# Comment on Giacomini, Kitagawa and Read's 'Narrative Restrictions and Proxies'

Lutz Kilian

**Working Paper 2117** **December 2021**

Research Department

https://doi.org/10.24149/wp2117

# Comment on Giacomini, Kitagawa and Read's 'Narrative Restrictions and Proxies'*

Lutz Kilian†

First draft: October 16, 2021
This version: December 2, 2021


## Abstract

In a series of recent studies, Raffaella Giacomini and Toru Kitagawa have developed an innovative new methodological approach to estimating sign-identified structural VAR models that seeks to build a bridge between Bayesian and frequentist approaches in the literature. Their latest paper with Matthew Read contains thought-provoking new insights about modeling narrative restrictions in sign-identified structural VAR models. My discussion puts their contribution into the context of Giacomini and Kitagawa's broader research agenda and relates it to the larger literature on estimating structural VAR models subject to sign restrictions.


**JEL code:** C11, C32, C52

**Keywords:** Structural VAR, single prior, multiple prior, posterior, joint inference, impulse response, narrative restrictions.

---

* The views expressed in the paper are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. I thank Atsushi Inoue, Mikkel Plagborg-Møller and Juan Rubio-Ramirez for helpful discussions. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

† Lutz Kilian, Federal Reserve Bank of Dallas, Research Department, 2200 N. Pearl St., Dallas, TX 75201, USA and CEPR. E-mail: kilian2019@gmail.com.

# 1 Introduction

In a series of recent studies, Raffaella Giacomini and Toru Kitagawa have developed an innovative new methodological approach to estimating sign-identified structural VAR models that seeks to build a bridge between Bayesian and frequentist approaches in the literature. Their latest paper with Matthew Reed contains thought-provoking new insights about modeling narrative restrictions in sign-identified structural VAR models. My discussion puts their contribution into the context of Giacomini and Kitagawa's broader research agenda and relates it to the larger literature on estimating structural VAR models subject to sign restrictions. It is useful to start with the conventional Bayesian approach to estimating sign-identified VAR models, which has remained the most widely approach used by practitioners interested for answering substantive economic questions.

# 2 Recent Critiques of Conventional Bayesian Priors

The conventional approach to estimating sign-identified structural VAR models is based on a uniform prior for the orthogonal matrix $Q$ (also known as a Haar prior) and a Gaussian-inverse Wishart prior for the reduced-form parameters $A$ and $\Sigma$, which denote the slope parameters and error covariance matrix, respectively (see, e.g., Rubio-Ramirez, Waggoner and Zha 2010; Arias, Rubio-Ramirez and Waggoner 2018; Antolin-Diaz and Rubio Ramirez 2018). An obvious concern is that any priors for the VAR model parameters may be unintentionally informative for the implied prior of the impulse response vector $\theta$, obtained by stacking the structural impulse responses of interest. This impulse response vector can be expressed as a nonlinear function $g(\cdot)$ of the model parameters $Q$, $A$ and $\Sigma$ such that $\theta = g(Q, A, \Sigma)$. By the change-of-variable method we know that any prior specification for $Q$, $A$ and $\Sigma$ may imply an informative prior for $\theta$.

The concern is that most researchers do not to derive this implied impulse response prior, and hence are unaware of how informative this impulse response prior is. Inoue and Kilian (2021a) show by example that in realistic settings even strong priors about $A$ and $\Sigma$ tend to be overturned by the data, when they differ substantially from the information in the likelihood. In contrast, whatever prior we assume for $Q$ will not be revised by the data, since the rotation matrix $Q$ does not enter the likelihood, and hence may result in unintentionally informative and economically implausible impulse response priors, regardless of the length of the estimation period. This observation has stimulated a growing literature on the estimation of sign-identified structural VAR models including Baumeister and Hamilton (2015, 2019, 2020), Gafarov, Meier and Montiel Olea (2018), Giacomini and Kitagawa (2021), Giacomini, Kitagawa and Read (2021), Giacomini, Kitagawa and Uhlig (2019), Granziera, Moon and Schorfheide (2018), Inoue and Kilian (2021a,b), Plagborg-Møller (2019), and Wolf (2020), among others.

### 2.1 The Baumeister and Hamilton Critique

Baumeister and Hamilton (2015, 2019) considerably sharpened the standard critique of the conventional approach to estimating sign-identified structural VAR models by making four specific claims:

1. Conventional priors are unintentionally informative for the implied impulse response prior.
2. The impulse response posterior is dominated by this unintentionally informative prior.
3. As a result, all empirical results obtained using the conventional approach should be discarded. This includes many highly influential studies published in leading journals.
4. The conventional approach should be replaced by an alternative class of priors they propose to remedy this problem.

These claims have received considerable attention in the literature (see, e.g., Giacomini and Kitagawa 2021; Watson 2020; Wolf 2020). As shown in Inoue and Kilian (2021a), however, none of these claims is correct. As to the first claim that conventional priors are by design unintentionally informative, Baumeister and Hamilton (2015) propose computing the distribution $h(\theta|A, \Sigma)$, where $\theta$ is the impulse response vector and $A$ and $\Sigma$ denote the slope parameters and error covariance matrix, respectively, as a measure of the prior distribution. They provide stylized examples where this "prior" distribution looks unintentionally informative. The problem is that the prior in the conventional approach, as exemplified by Uhlig (2005), Rubio-Ramirez, Waggoner and Zha (2010), Arias et al. (2018) and Antolin-Diaz and Rubio Ramirez (2018), does not condition on these parameters, but treats them as random with a Gaussian-inverse Wishart distribution. Thus, the numerical examples in Baumeister and Hamilton (2015) do not speak to the shape of the impulse response prior implied by the conventional approach. This is a question only addressed in subsequent work by Inoue and Kilian (2021a) who illustrate that in typical applications the impulse response priors underlying the conventional approach look completely different from those reported by Baumeister and Hamilton and much more reasonable. This is not to say that Inoue and Kilian disagree with the point that the impulse response prior is determined by the prior for $Q$ asymptotically. Rather their point is that the extent of the uncertainty about $\theta$ induced by the prior for $Q$ tends to be much smaller in practice than conjectured by Baumeister and Hamilton.

Baumeister and Hamilton provide no evidence in support of their second claim that the impulse response posterior tends to be dominated by the implicit impulse response prior, but Inoue and Kilian (2021a) demonstrate that in typical applications the impulse response posterior is dominated by the data rather than the prior. In other words, the substantive insights generated by these studies were not present in the prior, but are present in the posterior. Given that there is no support for Baumeister and Hamilton's first two claims, their call for disregarding all empirical evidence from previous studies is also moot.

As to their last claim, Baumeister and Hamilton (2015) fail to realize that the alternative class of priors they propose suffers from exactly the same potential drawback as the conventional approach. Their proposal, in its simplest form, is to work with the structural VAR presentation, $B_0 y_t = B_1 y_{t-1} + ... + B_p y_{t-p} + w_t$, where $y_t$ is the data and $w_t$ is a Gaussian iid error, and to impose priors on the parameters $B_0$, $B_1$, ..., $B_p$. However, it is readily apparent that an explicit prior on $B_0$, in particular, may be unintentionally informative for the prior on $\theta = f(B_0, B_1, ..., B_p)$, because $\theta$ is obtained based on a nonlinear transformation $f(\cdot)$ of the structural parameters. The same problem arises if, in addition, we specify a prior on some elements of $B_0^{-1}$. Thus, Baumeister and Hamilton's approach should never be used without examining the implied prior for $\theta$.

Inoue and Kilian (2021a) demonstrate that the structural parameter prior specified in Baumeister and Hamilton (2019), for example, implies an unintentionally informative and highly economically implausible impulse response prior, whereas related models based on the conventional prior for $A$, $\Sigma$ and $Q$ do not have that problem. Thus, there is no basis for discriminating against one approach at the expense of the other. This result is not surprising since the approach of specifying a prior on the structural model parameters in a simultaneous equation model is equivalent to implicitly specifying a prior on $A$, $\Sigma$ and $Q$. Effectively, Baumeister and Hamilton are specifying a single prior on $Q$ that differs from the conventional Haar prior. As with any single prior on $Q$, we need to consider the implications of that prior for the prior distribution of $\theta$ on a case-by-case basis.[1]

[1] Baumeister and Hamilton argue that eliciting a prior for the impact demand and supply elasticities, which can be expressed as a function of the elements of $B_0$, is more natural and more credible than eliciting sign restrictions on functions of elements of $B_0^{-1}$ and a prior for $A$, $\Sigma$ and $Q$. As discussed in Kilian (2021), this is not the case because extraneous elasticity estimates do not in general conform to the definition of an elasticity in the structural model, but do imply restrictions on $B_0^{-1}$. For example, the elasticity priors postulated in Baumeister and Hamilton's (2019) work are ad hoc and cannot be justified based on the extraneous evidence they cite.

In short, the analysis in Inoue and Kilian (2021a) shows that the concern over unintentionally informative $Q$ priors in the conventional approach has been greatly overstated. The conventional approach, far from being obsolete, remains a contender for applied work, if used with care.

## 2.2 A robust Bayesian approach

Giacomini and Kitagawa (2021) propose a creative solution to the problem of a single prior for $Q$ being potentially informative for the impulse response prior. Like Granziera. Moon and Schorfheide (2018), they take the position that we should not take a stand on the prior for $Q$. Instead, they propose an agnostic multiple-prior approach for $Q$, which implicitly allows for any possible prior for $Q$ and only requires taking a stand on the sign restrictions. Their approach thus is robust to the choice of the prior for $Q$ and yields two types of summary statistics. One is a consistent estimate of the identified set for each impulse response. This is a major improvement on the work of Granziera et al. (2018), which only provided robust pointwise confidence sets for the impulse responses. The other summary statistic is the pointwise credible set that is associated with this identified set.[2]

### 2.2.1 How to interpret the impulse response estimates

Giacomini and Kitagawa (2021, p. 1521) explicitly note that their paper does not intend to provide any normative argument as to whether one should adopt a single prior or multiple priors in set-identified models. Their stated goal is to offer new tools for inference that complement existing tools including both the conventional single-prior approach and the single-prior approach discussed in Baumeister and Hamilton (2015). The limitation of their approach is that it is makes it more difficult to interpret the estimates from an economic point of view. In practice, Giacomini and

[2] Not surprisingly, the robust pointwise credible sets proposed by Giacomini and Kitagawa tend to be wider than standard pointwise credible sets for impulse responses because they embody additional uncertainty about the prior for $Q$. In fact, their pointwise credible sets are asymptotically equivalent to the frequentist pointwise confidence set derived in Granziera, Moon and Schorfheide (2018), but Giacomini and Kitagawa's approach is applicable to a wider class of models, is more computationally convenient, and does not require strong assumptions about the distribution of the reduced-form impulse responses.

Kitagawa suggest focusing on the probability of a given impulse response being positive or negative and on whether the pointwise credible set includes zero. It is hard to think of concrete macroeconomic questions that can be addressed with these tools. For example, an applied researcher would not start by postulating that the response of real output to a monetary policy shock six months after the shock should be negative. Even if a researcher suspected that the response at that horizon should be negative, the same reasoning would suggest that other responses nearby should also be negative, invalidating pointwise inference. What the applied researcher actually would be interested in is the shape of the response functions and, more generally, the co-movement across response functions. The concern is that both the identified set and the robust credible set tend to accommodate a wide range of response functions with different shapes.[3]

In addition, the pointwise error bands reported in Giacomini and Kitagawa are not appropriate for inference about the vector of impulse responses, $\theta$, obtained by stacking the impulse responses of interest. While they acknowledge that their inference is only correct pointwise, at present there is no way of conducting joint inference that is robust to the choice of the prior for $Q$. It would be desirable for the authors to extend their method in that direction. This raises the question of whether Giacomini and Kitagawa's posterior distributions could instead be evaluated from a Bayesian point of view using the Bayes estimator of $\theta$ and joint credible sets, as discussed in Inoue and Kilian (2021a,b). Since Giacomini and Kitagawa are willing to make probability statements about individual impulse responses, presumably nothing should stand in the way of evaluating the joint probability distribution of $\theta$, which would greatly enhance the usefulness of their approach from an applied point of view.[4]

[3]For example, in Giacomini and Kitagawa's model I, which corresponds to the model in Granziera et al. (2018), the identified set in their Figure 1 accommodates persistent real output increases in response to a monetary tightening as well as persistent declines. It also accommodate oscillating response functions and effectively permanent changes in real output.

[4]Another way of addressing this problem may be the alternative multiple prior approach proposed in Giacomini, Kitagawa and Uhlig (2019), which incorporates a probabilistic belief in a specific single prior for $Q$ and a set of alternative priors in the neighborhood of this belief. Unlike the method in Giacomini and Kitagawa (2021), this

#### 2.2.2 The implicit impulse response prior

Giacomini, Kitagawa and Uhlig (2019) recognize that both the fully ambiguous multiple prior approach and the approach of relying on a single prior for $Q$ could be a poor approximation of the researcher's prior. This raises the question of how one would judge whether a given prior lines up with the researcher's prior views. It is well understood that in designing a prior we have to focus on one dimension of the structural VAR model. We cannot control priors simultaneously in multiple dimensions. Since there is general agreement among participants in this literature that the impulse response vector is the most important aspect of the structural VAR model to be estimated, it is natural to focus on the implied prior for $\theta$. Thus, a natural question is what the implicit prior for $\theta$ is in Giacomini and Kitagawa's approach and whether this prior reflects the prior views of the researcher.

Given the intensity with which the question has been debated in the recent literature, Giacomini and Kitagawa show remarkably little interest in the impulse response prior implied by their approach. Inoue and Kilian (2021a) show how this question can be addressed in the context of the conventional single-prior approach to estimating sign-identified structural VAR models as well as the single-prior approach of Baumeister and Hamilton (2015). In contrast, there are no such tools for the robust multiple prior Bayesian approach proposed by Giacomini and Kitagawa (2021). We frankly do not know what the implicit impulse response prior is of their Bayesian approach. Although one could apply the same tools they employ in evaluating the posterior to evaluate impulse response draws from the prior distribution of their model, these statistics do not lend themselves to judging the economic content of the prior. For example, if we plot the vertical bars representing the identified set under the prior, this set may be equally consistent with monetary policy shocks causing an increase or a decrease in real output or, for that matter, other shapes of the response

alternative approach generates a point estimate of the impulse responses under additively separable absolute loss.

function. Moreover, a likely outcome is that the change in real output could be permanent or transitory, notwithstanding a consensus in the literature that monetary policy shocks have no permanent effects on real output. Clearly, some of these responses may reflect the researcher's prior views, while others do not, making this model specification a poor approximation to the researcher's prior in general.

### 2.2.3 Lessons for Practitioners

The use of the robust multiple prior Bayesian approach in models with sign restrictions seems particularly worthwhile when the user prefers a frequentist interpretation of the estimates. They are no downsides to using this approach from a frequentist point of view. For Bayesian users, Giacomini and Kitagawa's approach is the only credible option when the identified set conditional on the reduced-form parameters is wide, as would typically be the case in partially identified models. In contrast, for more tightly restricted sign-identified models of the type commonly used in macroeconomics and energy economics, identified sets tend to be narrow under the conventional prior specification, so we can take advantage of the ability to compute Bayes estimates and to conduct joint inference. There is no compelling rationale for applying Giacomini and Kitagawa's approach in that setting.

## 3 Sign-identified models subject to narrative restrictions

Giacomini, Kitagawa and Read (2021a) extend the Bayesian analysis in Giacomini and Kitagawa (2021) to sign-identified models subject to narrative restrictions. Bayesian inference about narrative restrictions was introduced in Antolin-Diaz and Rubio-Ramirez (2018). Narrative restrictions (NRs) based on extraneous evidence play an important role in applied work in sharpening inference about impulse responses. NRs can take several forms in practice. For example, they may involve sign restrictions on structural shocks or restrictions on the relative magnitude of shocks in selected periods. Giacomini et al. also consider restricting historical decompositions in a similar fashion. They define the historical decomposition as the cumulative contribution of a sequence of realizations of a given structural shock between selected dates to the unexpected change in some variable over the same period. This is not the textbook definition of a historical decomposition, however, which refers to the cumulative contribution to date of all shocks of a given type to the change in some variable between selected dates (see Kilian and Lütkepohl 2017). Narrative restrictions on such historical decompositions are widely used in applied work (see, e.g., Zhou 2020). Although the analysis in Giacomini et al. presumably carries over to the latter type of NRs, it would have been useful to provide a comprehensive analysis.

Giacomini et al. make several distinct contributions. Their starting point is a critique of the conventional method of incorporating NRs in Antolin-Diaz and Rubio-Ramirez (2018). Their point is that, when the reduced-form parameters $A$ and $\Sigma$ are known, the likelihood conditional on the NRs will be maximized at the value of $Q$ that minimizes the ex ante probability that the NRs are satisfied. The posterior based on this conditional likelihood will therefore asymptotically place higher posterior probability on values of $Q$ that result in lower ex ante probability that the NRs are satisfied.

Giacomini et al. instead advocate constructing the likelihood without conditioning on the NR holding. For any value of the reduced-form parameters such that the data are compatible with the NR, there will be a set of values of $Q$ that satisfy the NR restrictions. The value of the unconditional likelihood will be the same for all values of $Q$ within this set. Thus, it may seem enough to simply replace the conditional likelihood by the unconditional likelihood when implementing the conventional method. Giacomini et al. argue that this would not be advisable. Their concern is that the conditional posterior of $Q$ will be proportionate to the conditional prior for $Q$ in these regions, so posterior impulse response inference may be sensitive to the choice of the conditional prior for $Q$, even asymptotically.

To address this concern, Giacomini, Kitagawa and Read propose a generalization of the method in Giacomini and Kitagawa (2021) that allows in addition for the presence of NRs. Unlike in Giacomini and Kitagawa (2021), the robust identified set in this setting is not a consistent estimate of the identified set and lacks a frequentist interpretation. However, the corresponding robust pointwise credible sets may be given a frequentist interpretation, when the numbers of NRs is small relative to the length of the estimation period. The proposed algorithm for constructing the unconditional likelihood mirrors Algorithm 2 in Antolin-Diaz and Rubio-Ramirez (2018), except that the importance sampler is skipped. Giacomini, Kitagawa and Read (2021b) make the case that this approach to frequentist inference is more reliable than estimating the structural VAR model based on narrative proxies capturing information about the sign of structural shocks on specific dates, at least when the number of NRs is small.

Giacomini et al.'s (2021a) empirical illustration is based on Uhlig's (2005) model of monetary policy shocks.[5] The central question for an applied user is what the benefits are of robust agnostic multiple prior inference in this setting compared to simply relying on the unconditional likelihood. The answer is not straightforward, because the algorithm proposed by Giacomini et al. conflates the potential benefits from reducing the sensitivity to the use of conventional priors with the potential cost of allowing for a wider range of $Q$ priors, not all of which may be reflect the researcher's prior, as discussed earlier. Put differently, when Giacomini et al. show that their robust pointwise credible sets are wider than the standard pointwise credible sets, we do not know whether this is because they saved us from using an unintentionally informative $Q$ prior or because they changed

[5] This choice is problematic. As formally demonstrated in Wolf (2020), the monetary policy shock is not identified in this model because the sign of the impact response of real output to a monetary policy shock is left unrestricted. It would be surprising, however, if the finding that robust pointwise credible sets tend to be wider than standard pointwise credible sets were not true more generally.

the researcher's implicit prior about the impulse response vector $\theta$.

As stressed by Giacomini et al., the problem of the $Q$ prior being unintentionally informative for the posterior in the context of NRs conceptually is exactly the same as in the case of the conventional prior for sign-identified models in Giacomini and Kitagawa (2021). In the latter case, Inoue and Kilian (2021a) showed by example that this problem tends to be negligible in many empirical applications because the substantive conclusions based on the posterior are not driven by the prior. By the same token, it is not obvious how serious the additional problem related to NRs identified in Giacomini et al. is in practice. Intuitively, one would expect this problem to be most severe in comparatively agnostic models such as models that are only partially identified in that only one of the structural shocks is identified. The more tightly identified the structural VAR model is, the smaller the identified impulse response set is likely to be, limiting the impact of the problem identified by Giacomini et al.

All of this, of course, remains speculation unless we have direct evidence on the practical significance of the problem identified in Giacomini et al. To address this point, I would encourage the authors to examine whether the substantive conclusions of studies in the literature that utilize NRs are driven by the use of the conventional uniform-Gaussian inverse Wishart prior or not. This could be accomplished by reexamining these models based on the unconditional likelihood and comparing the joint impulse posterior to the joint impulse response prior as in Inoue and Kilian (2021a). A case in point is the global oil market model of Kilian and Murphy (2014), as evaluated in Inoue and Kilian (2021a). This model combines narrative restrictions on the historical decompositions and sign restrictions on the impulse responses. Figure 1 compares the Bayes estimate of the impulse response vector $\theta$ based on the unconditional likelihood, as suggested by Giacomini et al., to the corresponding estimator under the conventional prior. Although there are some differences in the impulse response prior compared to the prior shown in Inoue and Kilian (2021a) based on the conditional likelihood, especially at short horizons, the posterior estimate of $\theta$ is effectively identical to that obtained following Antolin-Diaz and Rubio-Ramirez (2018). Moreover, the key substantive conclusion that the flow demand shock has large and persistent effects on global real activity and on the real price of oil clearly is not an artifact of the impulse response prior. Nor is the large and persistent effect of the storage demand shock on the real price of oil. Thus, we can say with confidence that none of the substantive conclusions based on this model are driven by the prior for $Q$, suggesting that in this case there would no reason to rely on the robust multiple prior approach of Giacomini et al.

## 4 Can we empirically assess the plausibility of identifying restrictions?

One argument in the paper that is attributed to Giacomini and Kitagawa (2021) is that we can judge the plausibility of the maintained identifying restrictions in a structural VAR model based on the fraction of posterior draws that are admissible in that they satisfy the identifying restrictions. That argument does not seem compelling. Suppose we have a set of economically motivated sign restrictions. Now we drop some of these sign restrictions, making the simplified model uncompelling from an economic point of view. Would this simplified model not by construction receive more "empirical support" than the correct model, as measured by the fraction of admissible models? In fact, the more informative the identification of the structural VAR model is and the tighter the identified set is as a result, the less "empirically plausible" it would seem by this metric. Thus, the procedure Giacomini et al. allude to seems designed to penalize the use of economically motivated identifying restrictions.

Similar concerns apply to Giacomini, Kitagawa and Read's claim that the plausibility of narrative restrictions may be judged based on the fraction of posterior draws that satisfy these restrictions. For example, they argue that the October 1979 restriction on the sign of the monetary policy shock is satisfied for every draw and, hence, should be considered highly plausible, whereas the extended set of narrative restrictions involving seven additional sign restrictions on the monetary policy shock on selected dates is regarded much less plausible with a fraction of only 53%. That argument seems odd. Clearly, when we have one NR nested in a larger set of NRs, holding constant the remaining identifying restrictions, the fraction of admissible models for the larger set of NRs cannot be higher than that for the single NR, which virtually ensures that the fraction of admissible models will decline with the addition of more NRs.[6]

A more sensible interpretation in my view is that both identifying restrictions on impulse responses and NRs are maintained assumptions. NRs ideally should be based on extraneous evidence, which makes them part of the data we condition on in fitting the structural model. There is no question that it is typically easier to fit the data without NRs. The point of imposing them is that they restrict the reduced-form parameter space of the model and hence the identified set of the impulse responses.

## 5 Conclusion

The authors deserve credit for clarifying the implications of widely used priors in that literature, in particular as they relate to narrative restrictions. Their work contributes to a better understanding of the merits and limitations of the conventional approach and provides tools that allow users to remain agnostic about parts of the model, albeit at the cost of remaining agnostic about the implicit impulse response prior. I highlighted several potential drawbacks of the proposed approach. One

[6] When imposing that the alternative NR that the shock to the federal funds rate in October 1979 was the largest policy shock in the estimation period, which seems eminently reasonable, the empirical plausibility drops to 17%, yet the impulse response functions remain so similar to the responses obtained based on the extended NR set that the authors do not bother to include the plot. It seems similarly odd for the empirical plausibility of two model specifications to differ substantially when they imply very similar estimates for the impulse responses we care about.

drawback is that the implied identified sets and robust credible sets may be too wide to allow substantive conclusions about the shape of response functions in practice. Another drawback is that inference currently is limited to individual impulse responses, which are of less interest in applied work than vectors of impulse responses. I also drew attention to the fact that the paper does not actually quantify the problem of unintentionally informative priors that it highlights. I showed by example that this problem may be negligible in practice, creating a trade-off between the robust approach proposed in this paper and conventional single prior approaches that not only make explicit the role of the prior and its impact on the posterior, but allow the construction of impulse response point estimates that can be readily interpreted.

**Additional references**


1. Antolin-Diaz, J., Rubio-Ramirez, J.F., 2018. Narrative sign restrictions for SVARs. American Economic Review 108, 2802-2839. https://doi.org/10.1257/aer.20161852

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

3. Baumeister, C., Hamilton, J.D., 2015. Sign restrictions, vector autoregressions, and useful prior information. Econometrica 83, 1963-1999. https://doi.org/10.3982/ECTA12356

4. Baumeister, C., Hamilton, J.D., 2019. Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks. American Economic Review 109, 1873-1910. https://doi.org/10.3982/ECTA12356

5. Baumeister, C., Hamilton, J.D., 2020. Drawing conclusions from structural vector autoregressions identified on the basis of sign restrictions. Journal of International Money and Finance,

forthcoming. https://doi.org/10.1016/j.jimonfin.2020.102250

6. Gafarov, B., Meier, M., Montiel Olea, J.L., 2018. Delta-method inference for a class of set-identified SVARs. Journal of Econometrics 203, 316-327. https://doi.org/10.1016/j.jeconom.2017.12.004

7. Giacomini, R., Kitagawa, T. 2021. Robust Bayesian inference for set-identified models. Econometrica 89, 1519-1556. https://doi.org/10.3982/ECTA16773

8. Giacomini, R., Kitagawa, T., Read, M., 2021a. Identification and estimation under narrative restrictions. Manuscript, University College London.

9. Giacomini, R., Kitagawa, T., Read, M., 2021b. Narrative restrictions and Proxies. Manuscript, University College London.

10. Giacomini, R., Kitagawa, T., Uhlig, H., 2019. Estimation under ambiguity. Manuscript, University College London.

11. Granziera, E., Moon, H.R., Schorfheide, F., 2018. Inference for VARs identified with sign restrictions. Quantitative Economics 9, 1087-1121. https://doi.org/10.3982/QE978

12. Inoue, A., Kilian, L., 2021a. The role of the prior in estimating VAR models with sign restrictions. Manuscript, Federal Reserve Bank of Dallas.

13. Inoue, A., Kilian, L., 2021b. Joint Bayesian inference about impulse responses in VAR models. Journal of Econometrics, forthcoming. https://doi.org/10.1016/j.jeconom.2021.05.010

14. Kilian, L. 2021. Understanding the estimation of oil demand and oil supply elasticities. Manuscript, Federal Reserve Bank of Dallas.

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. https://doi.org/10.1002/jae.2322

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

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

18. Uhlig, H., 2005. What are the effects of monetary policy on output? Results from an agnostic identification procedure. Journal of Monetary Economics 52, 381-419. https://doi.org/10.1016/j.jmoneco.200

19. 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. https://doi.org/10.1086/707179

20. Wolf, C. 2020. SVAR (mis-)identification and the real effects of monetary policy shocks. American Economic Journal: Macroeconomics 12, 1-32. https://doi.org/10.1257/mac.20180328

21. Zhou, X., 2020. Refining the Workhorse Oil Market Model. Journal of Applied Econometrics, 35, 130-140. https://doi.org/10.1002/jae.2743

**Figure 1: Estimates of Impulse Response Vector under Prior and under Posterior based on the Unconditional Likelihood in the Kilian and Murphy (2014) global oil market model**

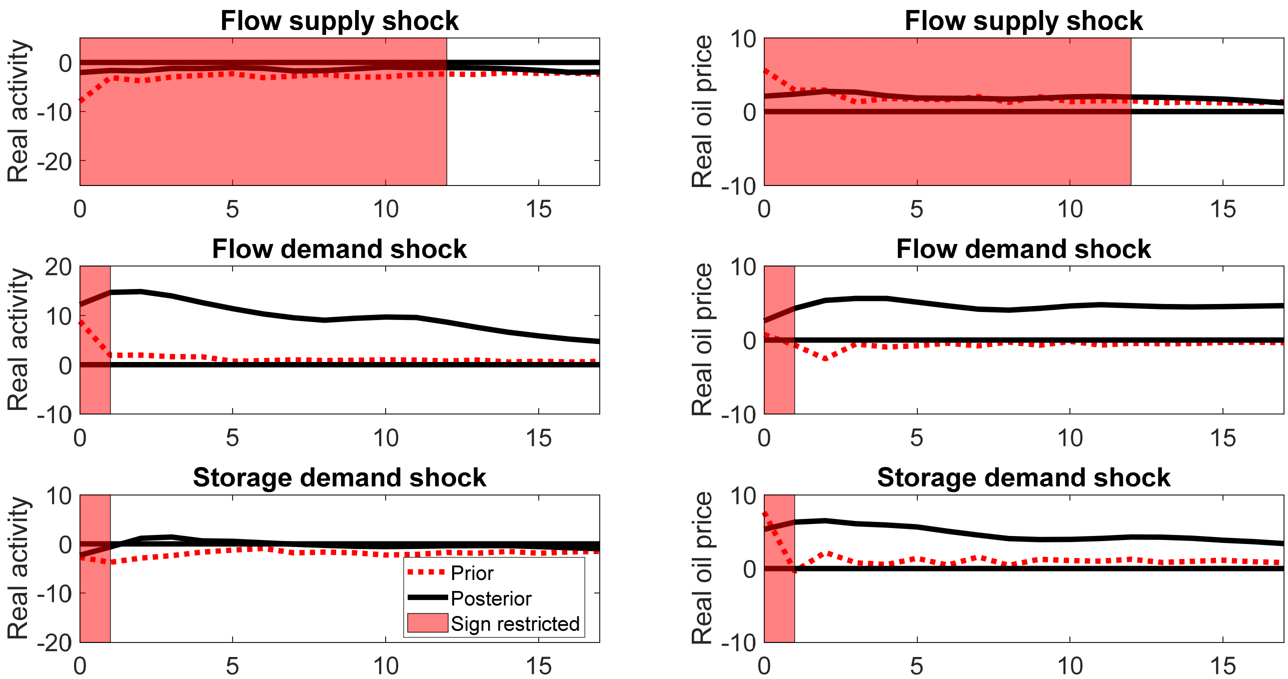


NOTES: Horizons at which sign restrictions are imposed are shown as shaded areas. The Bayes estimate under absolute loss, constructed as in Inoue and Kilian (2021b), is shown as solid lines. The corresponding estimate under the prior is shown as dotted lines.