly informative apart from

the sign restriction because we have not added any other information beyond the sign. Of course, the prior median in that case will no longer be centered on zero. The value of the truncated prior density increasing, as the value of the impact response moves away from zero, before decreasing for larger deviations, in contrast, implies a unimodal prior distribution with a peak away from zero that would require additional justification, making the prior strongly informative.

Moving from the impact responses to the responses at higher horizons, a researcher with a weakly informative prior would not wish to take a stand on the sign of the impulse responses at longer horizons either. Ideally, the central tendency of a weakly informative impulse response priors at longer horizons should follow the zero line with the probability mass distributed symmetrically across the zero line. For example, under absolute loss, the prior would be sign-neutral if the median is zero at all horizons greater than zero because in that case exactly 50% of the probability mass is in the positive region and 50% in the negative region at each horizon. Imposing sign restrictions at horizons greater than zero obviously violates sign neutrality. If that is all the information the prior embodies, one would expect the marginal prior density to smoothly decay, as the distance from zero grows, as in the case of the impact response.

Finally, narrative restrictions or combinations of sign and zero restrictions may affect the impulse response prior on impact and at longer horizons more generally. For example, in some cases narrative restrictions may simply tighten the prior, whereas in others they may affect the shape of the impulse response prior.

In practice, it is useful to consider the priors for the impact responses and the responses at longer horizons separately. For example, it can be shown that a prior may be only weakly informative about the responses at horizons greater than zero, while being strongly informative about the impact responses, or it may be weakly informative about the impact response but strongly informative about the responses at longer horizons. It is also worth pointing out that a white noise prior for

the slope parameters of the growth rate of a variable amounts to a random walk prior for the log-level of the variable. Thus, such a prior may be symmetric about the medians of the growth rate responses, but not symmetric about the medians of the cumulative responses. Conversely, a random walk prior for the log-level of a variable may be clearly informative about the persistence of log-level responses, but imply a symmetric prior for the growth rate responses. In general, the information content of a prior may differs depending on which response we are interested in.

3.2 Strongly informative impulse response priors

Another situation is that the researcher has a prior that is strongly informative about the impulse responses, perhaps motivated by the implications of economic theory or of natural experiments. For example, one might form an impulse response prior based on a dynamic stochastic general equilibrium model (see, e.g., Plagborg-Møller 2019) or one may be guided by historical evidence of how large or persistent the effects of a given shock are. Strongly informative impulse response priors exhibit a peak (and median) at impulse response values other than zero. Typically such a prior will reflect the researcher's personal beliefs informed by economic reasoning.

Alternatively, a researcher may play devils' advocate by using a strongly informative prior that is diametrically opposed to his or her own beliefs. The idea is that successfully overturning the opposing belief tends to be more persuasive than providing evidence for one's own views based on a prior that favors that view. Even when a negative response is perfectly reasonable given current scientific information, as Gelman et al. (2013) stress, a researcher might want the prior distribution to lean against his or her prior views, which amounts to requiring a higher standard of proof. Likewise, even when a researcher has views about the likely dynamic responses of inflation and the output gap to a monetary policy tightening, informed by macroeconomic theory, the researcher may prefer a weakly informative impulse response prior to a prior that commits to a

specific theoretical model or class of models.

3.3 Practical implications

While the ability of researchers to articulate an impulse response prior is widely accepted, the challenge is that users of the Gaussian-inverse Wishart-Haar prior do not control the implicit impulse response prior when working with the Gaussian-inverse Wishart-Haar prior. This explains why it is essential in applied work to make explicit the impulse response prior implied by the Gaussian-inverse Wishart-Haar prior. In practice, some impulse response priors will be reasonable, while others may not be. Without seeing the prior neither the author nor the reader of a study knows whether this prior agrees with one's personal views, which is problematic unless the identified set is sufficiently narrow for the substantive conclusions to be robust to the choice of the prior, as discussed in Section 5.

In the next section, we discuss four empirical examples that rely on the Gaussian-inverse Wishart-Haar prior. Each example is motivated by a different economic setting and relies on different identifying assumptions. We estimate the implied impulse response prior by simulation and illustrate how this prior depends on the identifying assumptions. As we will show, no two models share the same impulse response prior, making it important to analyze the prior on a case-by-case basis. We examine what the impulse response priors are in each of these studies, to what extent they are informative, and in what sense they are or are not economically reasonable.

4 Empirical illustration 1

Our first example is a stylized quarterly VAR(4) model of U.S. monetary policy including the output gap (*gap*), deflator inflation (π) and the federal funds rate (i). The model imposes the same impact sign restrictions as Ouliaris and Pagan (2016). These identifying restrictions have also been

employed in Peersman (2005) as part of a larger model.

$$\begin{pmatrix} u_t^{gap} \\ u_t^{\pi} \\ u_t^{i} \end{pmatrix} = \begin{bmatrix} + & + & - \\ - & + & - \\ - & + & + \end{bmatrix} \begin{pmatrix} w_t^{\text{cost}} \\ w_t^{\text{demand}} \\ w_t^{\text{monetary policy}} \end{pmatrix}. \tag{3}$$

When implementing the Gaussian-inverse Wishart-Haar prior, for now we follow the convention of centering the prior for A on an independent white noise mean. This results in a largely sign-neutral impulse response prior at longer horizons, as illustrated in Figure 1. Given the sign restrictions, the impact responses of the output gap and inflation rate are bounded by zero from above. In the absence of further specific knowledge about these impact responses, a reasonable view is that values of the impact responses that are far from zero must have lower probability than values closer to zero. For example, any macroeconomist would consider a very large decline in the impact response of the output gap to a 25 bp monetary tightening a priori less likely than an impact response closer to 0%. A similar argument can be made for inflation. Thus, a natural view would be that the density value should in general decline, the more the response differs from zero with most of the probability mass of the impact responses concentrated close to zero. From this point of view the two histograms in the upper panel Figure 1 are not unreasonable and compatible with a researcher using a prior about the responses that is weakly informative apart from the identifying sign restrictions. While the density of the impact response of inflation peaks for values slightly below zero rather than at zero, the difference is small and one would be hard pressed to find fault with this prior.

The lower panel in Figure 1 indicates that at horizons greater than zero the marginal prior distributions of the responses of inflation and the output gap are sign-neutral in that they are centered on zero with a roughly symmetric distribution of the probability mass about zero. This

prior is consistent with the researcher not being willing to take a stand on the sign of the responses, but it allows for departures from the zero line in either direction. The priors are fairly tight about the medians, as indicated by the 68% and 90% percentile intervals. This is not say that every user of this model will necessarily agree with this impulse response prior, but it highlights that - despite the fact that the Gaussian-inverse Wishart-Haar prior was developed mainly for computational reasons - the implied impulse response prior in practice may appeal to some users. If it does, there is nothing wrong with conducting posterior inference by conventional methods in this model.

Alternatively, one could have examined this model under a Gaussian-inverse Wishart-Haar prior with the prior for A centered on an independent random walk mean. Such priors have been popular in the forecasting literature because random walk priors tend to generate more accurate out-of-sample forecasts. While such a prior may not be natural in the current empirical example, replacing the white noise mean by a random walk mean is a useful thought experiment, nevertheless, to illustrate how assumptions about the reduced-form prior may affect the implicit impulse response prior. As shown in Figure 2, this change has no material effect on the shape of the marginal priors for the impact period, but it implies that the impact response under the prior is propagated across all horizons. In other words, we are taking a strong stand on the sign of the responses at longer horizons. We are also taking a stand on the inflation and output gap responses declining at the same time in expectation. The prior uncertainty greatly increases on impact to the point that the lower end of the 90% percentile region is not even contained in the plot. While the prior uncertainty about the responses increases sharply at longer horizons, as expected, it is not symmetric about zero. Thus, this alternative prior specification is strongly informative about the sign of these responses and should only be used when the user is comfortable with its implications for the impulse response prior. We conclude that a Gaussian-inverse Wishart-Haar prior with the prior of A centered on zero at all lags is the preferred prior specification for this model, if the user is aiming for a prior

that is only weakly informative about the sign of the responses of the model variables.

5 Empirical illustration 2

Next we consider a monthly structural VAR(12) model of the effects of gasoline price shocks on headline inflation and inflation expectations in Kilian and Zhou (2022). The variables include the log real price of U.S. motor gasoline ($rpgas$), CPI headline inflation (π) and the mean one-year inflation expectation in the Michigan Survey of Consumers (π^{exp}). The model explains variation in these three variables in terms of three mutually uncorrelated structural shocks: (1) A nominal gasoline price shock; (2) a shock to the CPI excluding gasoline; and an idiosyncratic shock to household inflation expectations that is not reflected in current prices. The identification relies on impact sign restrictions and exclusion restrictions.

$$\begin{pmatrix} u_t^{rpgas} \\ u_t^{\pi} \\ u_t^{\pi^{\text{exp}}} \end{pmatrix} = \begin{bmatrix} + & - & 0 \\ + & + & 0 \\ + & + & + \end{bmatrix} \begin{pmatrix} w_t^{\text{nominal gasoline price}} \\ w_t^{\text{CPI excl. gasoline}} \\ w_t^{\text{idiosyncratic inflation expectation}} \end{pmatrix} \tag{4}$$

The model is estimated based on a conventional uniform-Gaussian-inverse Wishart prior with the prior means of all slope parameters set to zero. Since identification is achieved by a combination of sign and exclusion restrictions on the impact responses, simulating the impulse responses requires the use of the importance sampler proposed in Arias et al. (2018).

For expository purposes, the nominal gasoline price shock has been normalized to imply a 10% increase in the real price of gasoline, which implies a shock to the nominal gasoline price of about 10% as well. While the marginal prior for the impact response of inflation expectations in Figure 3 with some goodwill may be viewed as only weakly informative by the same reasoning as in the first empirical example, the prior for the impact response of headline inflation shows a pronounced peak

near 0.2% and is clearly informative. The fact that this prior is unimodal with a peak away from zero, however, is consistent with economic reasoning. Given that consumer spending on gasoline and other motor fuel combined accounts for 2.89% of all consumer expenditures over the estimation period, according to the BEA, the direct effect of a 10% nominal gasoline price shock on headline inflation would be expected to be somewhat less than 0.289% on impact, which is not far from both the peak in the histogram of the impact response of inflation and the median response. Given that the uncertainty about the gasoline consumer expenditure share is two-sided, a unimodal prior for the impact response of inflation centered near 0.2% is not an unreasonable starting point. Since the prior medians at longer horizons are consistently near zero with the 68% and 90% percentile intervals either balanced in either direction or tilted in the economically more plausible direction of positive responses, as shown in the lower panel of Figure 3, we may conclude that the impulse response prior implied by the use of the Gaussian-inverse Wishart-Haar prior in Kilian and Zhou (2022) is at best modestly informative about the impulse responses.

The apparent oscillation in the endpoints of the 90% intervals at longer horizons reflects complex roots embodied in the white noise prior for the reduced-form parameter A. Unlike in the previous empirical example, the identifying restrictions do not eliminate these roots. Researchers concerned about this feature of the white noise prior may reduce the amplitude of the oscillation by increasing the lag decay rate from its default value of 2 or may eliminate the oscillation entirely by lowering λ. A lower λ also renders the percentile intervals more balanced about zero at longer horizons and reinforces the conclusion that the prior is only modestly informative.

6 Empirical Illustration 3

Our next example is the widely used monthly VAR(24) model of the global oil market originally proposed in Kilian and Murphy (2014), as updated and implemented in Zhou (2020). The model

variables are the percent change in global crude oil production ($\Delta prod$), an appropriate measure of the global business cycle (rea), the log real price of oil ($rpoil$), and the change in global crude oil inventories (Δinv). The structural shocks include a flow demand shock, a flow supply shock, a storage demand shock and a residual demand shock (implicitly defined as the complement to all other shocks) that captures, for example, changes in the intensity of oil use, as users substitute toward more or less oil-intensive technologies. The structural shocks are identified based on static and dynamic sign restrictions, complemented by elasticity bounds and additional narrative inequality restrictions. The model heavily relies on static sign restrictions.

$$\begin{pmatrix} u_t^{\Delta prod} \\ u_t^{rea} \\ u_t^{rpoil} \\ u_t^{\Delta inv} \end{pmatrix} = \begin{bmatrix} - & + & + \\ - & + & - \\ + & + & + \\ & & + \end{bmatrix} \begin{pmatrix} w_t^{\text{flow supply}} \\ w_t^{\text{flow demand}} \\ w_t^{\text{storage demand}} \\ w_t^{\text{residual demand}} \end{pmatrix} \tag{5}$$

The sign restrictions associated with the flow supply shock are applied not only on impact, but for the first 12 months. Further details of the construction of the data and of the elasticity bounds and narrative restrictions can be found in Zhou (2020). Whereas Zhou (2020) utilized an improper diffuse prior for the reduced-form parameters as in Uhlig (2005), here we employ a uniform-Gaussian-inverse Wishart prior specification analogous to that used in the earlier empirical examples. This change has no material effect on the posterior estimates.

For expository purposes, we examine the prior obtained without the narrative restrictions. In the interest of space we focus on the responses of global real activity and the real price of oil to flow supply, flow demand and storage demand shocks. Figure 4a shows that in this case the model assigns fairly similar informative priors for the impact responses of the real price of oil to the flow supply, flow demand and storage demand shock. The respective prior medians are all near

5% and the priors allow for considerable uncertainty about the impact effects. Thus, while these priors clearly are informative in that they embody the belief that all three shocks have nontrivial effects on the real price of oil with impact responses near zero unlikely, they largely treat these shocks symmetrically. The corresponding impact priors for global real activity put substantial probability mass on responses between 0 and 15 for the flow supply and the storage demand shock, consistent with the latter shock representing geopolitical news about oil supply disruptions rather than demand news which would be associated with higher future real activity (see, e.g., Alquist and Kilian 2010). Both priors taper off for larger deviations from zero. The same pattern is found for the flow demand shock, except that there is less probability mass on large increases before the prior tapers off.

Figure 4b extends this analysis to longer horizons. Here it becomes apparent that the impulse response prior favors persistently large effects from flow supply shocks at horizons greater than zero, which is a direct consequence of the dynamic sign restrictions on the responses to flow supply shocks. As a result, the prior medians in Figure 4b show a positive response of the real price of oil to a flow supply disruption in the oil market at all horizons and a similarly persistent negative response of global real activity. In contrast, except on impact, the prior medians of the real price of oil and real activity responses to the flow demand shock tend to be close to zero. Thus, this prior is consistent with the belief that shocks to the flow supply of oil and fears of oil supply disruptions play a greater role in driving the real price of oil than flow demand shocks associated with the global business cycle.

This prior closely resembles the traditional view in the oil market literature from the 1970s to the 1990s that the main determinants of the real price of oil are shocks to the flow supply of oil. In contrast, the more modern view that flow demand shocks associated with the global business cycle cause persistent fluctuations in the real price of oil and are the most important driver of the real

price of oil is downplayed by this prior. Thus, this prior stacks the deck against the modern view that oil demand shocks matter more for the price of oil and for global real activity than oil supply shocks. This conclusion is only strengthened when incorporating further narrative restrictions in the prior. The prior also allows for the view that shifts in oil price expectations associated with geopolitical events in the Middle East operating through storage demand shocks may rival actual oil supply disruptions as a determinant of the real price of oil. Researcher may choose this prior because it truly reflects their personal beliefs or they may adopt this prior precisely because they wish to persuade other researchers of the opposing view that global real activity is largely driven by flow demand shocks. The point of this example is again that despite the Gaussian-inverse Wishart-Haar prior having been chosen for computational efficiency, the implicit impulse response prior is not necessarily devoid of economic interpretation.

7 Empirical Illustration 4

Our last example is the classical VAR(12) model of monetary policy shocks in Uhlig (2005). Uhlig 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 six months. The response of real GDP remains unrestricted. All model variables are in levels. Rather than relying on improper reduced-form priors, as Uhlig (2005) did, we rely on a Gaussian-inverse Wishart-Haar prior with the prior mean of A set to independent random walks and the own-standard deviation of the first autoregressive lag, λ, set to 0.01. This does not change the substantive conclusion in Uhlig (2005) that the response of real GDP to a monetary tightening is indistinguishable from zero and possibly positive. Increasing the value of λ tends to increase the number of explosive draws. The upper left plot in Figure 5 suggests that the implied prior for the impact response of real GDP is only weakly informative with a prior median of -0.007. As the lower left plot shows, the prior

medians in the lower panel are also close to zero at all horizons. The prior allows for considerable uncertainty at longer horizons, as indicated by the widening 68% and 90% percentile intervals. Thus, this impulse response prior may be viewed as weakly informative.

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 funds rate in selected months. Adding these restrictions increases the concentration of the prior without affecting the central tendency, which remains close to zero at 0.002 on impact, or the approximate symmetry of the prior about zero. Most probability mass is concentrated in the range between -1 and $+1$ in the upper right plot in Figure 5, while the 68% and 90% percentile intervals at all horizons are much closer in the lower right plot. Thus, the narrative restrictions produce a tighter prior than the Uhlig (2005) model, but the impulse response prior remains weakly informative.

8 What if the implied impulse response prior is not agreeable?

The point of our analysis is not to persuade the reader that the priors in the empirical illustrations are representative of everyone's views. Rather our point is that these priors are not obviously nonsensical or unintentionally informative. Some researchers will find these priors agreeable, while others may not. Nor do we claim that our results are necessarily representative. Rather we stress the importance of authors making explicit the impulse response prior implied by the Gaussian-inverse Wishart-Haar prior, as discussed in this paper, and defending it.

This leaves the question of how to proceed if one finds oneself in disagreement with the implied impulse response prior. It may seem that the impulse response prior being unintentionally informative for a given model means that the substantive conclusions of this study are questionable

and should be discarded. Recent research shows that this is not the case necessarily. Inoue and Kilian (2025) provide evidence that the reported conclusions from sign-identified VAR models are unaffected by the use of the Haar prior if the identification is sufficiently tight. This condition may be evaluated empirically. Inoue and Kilian 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. Their approach builds on the estimate of the identified set of the impulse response proposed in Giacomini and Kitagawa (2021) 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 estimated 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 and nothing stands in the way of using the conventional Gaussian-inverse Wishart -Haar prior, even when the implied impulse response prior is unreasonable from an economic point of view. For example, the identified set of the impulse responses in Kilian and Zhou (2022) is quite tight, whereas that in Uhlig (2005) is not.

9 Concluding remarks

Several studies have voiced concerns that the Gaussian-inverse Wishart-Haar prior typically used in estimating sign-identified structural VAR models, as popularized by Rubio-Ramirez et al. (2010), Arias et al. (2018) and Antolin-Diaz and Rubio-Ramirez (2018), building on Uhlig (2005), may be unintentionally informative. As noted by Watson (2020), only "good" impulse response priors

lead to "good" inference in such models, where a "good" impulse response prior is generally understood to mean a proper prior that is not unintentionally economically informative. We proposed appropriate tools to differentiate good impulse response priors from bad ones on a case-by-case basis and illustrated their use in a range of empirical applications. We made the case that applied researchers should routinely report the impulse response prior implied by their econometric model to allow readers to discount studies based on priors they disagree with. Even if one disagrees with the prior used in a given study, however, it can be shown that the choice of the parameter prior becomes irrelevant when the identification of the structural responses is sufficiently tight. Thus, posterior inference based on the conventional prior used in estimating sign-identified structural VAR models is justified unless the identified set of the structural impulse responses is wide and the implicit impulse response prior is unreasonable.

Acknowledgments

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 thank John Landon-Lane for helpful discussions. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

1. Alquist, R., Kilian, L., 2010. What do we learn from the price of crude oil futures? Journal of Applied Econometrics 25, 539-573.

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

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

4. Arias, J., Rubio-Ramirez, J.F., Waggoner, D.F., 2025. Uniform priors for impulse responses. Econometrica, forthcoming.

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

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

8. Gelman, A., Carlin, J.B., Stern, H.S., Dunson, D.B., Vehtari, A., Rubin, D.B., 2013. Bayesian Data Analysis. 3rd ed. Chapman & Hall, Boca Raton, FL.

9. Giacomini, R., Kitagawa, T. 2020. Robust Bayesian inference for set-identified models. Econometrica 89, 1519-1556.

10. Gustafson, P., 2009. What are the limits of posterior distributions arising from nonidentified models and why should we care? Journal of the American Statistical Association 104, 1682-1695.

11. Inoue, A., Kilian, L., 2022. Joint Bayesian inference about impulse responses in VAR models. Journal of Econometrics 231, 457-476.

12. Inoue, A., Kilian, L., 2025. When Is the Use of Gaussian-inverse Wishart-Haar Priors Appropriate? Journal of Political Economy, forthcoming.

13. Karlsson, S., 2013. Forecasting with Bayesian vector autoregressions. Handbook of Economic Forecasting, vol 2B, ed. by G. Elliott and T. Timmermann. North Holland, Elsevier, 791-897.

14. Kilian, L., 2025. Impulse response diagnostics for priors on parameters in structural vector autoregressions. Working Paper, Federal Reserve Bank of Dallas.

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

16. 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.

17. Kilian, L., Zhou, X. 2022. Oil Prices, Gasoline Prices and Inflation Expectations: A New Model and New Facts. Journal of Applied Econometrics, 37, 867-881.

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. Poirier, D. 1998. Revising beliefs in nonidentified models. Econometric Theory 14, 483-509.

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

23. Sims, C.A., Zha, T., 1999. Error bands for impulse responses. Econometrica 67, 1113-1156.

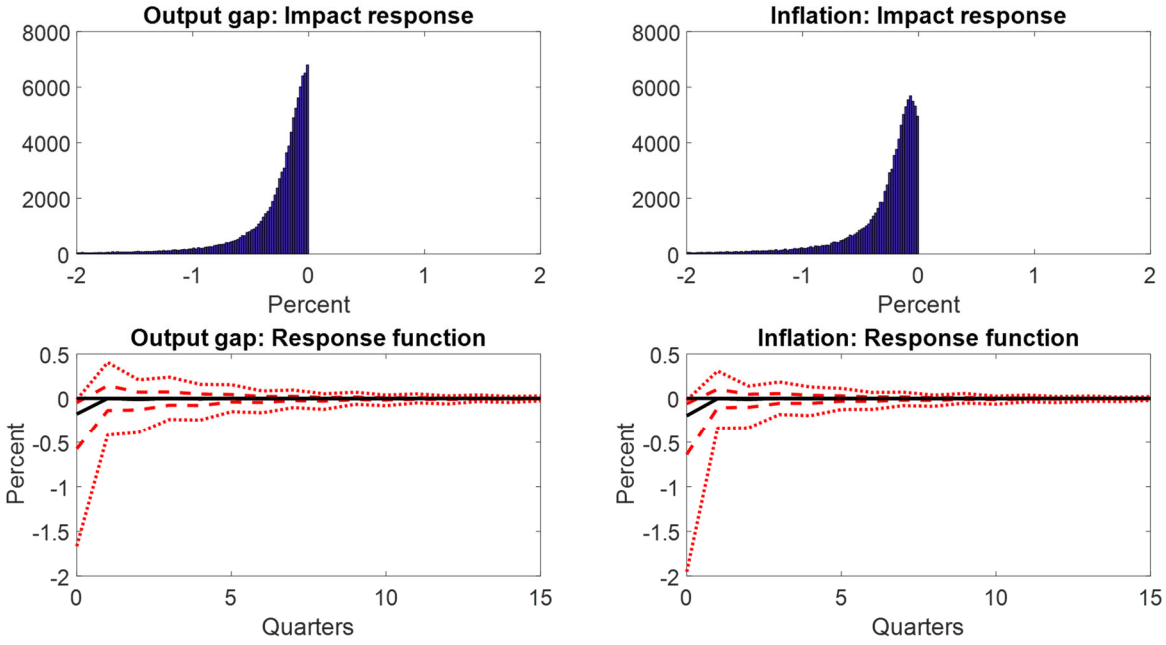
24. Uhlig, H., 1994. What macroeconomists should know about unit roots: A Bayesian perspective. Econometric Theory 10, 645–671.

25. 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.

26. 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.

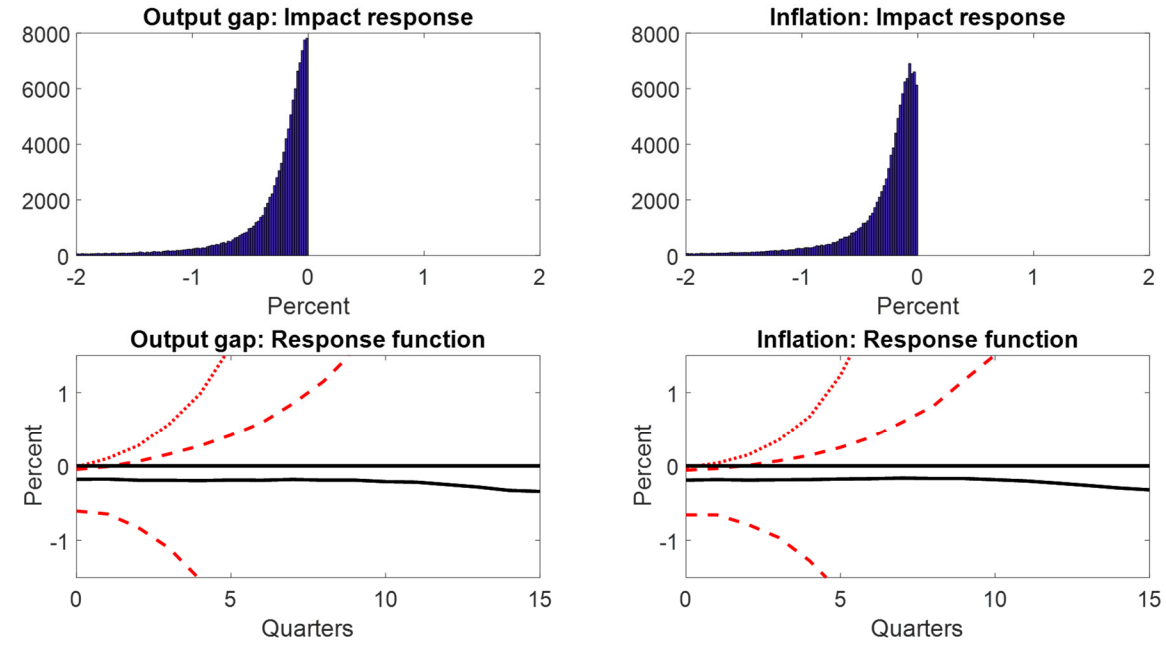
27. Zhou, X., 2020. Refining the workhorse oil market model, Journal of Applied Econometrics 35, 130-140.

Figure 1: Marginal prior about responses to 25 bp monetary policy shock in Ouliaris and Pagan (2016)



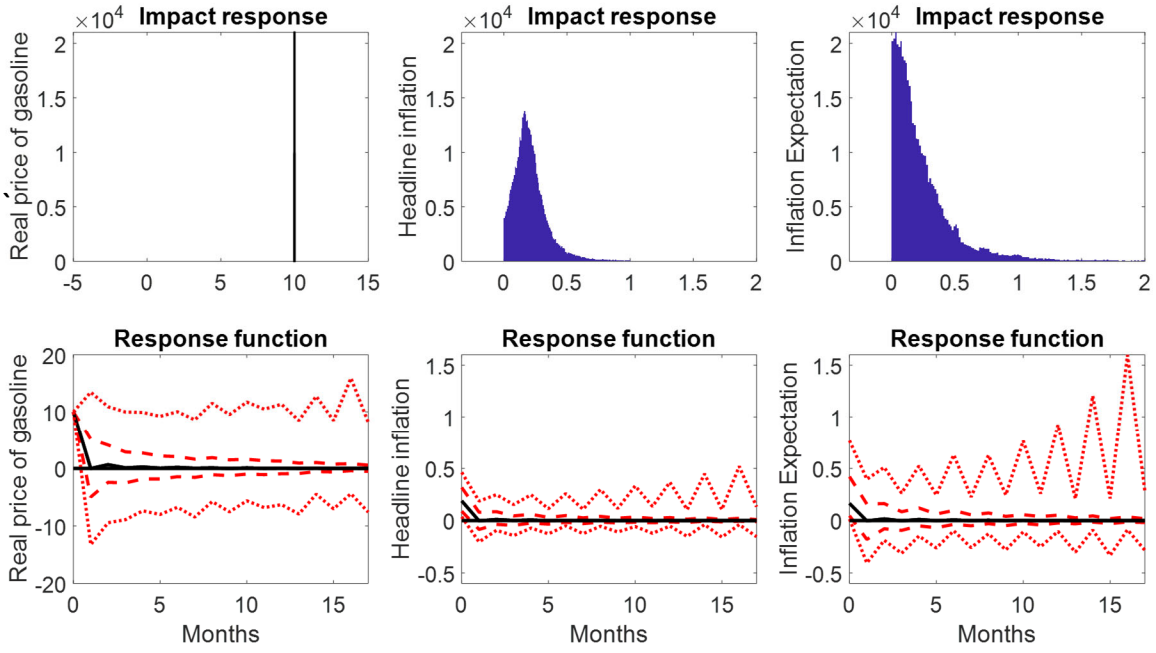
NOTES: Based on white noise prior mean for A. The lower panel shows vectors of pointwise prior medians and the 68% and 90% percentile intervals of the marginal prior distributions.

Figure 2: Marginal prior about responses to 25 bp monetary policy shock in Ouliaris and Pagan (2016)



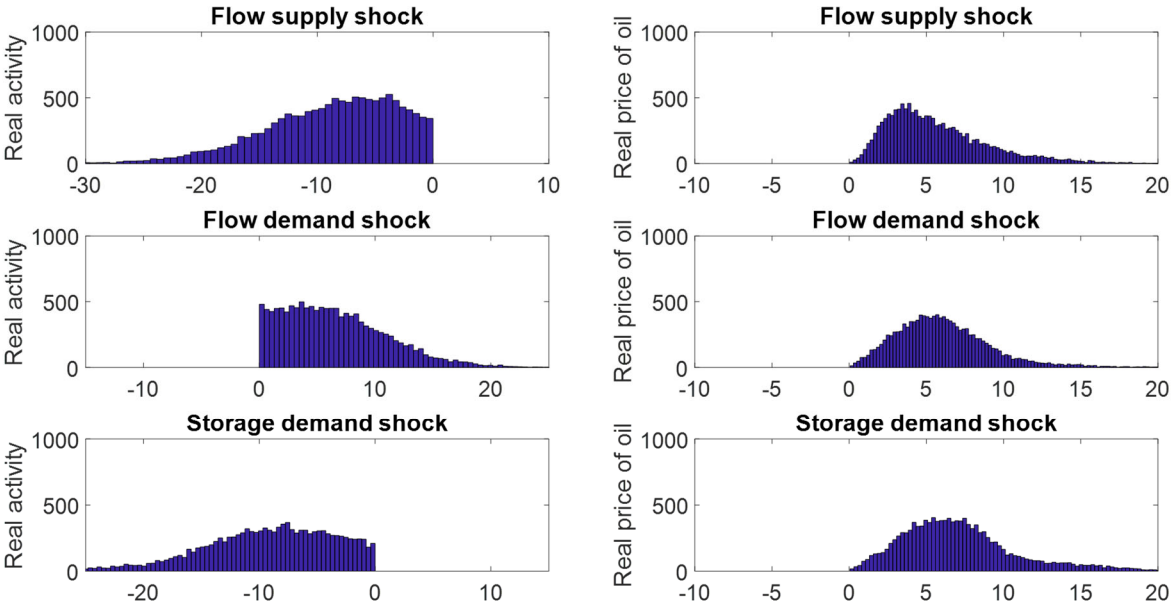
NOTES: Based on random walk prior mean for A. The lower panel shows vectors of pointwise prior medians and the 68% and 90% percentile intervals of the marginal prior distributions.

Figure 3: Marginal prior about responses to 10% gasoline price shock in Kilian and Zhou (2022)



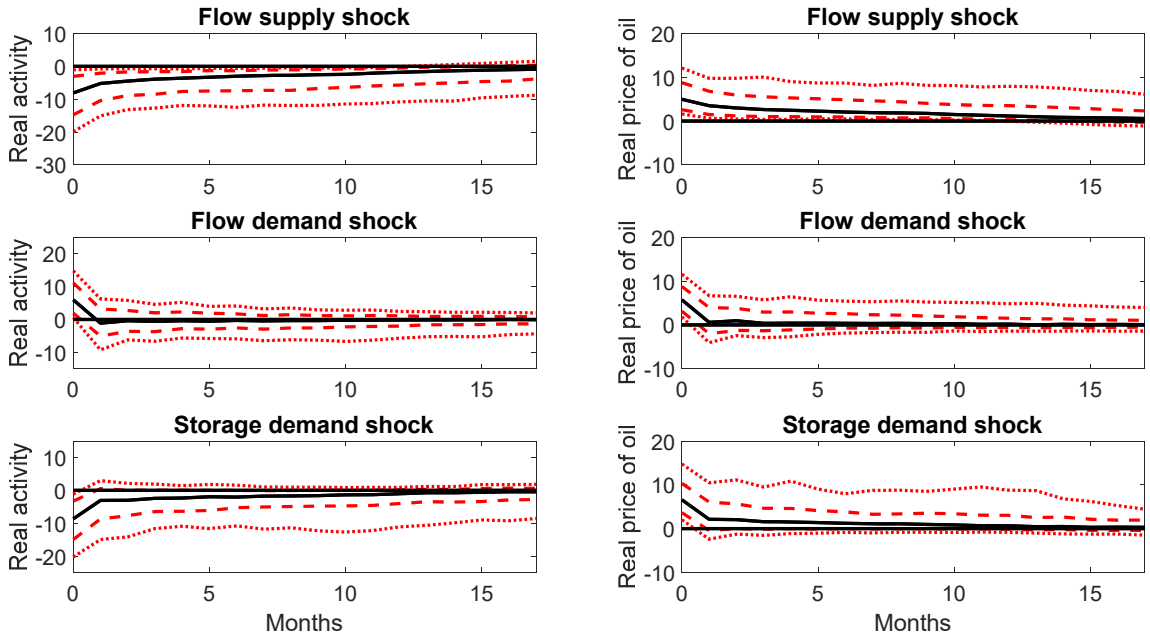
NOTES: Based on white noise prior mean for A. The lower panel shows vectors of pointwise prior medians and the 68% and 90% percentile intervals of the marginal prior distributions.

Figure 4a: Marginal priors of impact impulse responses in Kilian and Murphy (2024)



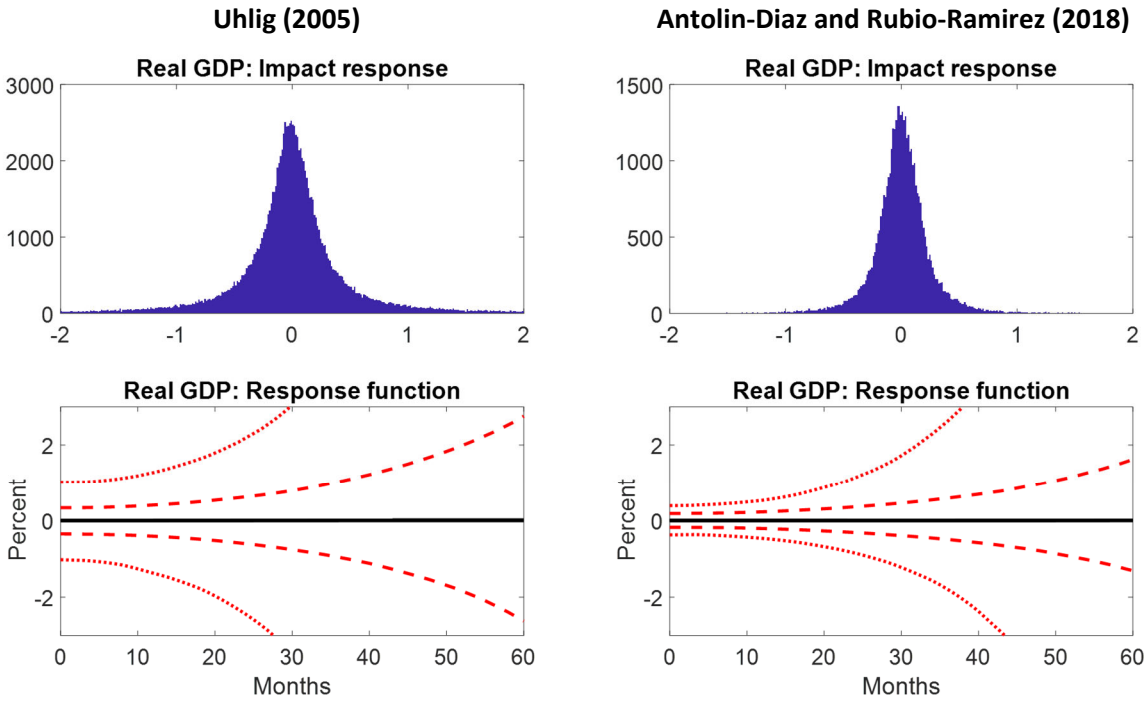
NOTES: Based on white noise prior mean for A.

Figure 4b: Marginal impulse response priors in Kilian and Murphy (2014)



NOTES: Based on white noise prior mean for A. The plot shows vectors of pointwise prior medians and the 68% and 90% percentile intervals of the marginal prior distributions.

Figure 5: Prior about real GDP response to 25bp monetary tightening



NOTES: Based on random walk prior mean for A. The lower panel shows vectors of pointwise prior medians and the 68% and 90% percentile intervals of the marginal prior distributions.