Federal Reserve Bank of Dallas

Impulse Response Diagnostics for Priors on Parameters in Structural Vector Autoregressions

Lutz Kilian

Working Paper 2507 **February 2025**

Research Department

https://doi.org/10.24149/wp2507

Impulse Response Diagnostics for Priors on Parameters in Structural Vector Autoregressions*

Lutz Kilian†

February 15, 2025

Abstract

Structural impulse response functions may be estimated based on priors about the parameters of the structural VAR presentation. Even when such priors appear seemingly reasonable, they may imply an unintentionally informative prior for the structural impulse responses. Rather than pretending that the posterior of the impulse responses does not depend on this prior, the proposal in this paper is to verify that the prior distribution of the vector of impulse responses of interest is not unintentionally informative. Moreover, if the impulse response prior is intentionally informative, this point must be conveyed, so the reader can properly evaluate the reported conclusions. This paper discusses easy-to-use diagnostic tools that help practitioners address these concerns.

JEL: C11, C32, C52, Q43

Keywords: VAR, prior, posterior, impulse response, inference.

---

*The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas or the Federal Reserve System. The author did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

†Lutz Kilian, Research Department, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, TX 75201, USA. Email: lkilian2019@gmail.com.

1. Introduction

The conventional approach to estimating sign-identified vector autoregressive (VAR) models involves specifying a conjugate Gaussian-inverse Wishart prior for the parameters in the reduced-form VAR representation and a Haar prior for the orthogonal rotation matrix. Posterior inference in such models is justified if the identification is sufficiently tight (see Inoue and Kilian 2025). Baumeister and Hamilton (2015) proposed to replace this conventional prior by a prior on the parameters of the structural VAR representation that may embody sign restrictions as well as short- and long-run identifying restrictions. This alternative Bayesian approach has been employed in a number of recent studies including Baumeister and Hamilton (2019), Herrera and Rangaraju (2020), Casoli et al. (2024), and Aastveit et al. (2025).

Of particular importance are the priors for the parameters that govern the contemporaneous relationship between the model variables. Sometimes these priors may be motivated based on related theoretical work, but more typically they are specified in an ad hoc fashion. A case in point is economic models in which some of these parameters represent the product of -1 and the price elasticities of demand or supply or the income elasticity. For example, building on Baumeister and Hamilton (2019), many studies choose to represent the marginal prior distribution of the demand and supply elasticities as truncated Student-t distributions.[1]

Even when elasticity priors appear seemingly reasonable at first sight, they may imply a counterintuitive prior distribution on the structural impulse responses that the user is not aware of. The reason is that structural impulse responses can be written as a nonlinear function of the structural model parameters. This means that even a seemingly sensible prior on the structural

[1] It may seem that one should be able to derive a prior distribution for these parameters from extraneous empirical elasticity estimates. This is not the case. Not only are elasticity estimates defined in accordance with the model hard to find, but, even when a range of estimates is available, this tells us nothing about the functional form of the prior. As a result, researchers tend to rely on introspection when specifying the elasticity priors.

parameters may imply an unintentionally informative prior about the impulse responses. Rather than pretending that the posterior of the impulse responses does not depend on this prior, the suggestion in this paper is to evaluate the joint prior distribution of the vector of impulse responses of interest to make sure that the prior for the structural VAR model parameters is not unintentionally informative. Moreover, if the impulse response prior is intentionally informative about the sign, shape or magnitude of the response functions, this point must be conveyed, so the reader can properly evaluate the reported conclusions based on the posterior. The paper provides easy-to-use diagnostic tools that help practitioners evaluate the impulse response prior implied by the prior on the structural model parameters. The use of these tools is illustrated based on the widely cited global oil market model of Baumeister and Hamilton (2019). Several subsequent empirical applications have built on their model and prior specification, making this application of broader interest.

2. Priors on the Parameters in Structural VAR Models

The ultimate objective of fitting structural VAR models is typically not to estimate the structural model parameters, but to quantify how the model variables respond to structural shocks. This involves inverting the structural VAR representation to recover the structural impulse responses. Consider a model of order p for a $K \times 1$ vector of time series data, y_t, with mutually uncorrelated iid Gaussian innovations, w_t, and $K \times K$ dimensional matrices, $B_0, B_1, ..., B_p$, such that

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

Imposing explicit prior distributions on the parameters in $B_0, B_1, ..., B_p$ is not equivalent to specifying an explicit prior on the vector of structural impulse responses, which is defined by the nonlinear transformation $\theta = g(B_0, B_1, ..., B_p)$. Even if a prior on $B_0, B_1, ..., B_p$ (or, alternatively, on $B_0^{-1}, B_1, ..., B_p$) may be defended on economic grounds, the prior on θ obtained by applying the change-of-variable method may be unintentionally informative. In practice, this prior may be simulated by drawing from the prior distribution of the structural model parameters. A useful starting point is to report the central tendency of the prior for θ as the value of θ that minimizes in expectation the loss function of the user, as discussed in Inoue and Kilian (2022), before considering the distribution.

3. Empirical Example

The empirical example is a global oil market VAR model discussed in Baumeister and Hamilton (2019).[2] My analysis serves to make explicit the implications of Baumeister and Hamilton's prior on the structural model parameters that were not apparent or discussed in the original paper. The four structural shocks are an oil supply shock, an oil-market specific demand shock (referred to as the consumption shock), an oil inventory demand shock and a shock to global economic activity. For details on the data, loss function, prior, and model specification the reader is referred to the original source. In the interest of space, I focus on the responses of global real activity and of the real price of oil.

Figure 1 shows as solid lines the central tendency of the prior for θ implied by the baseline model in Baumeister and Hamilton (2019). The prior distributions are generated using the same Metropolis-Hastings algorithm as in the original study, except that the proposal density is centered on the prior mode. It is useful to consider each shock in isolation. First, the only shock that is a priori expected to substantially raise the real price of oil is the oil supply shock. In contrast, the response of the real price of oil to the demand shocks in the model is either quite

[2] The current paper is not concerned with the merits of this model, but uses it merely as an illustrative example. For further discussion of this model, its robustness and its limitations see, e.g., Herrera and Rangaraju (2020), Kilian (2022), and Braun (2023).

small, near zero, or even negative. An unexpected oil supply disruption causes a persistent increase in the real price of oil by 2.5% on impact and by 6.6% after 12 months. The corresponding decline in global real activity gradually approaches -1.8% after 12 months and reaches -3.1% after 18 months. This prior is consistent with the authors being strong believers in the importance of oil supply shocks, but was not made explicit in their paper.

Second, an unexpected increase in global economic activity is associated with a sustained increase in the level of real activity over time that reaches 3.8% after 18 months. An obvious question is how to reconcile the persistently positive and increasing response of economic activity over 18 months with the negligible or even negative response of the real price of oil. This pattern is clearly at odds with conventional views about the relationship between economic expansions and the real price of oil (e.g., Barsky and Kilian 2001; Hamilton 2009). The reason is that higher economic activity means higher demand for oil, which has to raise the real price of oil. Even if one believed economic activity shocks to be unimportant for the real price of oil, one would not expect a negative price response to a positive shock to economic activity. This feature of Baumeister and Hamilton's (2019) prior was clearly unintended.

Third, we know a priori that an oil-market specific shock to the demand for oil that raises the real price of oil, must lower global real activity. The response of the real price of oil to such a shock in Figure 1 is modestly positive on impact and becomes essentially zero after a few months. This is not an unreasonable prior view, except that it raises the question of how this shock at the same time raises global real activity by 0.7% after 18 months.

Fourth, it is not clear how to reconcile the response of the real price of oil to a storage demand shock with standard economic theory for storable commodities, which implies a jump in the price on impact followed by a gradual decline.

Thus, the central tendency of the joint impulse response prior implied by Baumeister and Hamilton's prior on the structural VAR model parameters is economically implausible in several dimensions. Figure 1 also shows as dotted lines the corresponding central tendency of the joint impulse response posterior. The data substantially revise the prior. For example, the response of global real activity to an unexpected oil supply disruption is revised up from -3.1% after 18 months to -0.1%, while the corresponding response of the real price of oil shrinks to 2.9%. There is little evidence of a strong recessionary impact on the global economy and the oil price response is more modest. Likewise, the responses to an economic activity shock now look more reasonable than under the prior. The response of the real price of oil peaks after 6 months at 5.1% and remains persistently high at longer horizons. The response of global real activity after 18 months is 1.8% rather than 3.8%. The responses to the consumption demand shock diminish with the puzzling increase of real activity under the prior overturned by the data. The response of the real price of oil to the inventory demand shock no longer grows over time.

This evidence raises the concern that the Bayes estimate of the impulse response vector in Figure 1 has been pulled in the direction of the central tendency of the prior. One may object that the impulse response prior in this example is so diffuse that the central tendency of the prior may be ignored, but how informative the prior is, is not only determined by the support of the prior, but also by how the probability mass is distributed over the support of the prior. For example, one has to wonder to what extent the positive response of the real price of oil to economic activity shocks is underestimated because the prior assigns considerable probability mass to negative values of this response.

While visualizing the extent to which the joint prior distribution is informative about the sign, shape, and magnitude of the response functions is not straightforward, some insights may be gained by reporting the marginal impulse response prior distributions. Figure 2 shows that there is a 60% prior probability of a negative response of the real price of oil to an unexpected increase in economic activity at horizon 4, so the implicit impulse response prior distribution is tilted toward negative values. At the same time, the response of the real price of oil to the oil supply shock shows an 89% probability of a positive response to a negative oil supply shock with substantial probability mass on large responses, indicating that the prior strongly favors large price effects from oil supply shocks at the expense of economic activity shocks. Thus, the key finding of that paper that oil demand shocks are less important for oil price fluctuations than previously thought and oil supply shocks are more important may be an artifact of the prior on the structural parameters, which assigns more probability mass to this outcome than previous studies.

Earlier research has shown that allowing for large values of the price elasticity of oil supply, as the Baumeister-Hamilton prior does, raises the relative importance of oil supply shocks for explaining oil price fluctuations (see Kilian and Murphy 2012). Likely, this problem could be avoided by modifying the initial parameter prior, for example, by reducing the upper bound on the price elasticity of oil supply until the central tendency of the implied impulse response prior becomes more reasonable. Indeed, related evidence in Herrera and Rangaraju (2020) and Braun (2023) suggests that this feature explains much of the difference between the posterior results in Baumeister and Hamilton (2019) and in earlier research.

Concluding remarks

This paper presented tools that help understand the implications of priors on structural VAR parameters for the impulse response prior. These tools allow users to detect and correct unintentionally informative priors and help make explicit prior views about the impulse responses that otherwise would not be readily apparent.

References

Aastveit, K.A., Bjørnland, H.C., Cross, J.L., 2025. Inflation expectations and the pass-through of oil prices. Review of Economics and Statistics 105, 733-743.

Barsky, R.B., Kilian, L., 2001. Do we really know that oil caused the Great Stagflation? A monetary alternative. NBER Macroeconomics Annual, 137-183.

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

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.

Braun, R., 2023.The importance of supply and demand for oil prices: Evidence from non-Gaussianity. Quantitative Economics 14, 1163-1198.

Casoli, C., Manera, M., Valenti, D., 2024. Energy shocks in the Euro area: Disentangling the pass-through from oil and gas prices to inflation. Journal of International Money and Finance, 147, 103154.

Hamilton, J.D. 2009. Causes and consequences of the oil shock of 2007-08. Brookings Papers on Economic Activity, 1, 215-259.

Herrera, A.M., Rangaraju, S.K., 2020. The effects of oil supply shocks on U.S. economic activity: What have we learned? Journal of Applied Econometrics 35, 141-159.

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

Inoue, A., Kilian, L., 2025. When is the use of Gaussian-inverse Wishart-Haar priors appropriate? Journal of Political Economy, forthcoming.

Kilian, L., 2022. Facts and fiction in oil market modeling. Energy Economics 110, 105973.

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

Figure 1: Central tendency of impulse response prior and posterior

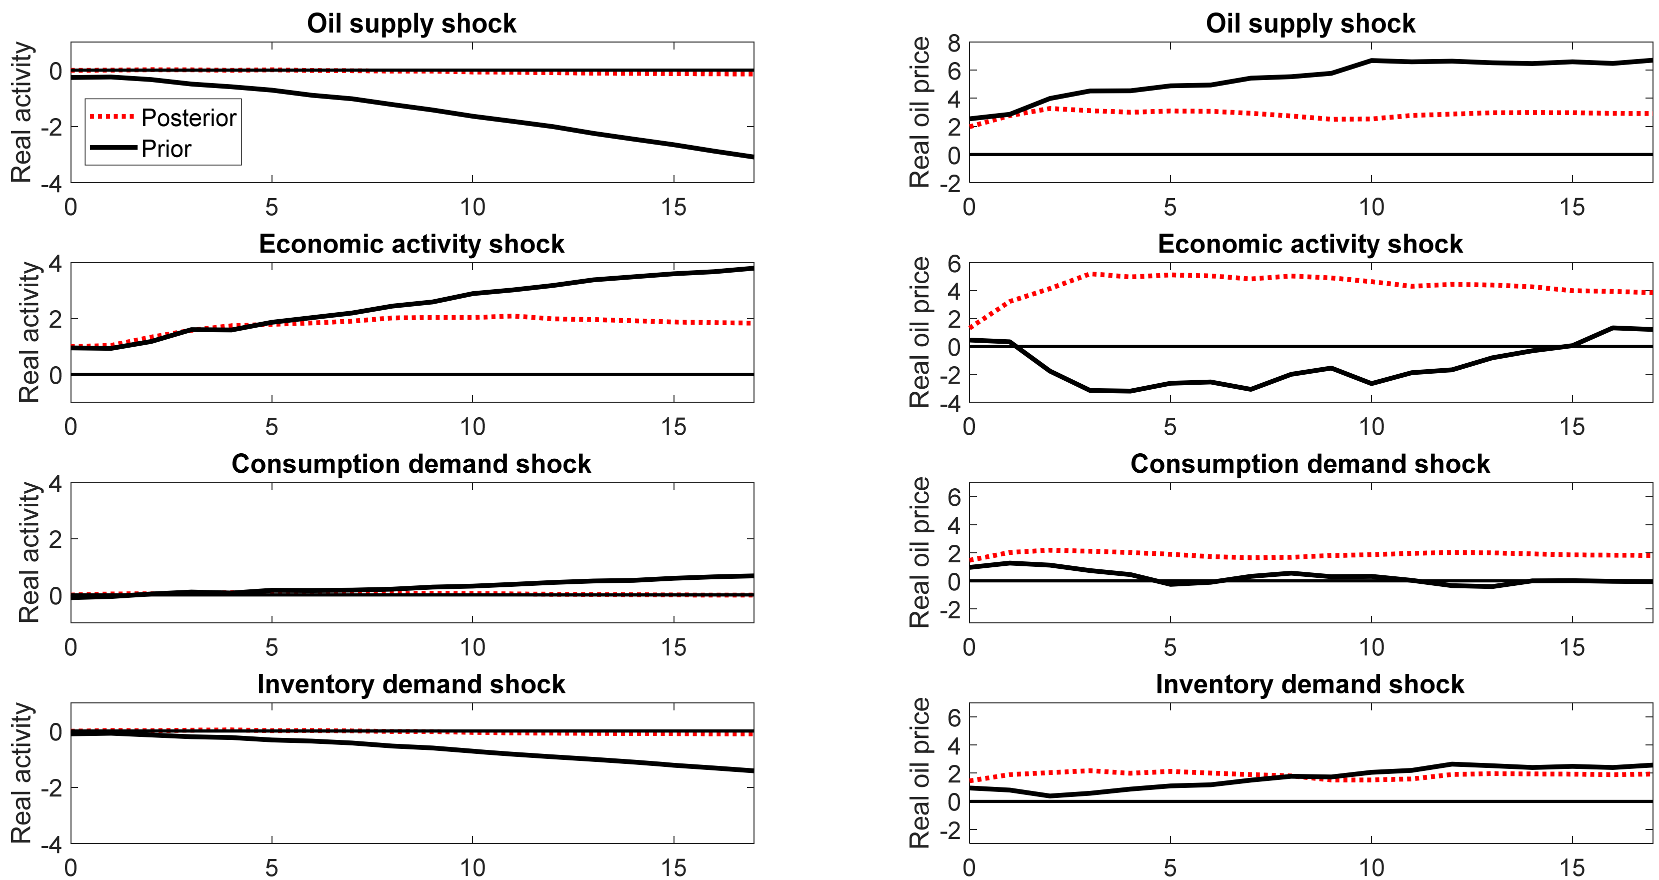

NOTES: Bayes estimate of θ under additively separable loss based on one million draws (dotted line) and the corresponding measure of central tendency under the prior (solid line) based on the baseline VAR model specification in Baumeister and Hamilton (2019).

Figure 2: Marginal prior distributions for response of real price of oil at horizon 4

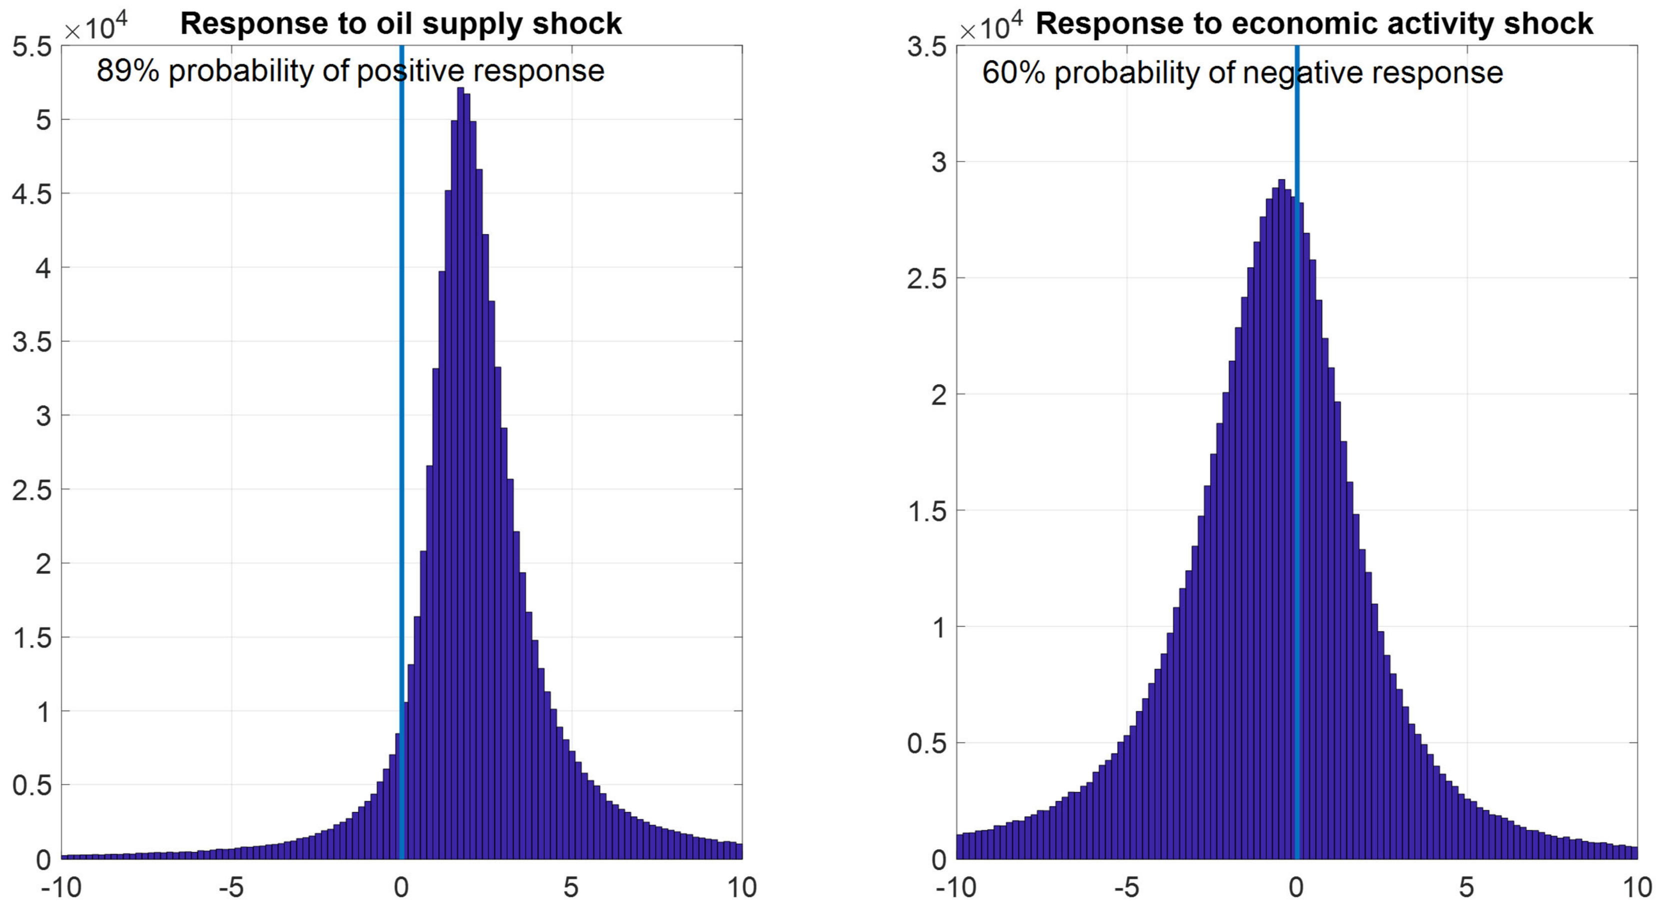

NOTES: Based on baseline VAR model specification in Baumeister and Hamilton (2019).