연세대학교 상경대학
**경제연구소**
Economic Research Institute
Yonsei University

서울시 서대문구 연세로 50
50 Yonsei-ro, Seodaemun-gu, Seoul, Korea
TEL: (+82-2) 2123-4065 FAX: (+82-2) 364-9149
E-mail: yeri4065@yonsei.ac.kr
http://yeri.yonsei.ac.kr/new

# An empirical assessment of the influence of informative rotation prior in the sign-identified SVAR model

**Hyeon-seung Huh**
*Yonsei University*

**David Kim**
*University of Sydney*

June 2025
2025RWP-246

**Economic Research Institute**

**Yonsei University**

# An empirical assessment of the influence of informative rotation prior in the sign-identified SVAR model

Hyeon-seung Huh
School of Economics
Yonsei University
Republic of Korea
hshuh@yonsei.ac.kr

David Kim
School of Economics
University of Sydney
Australia
Kim.david@sydney.edu.au

## Abstract

In the sign-identified Bayesian SVAR model, the standard setup usually postulates a Haar prior for the rotation matrix. However, the rotation matrix does not enter the likelihood, and its prior is never updated by data. A key implication is that the Haar prior rotation matrix can be unintentionally informative about posterior inference, despite having no relationship with economic interpretations or data. We show empirically how Haar prior rotation matrix could affect the results in the context of two well-known models: Baumeister and Hamilton (2018) and Peersman and Straub (2004, 2009). For both models, the histograms of accepted impact responses are shown to reflect closely the histograms of accepted rotation matrices. Although sampling uncertainty is updated by the data, it barely contributes to determining the set of accepted impact responses compared to the uncertainty about the rotation matrix, explaining why the histograms between the accepted impact responses and the accepted rotation matrices are similar in shape. To a lesser extent, the influence of the rotation matrix is carried over to subsequent responses where additional sampling uncertainty arises. Our results reinforce the argument that the rotation prior can affect the distribution of accepted responses, possibly leading to erroneous inferences.

*Key Words*: Structural vector autoregressions, Sign restrictions, Haar prior, Rotation matrix, Informativeness

*JEL Classification*: C32, C36, C51, E32, E52

## 1 Introduction

Using sign restrictions approach to identify shocks in structural VAR (SVAR) models has become popular, particularly as an alternative to the traditional identification schemes using parametric restrictions, such as short- or (and) long-run zero restrictions.[1] This approach involves rotating an initial set of orthogonal shocks in the SVAR to form a new set of orthogonal shocks from which impulse responses are obtained. These are retained if the impulse responses are in accordance with the sign restrictions and discarded otherwise. The procedure is repeated by taking another rotation of the initial orthogonal shocks from which another set of impulse responses is obtained and judged by the sign restrictions. Most sign-identified SVAR models adopt a Bayesian approach, and the standard setup usually postulates a normal-inverse Wishart prior for the reduced-form parameters and a Haar prior for the rotation matrix.

The Haar prior rotation matrix has a uniform distribution over the set of possible orthogonal matrices, implying that each of them is equally probable. For this reason, it is viewed as uninformative and arbitrary. However, elements in the matrix are not uniformly distributed, as demonstrated by studies including Shoemake (1995) and Baumeister and Hamilton (2015).[2] More importantly, the rotation matrix does not enter the likelihood and its prior is never updated by data, implying that the rotation prior can be unintentionally informative about posterior inference, despite having no bearing on economic interpretations or data (Moon and Schorfheide, 2012; Baumeister and Hamilton, 2015, 2018, 2020; Giacomini and Kitagawa, 2021). For example, Baumeister and Hamilton (2015, 2018) present empirical applications in which the impulse response posteriors are driven by the Haar prior rotation matrix and caution that the density of accepted responses, along with summary measures such as the median and other percentiles, may yield misleading conclusions.

In contrast, recent work by Inoue and Kilian (2020, 2024), Kilian (2022), and Arias *et al*. (2022) suggest that although the Haar prior may be unintentionally informative, this typically does not affect the impulse response posteriors. The uncertainty about the rotation matrix is dominated by the uncertainty about the reduced-form parameters, which is updated by the data. As such, the

[1] See Fry and Pagan (2011) and Kilian and Lütkepohl (2017) for a survey.
[2] Only exception is the three-variable model in which all elements in the Haar prior rotation matrix are uniformly distributed.

impulse response posteriors are driven largely by the data rather than the prior. Inoue and Kilian (2020, 2022) and Kilian (2022) argue that the quantitative importance of this observation greatly strengthens in models that employ many sign restrictions, as is often the case in applied work. They suggest that in such applications, concerns about the Haar prior for the rotation matrix have been much overstated, and alternative estimation methods are not typically required.

This paper revisits the informativeness of the Haar prior, which remains largely debatable. We take two particular studies proposing a way of examining the influence of the rotation prior in deriving impulse response posteriors. On the one hand, Inoue and Kilian (2022) compare the prior and posterior distributions of the accepted impulse responses. When both distributions differ substantively, they interpret this as evidence that the rotation prior is not influential for posterior inference. On the other hand, Giacomini *et al*. (2022) contend that comparing prior and posterior distributions does not address the issue because differences between the two can simply reflect the fact that the priors for the reduced-form parameters have been updated. They argue that one should assess whether changes in the prior for the rotation matrix affect the posterior, that is, posterior sensitivity to the choice of prior. However, Inoue and Kilian (2022) and Rubio-Ramírez (2022) maintain that the choice of the rotation prior is largely inconsequential for the set of accepted responses, since the posterior is dominated by the data in models with sign restrictions.

Given the contrasting views, we take a different route for assessing the influence of the Haar prior rotation matrix. By construction, uncertainty in the impulse response posterior is a combination of two sources: (1) sampling uncertainty associated with the reduced-form parameters and (2) identification uncertainty that comes purely from random draws for the rotation matrix. Our empirical assessment is to explore the relative contribution of the two sources of uncertainty. Unlike previous studies, we propose this task by employing as a metric the rotation matrix itself that generates the responses accepted by sign restrictions. Specifically, the density of accepted responses is compared with the density of the corresponding rotation matrices. If both are of a similar shape, the uncertainty in the accepted responses reflects, by implication, primarily the uncertainty in the rotation matrix, which is randomly drawn unrelated to the data. We show that this occurs mainly because the contribution of sampling uncertainty is scant and cannot be sufficient to overrule the influence of the rotation matrix for determining the set of accepted responses. Consequently, the rotation prior is informative for posterior inference.

For empirical exploration, we utilize two well-known models: a three-variable model of Baumeister and Hamilton (2018) and a five-variable model of Peersman and Straub (2004, 2009). In the former, all shocks are identified with a full set of sign restrictions in line with Inoue and Kilian (2022), suggesting the imposition of sign restrictions as tightly as possible to eschew the informativeness of the rotation matrix. In contrast, the latter adopts the so-called agnostic identification strategy proposed by Uhlig (2005), which imposes as few sign restrictions on the responses as necessary to identify shocks and remains agnostic about other responses, so as to be consistent with a wide range of theoretical models. While it is beyond the scope of this study to judge which approach is more prudent, we may also be able to gauge whether the influence of the rotation matrix can vary depending on whether the model is tightly identified by signs.

The rest of the paper is structured as follows. Section 2 provides a summary of the models in Baumeister and Hamilton (2018) and Peersman and Straub (2004, 2009). Section 3 briefly discusses key features of the Haar prior rotation matrix. Section 4 shows whether and how their findings in the two models have been influenced by the informativeness of the rotation matrix. Section 5 concludes.

## 2 Models and identification schemes

### 2-1. Baumeister and Hamilton model

Baumeister and Hamilton (2018, hereafter referred to as 'BH') consider a three-variable U.S. VAR model that consists of the output gap (Y), annual PCE deflator inflation (π), and the interest rate (R).[3] The output gap is defined as the log difference between real GDP and potential real GDP as estimated by the Congressional Budget Office. The dataset is quarterly over the sample period 1986:1 to 2008:3, and the VAR model is estimated with four lags in levels and a constant. Using the standard sign-restriction technique, BH identified only a monetary policy shock, but we extend to identify the other two shocks in the model based on the sign restrictions suggested by Fry and Pagan (2011). The sign restrictions are reproduced in Table 1. They identify the structural shocks as either an aggregate demand (AD), a cost-push (CP) or a monetary policy (MP) shock. The

[3] The data are available on the James Hamilton's website at https://econweb.ucsd.edu/~jhamilto/.

restrictions are the signs of impact responses to positive shocks, which can be found in a typical aggregate demand and supply diagram.

2-2. Peersman and Straub model

Peersman and Straub (2004 and 2009, hereafter 'PS') estimate a five-variable VAR model using the Euro Area data for the sample period 1982:Q1 to 2002:Q4. The variables are real GDP (Y), the GDP deflator (P), the interest rate (R), hours worked (H), and real wages (W). The VAR model is estimated in log levels, except the interest rate.[4] It also includes a constant and a time trend, and the number of lags is set at three. PS considered four structural shocks: a technology (TN) shock, a monetary policy (MP) shock, an aggregate demand (AD) shock, and a labor supply (LS) shock.[5] Table 2 reproduces the sign restrictions of PS that are used to identify the structural shocks. The sign restrictions are applied to the responses of the interest rate for the impact horizon and to the responses of all other variables for the impact and subsequent three horizons. Again, these sign restrictions were derived from a typical aggregate demand and aggregate supply diagram or well established DSGE models. PS adopt a minimum set of sign restrictions to be as agnostic as possible and avoid sign restrictions that are incompatible between the New Keynesian and Real Business Cycle models. In particular, the responses of hours worked are left unconstrained by sign restrictions, as the New Keynesian and Real Business Cycle models often suggest different implications, particularly in response to a technology shock.

2-3. A general model setup

Let $z_t$ be a vector of $n$ variables. Once a VAR model for $z_t$ is estimated, the reduced-form vector moving average representation may be expressed as[6]

$$z_t = D(L)e_t \tag{1}$$

where $D(L) = (I + D_1L + D_2L^2 + \ldots)$, $L$ is the lag operator, and $e_t$ is a vector of reduced-form errors, distributed as $N(0,\ \Omega)$, so that the errors have a mean of zero and are contemporaneously correlated. An initial set of orthogonal shocks is formed as $\varepsilon_t = P^{-1}e_t$ where $\Omega = PP'$. In this study, the matrix $P$ is obtained using a Cholesky decomposition of $\Omega$. Equation (1) is written as

[4] The data were obtained directly from Gert Peersman at Ghent University of Belgium.
[5] In their 2009 paper, they only identified a technology shock empirically but identified the remaining three shocks as well in its working paper version of 2004. We follow the latter and individually identify four underlying shocks in total.
[6] Deterministic terms, such as intercept and trend, are suppressed for brevity of illustration.

$$z_t = D(L)PP^{-1}e_t = D(L)P\varepsilon_t \tag{2}$$

The initial orthogonal shocks are rotated using the rotation matrix $Q$ of dimension ($n$x$n$), which has the property that $QQ' = Q'Q = I_n$ (i.e., orthonormal), to form a new set of orthogonal shocks. With this, equation (2) becomes:

$$z_t = D(L)PQQ'\varepsilon_t \tag{3}$$

where the new set of orthogonal shocks is $\eta_t = Q'\varepsilon_t = Q'P^{-1}e_t$, and the impulse responses to them are given by:

$$C(L) = D(L)PQ \tag{4}$$

and the impact response is:

$$C(0) = PQ \qquad (\because\ D(0) = I_n) \tag{5}$$

For a Bayesian extension, we adopt the standard setup, which postulates a normal-inverse Wishart prior for the reduced-form VAR parameters and a Haar prior for the rotation matrix $Q$. Based on this, simulating the model generates a large number of posterior responses for $C(L)$ in equation (4). The draw is retained if the set of responses satisfies the sign restrictions and discarded otherwise.

## 3 Rotation algorithms

When relying solely on sign restrictions for identification, the most common approach for generating the rotation matrix $Q$ is the algorithm proposed by Rubio-Ramírez *et al*. (2010). This method involves performing a QR decomposition on an ($n$x$n$) matrix $X$, whose elements are randomly drawn from a $N$(0, 1) distribution. From the QR decomposition of $X$=$QR$, $Q$ is an orthogonal matrix, and $R$ is an upper triangular matrix with its diagonal elements being positive to ensure uniqueness.[7] Panel (A) of Figure 1 depicts the empirical density distribution of a typical element in the $Q$ rotation matrix, when the model has the number of variables at $n$=2, 3, 4, 5, and 6, respectively, generated from 100,000 draws. Elements in the $Q$ rotation matrices exhibit values with more mass at the center when $n$>3, and with more mass at the tails when $n$=2. That is, some values are generated more frequently than others. The distribution is uniform only in the case of $n$=3. For each case of $n$, all remaining elements in the $Q$ rotation matrix have density distributions

[7] An alternative algorithm for generating the rotation matrix $Q$ is due to Fisher and Huh (2020) which utilizes Givens rotations. It produces equivalent results to the algorithm of Rubio-Ramírez *et al.* (2010).

of identical shape to the typical one in the figure. In the case of $n$=3, for example, all nine elements have the same uniform distribution. As the number of variables increases from $n$=6, the distributions continue to be more peaked at the center (not shown).

Both BH and PS adopt the $Q$ rotation matrix as a prior for the respective models of $n$=3 and $n$=5. We now examine the influence of this Haar prior rotation matrix on the posterior responses. In doing so, we focus on the impact responses of the variables, i.e., $C(0) = PQ$ in equation (5). The main reason is that responses at all horizons depend on the impact responses by the relationship of $C(j) = D(j)C(0)$ from equations (4) and (5), where $C(j)$ and $D(j)$ represent the responses to the structural shocks $\eta_t$ and the reduced-form errors $e_t$, respectively, at horizon $j$. By construction, the randomness of $C(0)$ comes from a combination of two sources: (1) sampling uncertainty in $P$ associated with the reduced-form parameters and (2) identification uncertainty that comes purely from a random rotation matrix $Q$.

We first look at the complete set of candidate impact responses, $C(0)$, that were generated prior to the application of sign restrictions to decide whether a set of responses should be accepted or not. Panels (B) and (C) of Figure 1 depicts a typical histogram of the candidate impact responses for the BH and PS models, respectively. In each model, the histogram of the candidate impact responses closely mimics in shape that of the element in the rotation matrix (i.e., $n$=3 for BH and $n$=5 for PS) in Panel (A).[8] Because the shocks were not identified economically in the absence of sign restrictions, all other candidate impact responses exhibit histograms that are of the same as the typical histogram in Panels (B) and (C), respectively, for the BH and PS models.[9] Consequently, their histograms closely resemble the corresponding histogram of the rotation matrices in Panel (A) across the models. This indicates that the candidate impact responses depend primarily on the $Q$ rotation matrix, and not on the matrix $P$, which is updated by the data. The next issue is how the imposition of sign restrictions can affect the influence of the rotation matrix in determining the set

[8] Note that elements in the rotation matrix range from -1 to 1 by definition.

[9] While all candidate impact responses exhibit histograms of the same shape, the location and width can differ depending on the variable. For example, the histograms of candidate impact responses for variable A are identical across any of the (unidentified) shocks, and the histograms of candidate impact responses for variable B are the same in shape as those for variable A, except that the location and width of the candidate impact responses can differ between the two variables. This is because the reduced-form errors associated with variables have different sizes. A complete plot of the candidate impact responses is available upon request. For a formal proof, see Baumeister and Hamilton (2015).

of accepted responses.

## 4 The role of sign restrictions

4-1. Baumeister and Hamilton (BH) model

The left panel of Figure 2 displays the median (50th percentile) of the accepted responses to the shocks (solid line), together with the 16th and 84th percentile bands of the accepted responses (shaded). The responses align with the predictions of a standard aggregate demand and supply model. In Figure 3, Panel (A) reports the histograms of *all* accepted responses on impact (shaded), whereas Panel (B) shows the histograms of the elements in the rotation matrix that generated the responses accepted by the sign restrictions (shaded). Comparing the two panels, one can see that the histograms of the impact responses closely mirror the histograms of the respective elements in the rotation matrix. That is, the impact responses of output gap to the three structural shocks exhibit histograms of similar shape to elements (1,1), (1,2), and (1,3) in the rotation matrix, respectively. The same holds for the impact responses of inflation and the interest rate, which share similar histograms with elements in the second and third rows of the rotation matrix, respectively. As such, the impact responses reflect primarily the rotation matrix, which was randomly drawn with no relation to the data. This echoes the critique by Baumeister and Hamilton (2015, 2018, 2020) regarding the influence of the rotation prior in deriving the set of accepted responses. From the relationship of $C(0) = PQ$, the similarity between the impact response and the rotation matrix also indicates that sampling uncertainty may have played only a limited role, the point we will revisit with more details.

The influence of the rotation matrix can also be revealed using the approach of Baumeister and Hamilton (2018, 2020), which is to analyze the relative contribution of the two sources of uncertainty in the impact response, $C(0) = PQ$, by shutting down the sampling uncertainty in $P$. We conduct a similar exercise by fixing the reduced-form VAR parameters in (1) at their OLS estimates for every random draw of the parameter values. Shutting down the sampling uncertainty this way leaves the rotation matrix as the sole generator of uncertainty. With some abuse of terminology, we refer to this exercise as non-Bayesian setup, as opposed to the Bayesian setup that takes account of both sampling and identification uncertainties. Panels (A) and (B) of Figure 3 report the results generated from the non-Bayesian setup (by the solid line). Since the sampling uncertainty was removed, no discernible difference in shape is found between the histograms of the impact responses and those of the respective elements in the rotation matrix. Importantly, all results are indistinguishable between the non-Bayesian and Bayesian setups. Panel (A) displays that the histograms of the impact responses under the non-Bayesian setup are similar to those of the impact responses under the Bayesian setup. Panel (B) produces the rotation matrices, the shapes of which are also indistinguishable across the two setups. By construction, our Bayesian and non-Bayesian setups are only distinguished by whether sampling uncertainty is permitted. Thus, the similarity of findings between the two setups implies that sampling uncertainty in the Bayesian setup plays little role compared to the uncertainty about the rotation matrix. Since the non-Bayesian setup assumes no sampling uncertainty, its impact responses are solely determined by the rotation matrix. Due to the lack of sampling uncertainty, the impact responses under the Bayesian setup are also driven primarily by the rotation matrix. Hence it is supposedly the similar rotation matrices between the non-Bayesian and Bayesian setups that produce the similar histograms for the impact responses as well.

Another point of interest from Panel (B) of Figure 3 is that the histograms of the rotation matrices are congruous with the sign restrictions, except possibly at the zero boundary. Looking at the first column, the histograms have values mostly positive satisfying the sign restrictions for the AD shock. Similarly, the values of the histograms in the second and third columns are mostly consistent with the sign restrictions for the CP and MP shocks, respectively. In light of this observation, we undertake an experiment with a hypothetical case in which the sign restrictions are imposed on the rotation matrix, rather than on the impact response. Under this identification scheme, a draw is accepted when its rotation matrix exhibits elements (1,1), (2,1), and (3,1) consistent, respectively, with the three sign restrictions for the AD shock, elements (1,2), (2,2), and (3,2) for the CP shock, and elements (1,3), (2,3), and (3,3) for the MP shock. This randomly drawn rotation matrix with the sign restrictions imposed on it is purely an artifact and conveys no relationship with economic interpretations or data.

Figure 4 reports the resulting histograms for the accepted rotation matrices and the corresponding impact responses. Since the sign restrictions have been imposed on the rotation matrix, Panel (B) shows that the accepted rotation matrices are exactly the same between the non-Bayesian and Bayesian setups. Panel (A) shows that the two setups exhibit similar histograms for the impact responses as well. As before, the impact responses under the Bayesian setup are determined mainly by the rotation matrix, while the contribution of sampling uncertainty is inconsequential. More striking is that the histograms of the impact responses are mostly consistent with the sign restrictions, even if no restrictions were imposed on the responses. In fact, they are similar to the histograms of the impact responses in Panel (A) of Figure 3, where the sign restrictions were imposed on the impact responses. A comparison of the Panels (B) of Figures 3 and 4 also shows that the rotation matrices yield similar histograms across the two identification schemes. This demonstrates that imposing sign restrictions on the impact responses is practically tantamount to imposing those same restrictions directly on the rotation matrices. A key implication is that the histograms of the impact responses from the former scheme resemble those of the respective elements in the rotation matrix from the latter scheme. The rotation matrices in the latter scheme were randomly drawn and only constrained by the sign restrictions in an arbitrary manner. Since they convey no relationship with economic interpretations or data, so do the impact responses in the former scheme. This presents strong evidence supporting the proposition that the rotation prior wields influence over the set of accepted responses. Again, the underlying reason is that sampling uncertainty is, at best, of second order influence relative to the uncertainty about the rotation prior. Due to the inherent weakness of sampling uncertainty, the impact responses and the rotation matrices show little variation irrespective of whether the sign restrictions are imposed on the rotation matrices or the impact responses. Furthermore, since the impact responses are similar across the two identification schemes, the same would be true for the responses at all other horizons by the relationship $C(j) = D(j)C(0)$ from equations (4) and (5). Figure 2 illustrates this point, as the responses from the scheme imposing the sign restriction on the impact responses in Panel (A) are almost identical to the responses from the scheme imposing the sign restrictions on the rotation matrices in Panel (B) over the horizons and across the shocks.

Finally, we examine whether the influence of the rotation matrix is carried over to the responses at horizons other than the impact one in the BH model. As discussed before, the responses at other horizons depend on the impact responses from the relationship of $C(j) = D(j)C(0)$. However, the relationship also indicates that there comes an additional source of sampling uncertainty associated with the lagged reduced-form parameters of $D(j)$, for $j > 0$. To gauge how sampling and identification uncertainties fare, Figure 5 displays the histograms of the accepted responses at the horizon of one, i.e., $C(1) = D(1)C(0)$. Compared with the case of impact responses, the histograms yield some features that are different between the Bayesian and non-Bayesian setups, illustrating the presence of additional sampling uncertainty. Nonetheless, the rotation matrix reported in Panel (B) of Figure 3 continues to exert influence in several responses: namely, the responses of output gap to AD, CP, and MP shocks, and the responses of the interest rate to CP and MP shocks. This suggests that sampling uncertainty at the horizon of one may still not be sufficient enough to overrule the influence of the rotation matrix in determining the accepted responses.

### 4-2. Peersman and Straub (PS) model

The model of Peersman and Straub (2004, 2009) is replicated, and Figure 6 reports the median (50th percentile) of the accepted responses to the shocks, together with the 16th and 84th percentile bands. For economic interpretations of these responses, we refer the reader to the original papers to save space. Moving to the main issue, examining the influence of the rotation matrix is carried out in a similar fashion as before. In Figure 7, Panel (A) shows the histograms of all accepted responses on impact (shaded), whereas Panel (B) reports the histograms of elements in the rotation matrix that generated the responses accepted by the sign restrictions (shaded). Most cases indicate that the histograms of the impact responses resemble the histograms of the respective elements in the rotation matrix. Again, the impact responses reflect primarily the rotation matrix. This result is particularly interesting because the PS model imposes sign restrictions on the impact and subsequent three horizons for output, prices, hours worked, and wages simultaneously. That is, the rotation matrices were retained only when the responses of these variables also satisfy the sign restrictions over the three subsequent horizons, in which additional sampling uncertainty arises. This could act to weaken the correspondence between the impact responses and the rotation matrices. For the PS model, it appears that the influence of the rotation matrix is strongly present in the impact responses for all the variables.

The figure also depicts the results under the non-Bayesian setup (i.e., sampling uncertainty being muted). As anticipated, the histograms of the impact responses and the histograms of the respective elements in the rotation matrix exhibit similar shapes. Panels (A) and (B) also illustrate that they are similar to the corresponding histograms of the impact responses and the rotation matrices under the Bayesian setup. The similarity of findings between the two setups suggests that sampling uncertainty in the Bayesian setup contributes little relative to the uncertainty about the rotation matrix in determining the accepted impact responses. Due to this apparent weakness of sampling uncertainty, the impact responses under the Bayesian setup are driven primarily by the rotation matrix. Overall, the implications are virtually the same as what we have found for the BH model, even though the sign restrictions were imposed beyond the impact horizon.

Panel (B) of the figure also indicates that the histograms of the rotation matrices mostly conform the sign restrictions, except possibly at the zero boundary. As before, we undertake an illustrative experiment, which applies the same set of sign restrictions to the rotation matrix instead. Figure 8 reports the resulting histograms for the accepted rotation matrices and the corresponding impact responses. With the sign restrictions imposed on the rotation matrix, the accepted rotation matrices should be identical between the Bayesian and non-Bayesian setups, as displayed in Panel (B). Panel (A) also shows that the two setups produce very similar histograms for the impact responses. This replicates the result under the Bayesian setup that there is little sampling uncertainty relative to the uncertainty about the rotation prior. Furthermore, the histograms of the impact responses are mostly consistent with the sign restrictions, although no sign restrictions were imposed on them. In fact, most histograms for the impact responses are similar to those of the impact responses in Figure 7, which were obtained from imposing the sign restrictions on the impact and three subsequent responses. There are only a small number of exceptions with virtually none for the responses of the interest rate, hours worked, and wages. For the Bayesian setup, it is the weakness of sampling uncertainty that renders the impact responses similar between the scheme imposing sign restrictions on the rotation matrices and the scheme imposing sign restrictions on the responses. Further, there are many cases that the impact responses in the latter scheme (Panel (A) of Figure 7) exhibit histograms that are similar in shape to the histograms of the rotation matrices in the former scheme (Panel (B) of Figure 8). For these cases, the impact responses reflect the rotation matrices which bear no relationship with economic interpretations or data by construction. The effects are pronounced in the responses of the interest rate, hours worked, and wages. The evidence again consolidates that the rotation prior is influential for the set of accepted responses. This is robust to the difference that one scheme imposes sign restrictions on the rotation matrices, while the other scheme imposes sign restrictions on the responses over several horizons.

Based on the relationship of $C(j) = D(j)C(0)$ in equations (4) and (5), we also gauge whether the influence of the rotation matrix is carried over to the responses at the horizon of one, $C(1) = D(1)C(0)$, where an additional sampling uncertainty arises from the lagged reduced-form parameters of $D(1)$. In the PS model, the rotation matrices were retained when the responses of output, prices, hours worked, and wages also fulfil the sign restrictions over the subsequent three horizons. Hence, there is an additional channel that the rotation matrix could be informative for subsequent responses. Figure 9 displays the histograms of the accepted responses at the horizon of one. Compared to the case of the impact responses, they exhibit more disparities between the non-Bayesian and Bayesian setups, reflecting additional sampling uncertainty. Yet, the rotation matrix reported in Panel (B) of Figure 7 still exerts influence in many responses. There is a relatively strong similarity between the histograms of the responses at the horizon of one and the histograms of the respective elements in the rotation matrix except for a few responses. The PS model also suggests that sampling uncertainty is not significant enough to overrule the influence of the rotation matrix in determining the set of accepted impulse responses.

## 5 Conclusion

While the sign-identified SVAR model has gained increasing popularity, it only achieves set identification, posing new methodological challenges. A key methodological issue is that since the rotation matrix does not enter the likelihood, and its priors are never updated by data, its influence does not vanish asymptotically. Thus, the rotation prior can be unintentionally informative about posterior inference, despite having no relationship with economic interpretations or data.

The present paper empirically explores the influence of the rotation prior in deriving impulse response posteriors, which remains contentious in the literature. By construction, the uncertainty in the impulse response posterior is a combination of two sources: sampling uncertainty, which is updated by the data, and uncertainty about the rotation matrix, which is not. Our empirical assessment is to expose the relative contribution of the two sources of uncertainty by using the rotation matrices that generate the impulse responses accepted by sign restrictions. Another metric in use is how the results change in the absence of sampling uncertainty. We took two well-known models for the empirical assessment: a three-variable model of Baumeister and Hamilton (2018) and a five-variable model of Peersman and Straub (2004, 2009).

For both models, the histograms of the accepted impact responses closely mirror those of the corresponding elements in the rotation matrix. This indicates that the accepted impact responses are driven primarily by the arbitrary rotation matrix. Although sampling uncertainty is updated by the data, its contribution was minimal and insufficient to outweigh the influence of the rotation prior in deriving the set of accepted impulse responses. This weakness of sampling uncertainty allows the rotation prior to dominate, resulting in histograms of the accepted responses that closely resemble those of the rotation matrices. Moreover, the influence of the rotation prior extends beyond impact responses, affecting responses at other horizons where additional sampling uncertainty arises. Our empirical results reinforce the critique concerning the undue influence of the rotation prior. This finding holds irrespective of whether the model is tightly identified using sign restrictions as in Baumeister and Hamilton (2018), or more loosely identified with a minimal set of restrictions following an agnostic approach as in Peersman and Straub (2004, 2009). Our experiments suggest that sign restrictions alone may be too weak to mitigate the influence of the rotation prior on the resulting impulse responses.

Several approaches have been proposed to address the issue of an informative rotation prior, with three drawing particular attention. First, Baumeister and Hamilton (2015, 2018) suggest acknowledging the informativeness of priors in sign-identified models and making this information content explicit. For example, they specify priors directly on the structural coefficients of the SVAR model in the form of a truncated *t*-distribution, while dispensing with the rotation matrix. Second, Giacomini and Kitagawa (2021) propose introducing ambiguity in the rotation matrix by replacing the single prior for the rotation matrix with a class of all priors that satisfy the identifying restrictions. This multiple-prior approach eventually makes the posterior inference free from the choice of the rotation matrix and reliant on a prior for the reduced-form parameters of the VAR models and on a set of identifying restrictions. The final approach calls for additional identifying information beyond sign restrictions, which may take the form of narrative restrictions, external instruments, elasticity bounds, selective exclusion restrictions, or shape restrictions. According to Inoue and Kilian (2024), the influence of the rotation prior tends to be negligible in models that are tightly identified by sign restrictions in conjunction with additional identifying information. However, it remains an ongoing issue how these three approaches can be effectively developed into a practical framework to complement and improve conventional sign-identified VAR models.

## References


Arias, J., J. Rubio-Ramírez, D. Waggoner, 2022, Uniform priors for impulse responses, Federal Reserve Bank of Philadelphia Working Paper #22-30.

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

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

Baumeister, C., J. Hamilton, 2020, Drawing conclusions from structural vector autoregressions identified on the basis of sign restrictions, *Journal of International Money and Finance* 109, #102250.

Fisher, L., H-S. Huh, 2020, Combining sign and parametric restrictions in SVARs by utilising Givens rotations, *Studies in Nonlinear Dynamics & Econometrics* 24, 1-19.

Fry, R., A. Pagan, 2011, Sign restrictions in structural vector autoregressions: A critical review, *Journal of Economic Literature* 49, 938-960.

Giacomini R., T. Kitagawa, 2021, Robust Bayesian inference for set-identified models, *Econometrica* 89(4), 1519-1556.

Giacomini, R. T. Kitagawa, M. Read, 2022, Narrative restrictions and proxies: Rejoinder, *Journal of Business & Economic Statistics* 40, 1438-1441.

Inoue, A., L. Kilian, 2020, The role of the prior in estimating VAR models with sign restrictions, Federal Reserve Bank of Dallas, Working Paper No. 2030.

Inoue, A., L. Kilian, 2024, When Is the use of Gaussian-inverse Wishart-Haar priors Appropriate?, Federal Reserve Bank of Dallas, Working Paper No. 2404.

Kilian, L., 2022, Comment on Giacomini, Kitagawa, and Read's "Narrative restrictions and proxies", *Journal of Business & Economic Statistics*, 40:4, 1429-1433,

Kilian, L., H. Lütkepohl, 2017, *Structural Vector Autoregressive Analysis*, Cambridge University Press, Cambridge, United Kingdom.

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

Peersman G., R. Straub, 2004, Technology shocks and robust sign restrictions in a EURO area SVAR, European Central Bank Working Paper #373.

Peersman G., R. Straub, 2009, Technology shocks and robust sign restrictions in a EURO area SVAR, *International Economic Review* 50, 727-750.

Rubio-Ramírez, J., 2022, Comments on “Narrative restrictions and proxies” by Giacomini, Kitagawa, and Read, *Journal of Business & Economic Statistics* 40, 1426-1428.

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

Shoemake, K., 1995, Uniform random rotations, in *Graphics Gems III*, Kirk, D. (ed.), Academic Press, 124-132.

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.

Table 1: Sign restrictions for Baumeister and Hamilton (2018)

| Shock\Variable | Output gap (Y) | Inflation (π) | Interest rate (R) |
|---|---|---|---|
| AD | ≥0 | ≥0 | ≥0 |
| CP | ≤0 | ≥0 | ≥0 |
| MP | ≤0 | ≤0 | ≥0 |

*Notes*: AD denotes an aggregate demand shock, CP a cost-push shock, and MP a monetary policy shock. The designation '≥0' denotes a non-negative response whereby the variable does not fall in response to the shock, while '≤0' denotes a non-positive response such that the variable does not rise in response to the shock. The sign restrictions are applied to the responses of the variables for the impact horizon.

Table 2: Sign restrictions for Peersman and Straub (2004, 2009)

| Shock\Variable | Output (Y) | Prices (P) | Interest rate (R) | Hours (H) | Wages (W) |
|---|---|---|---|---|---|
| TN | ≥0 | ≤0 | | | ≥0 |
| LS | ≥0 | ≤0 | | | ≤0 |
| AD | ≥0 | ≥0 | ≥0 | | |
| MP | ≤0 | ≤0 | ≥0 | | |

*Notes*: TN denotes a technology shock, LS a labor supply shock, AD an aggregate demand shock, and MP a monetary policy shock. The designation '≥0' denotes a non-negative response such that the variable does not fall in response to the shock, while '≤0' denotes a non-positive response such that the variable does not rise in response to the shock. The blank denotes an unrestricted response. The sign restrictions are applied to the responses of the interest rate for the impact horizon and to the responses of all other variables for the impact and subsequent three horizons.

Figure 1: The rotation and impact response matrices in the absence of sign restrictions

(A) Density distribution of a typical element in the rotation matrix

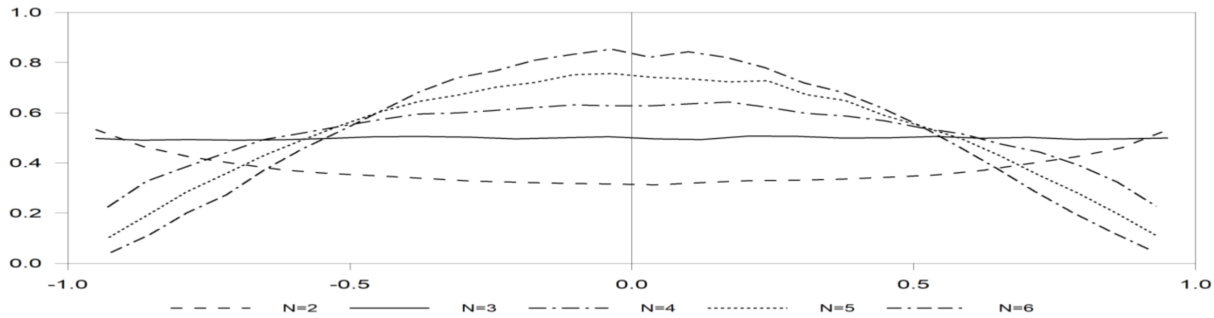


(B) Histogram of a typical element in the impact response matrix for the BH model

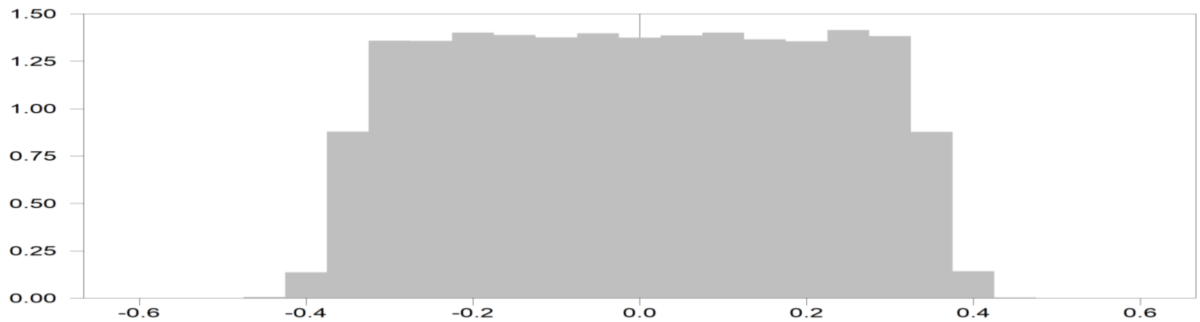


(C) Histogram of a typical element in the impact response matrix for the PS model

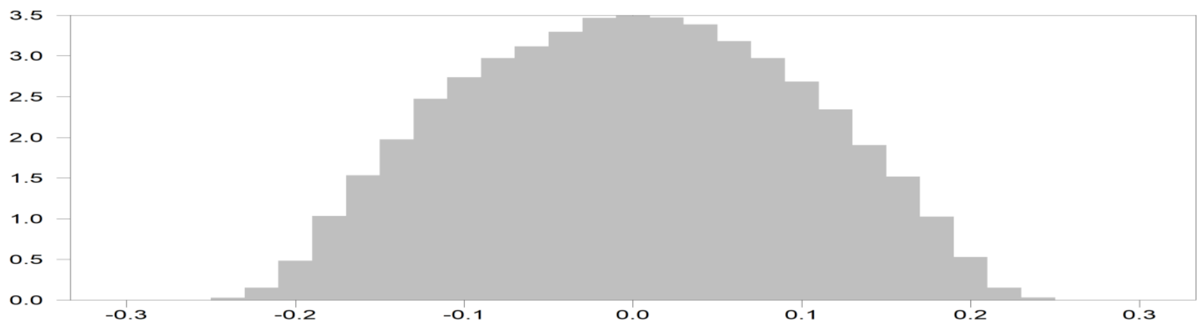

Figure 2: Responses of the variables to the shocks for the BH model

(A) Sign restrictions on the responses (B) Sign restrictions on the rotation matrix

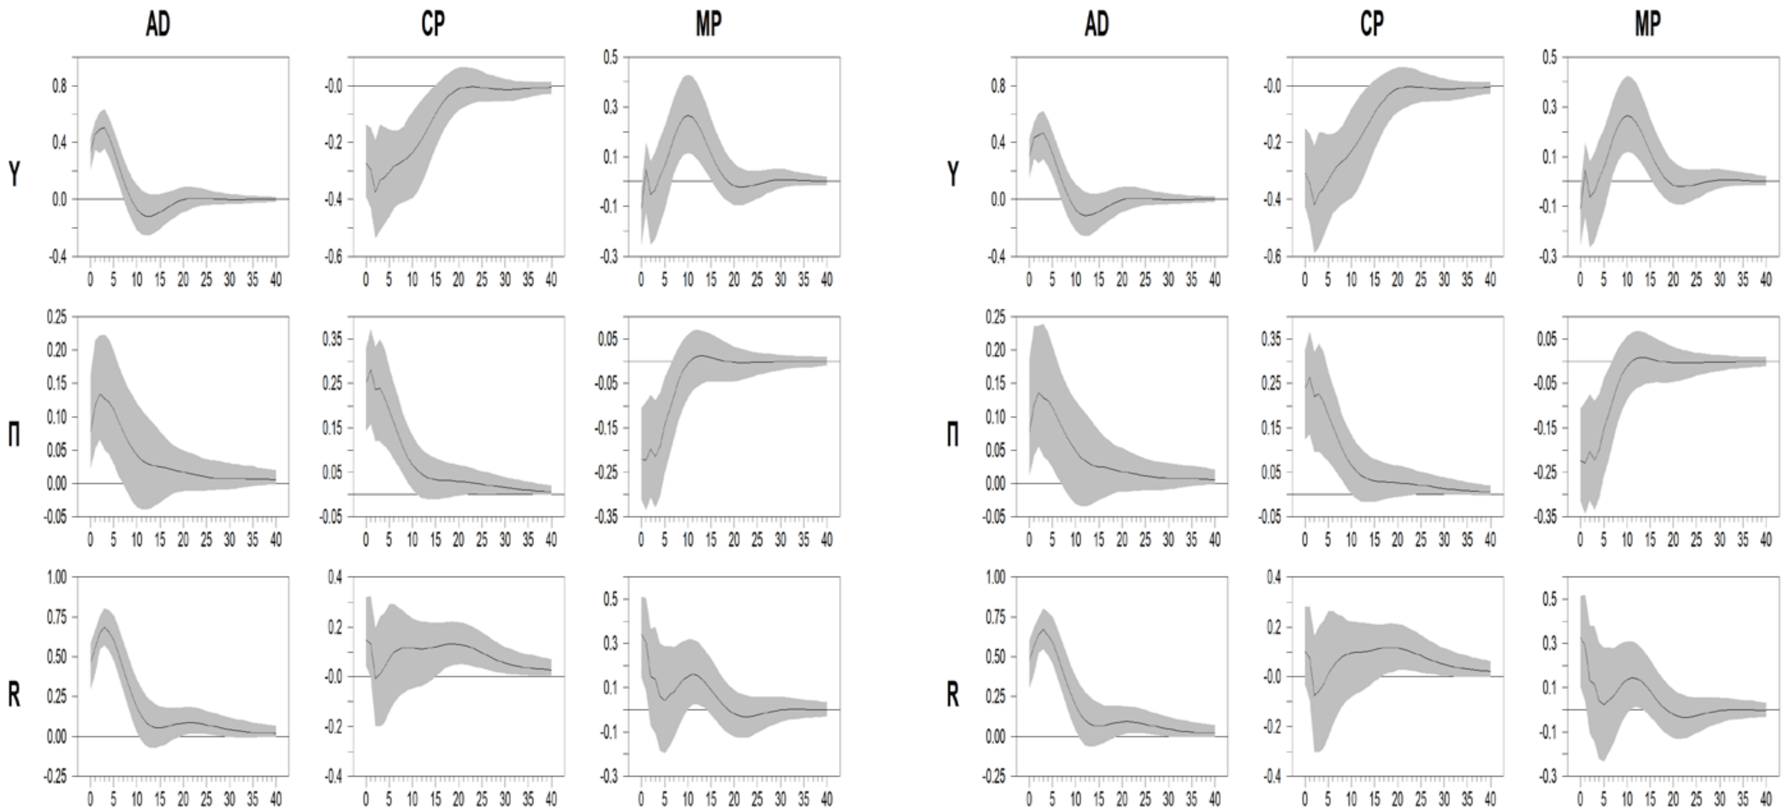


Note: The solid line and the shading are the median responses and the 16%–84% band responses, respectively.

Figure 3: Histograms of the impact responses and elements of accepted rotation matrices for the BH model

(A) Impact response matrix

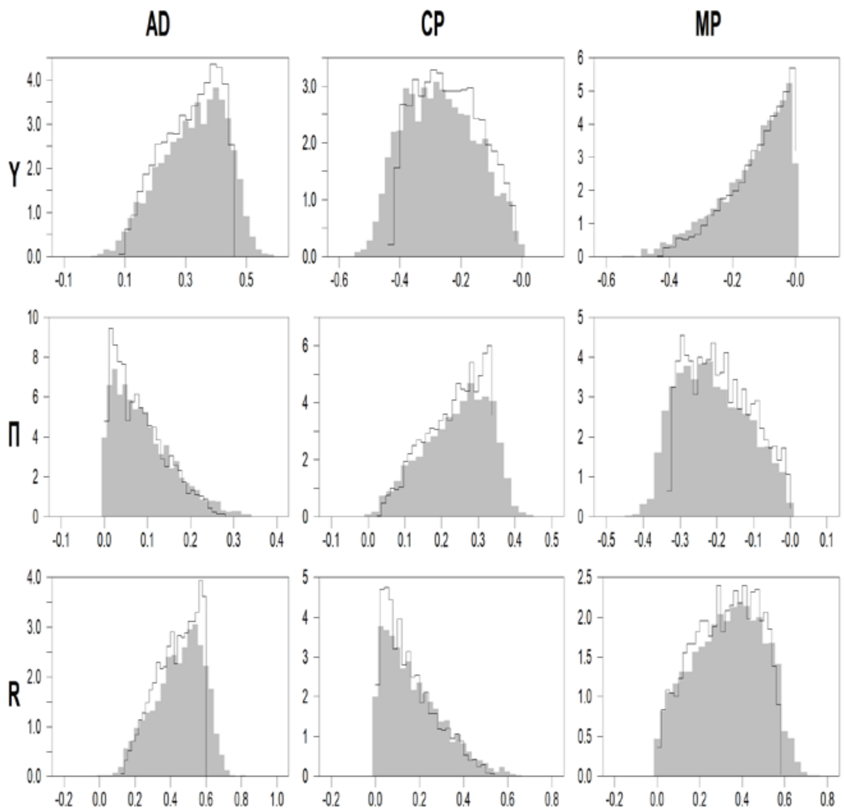


(B) Accepted rotation matrix

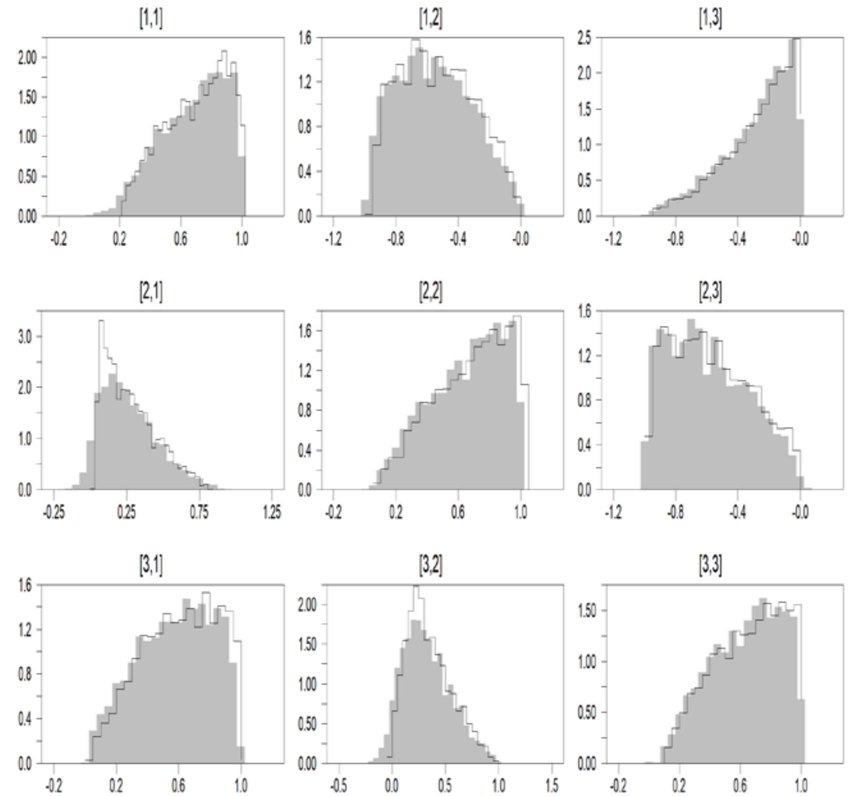


Notes: The shading (solid line) refers to histograms generated under the Bayesian (non-Bayesian) setup. In the right panel, element $(i, j)$ refers to the element in the *ith* row and *jth* column of a matrix.

Figure 4: Histograms of the impact responses and elements of accepted rotation matrices for the BH model when the sign restrictions are imposed on the rotation matrix

(A) Impact response matrix

(B) Accepted rotation matrix

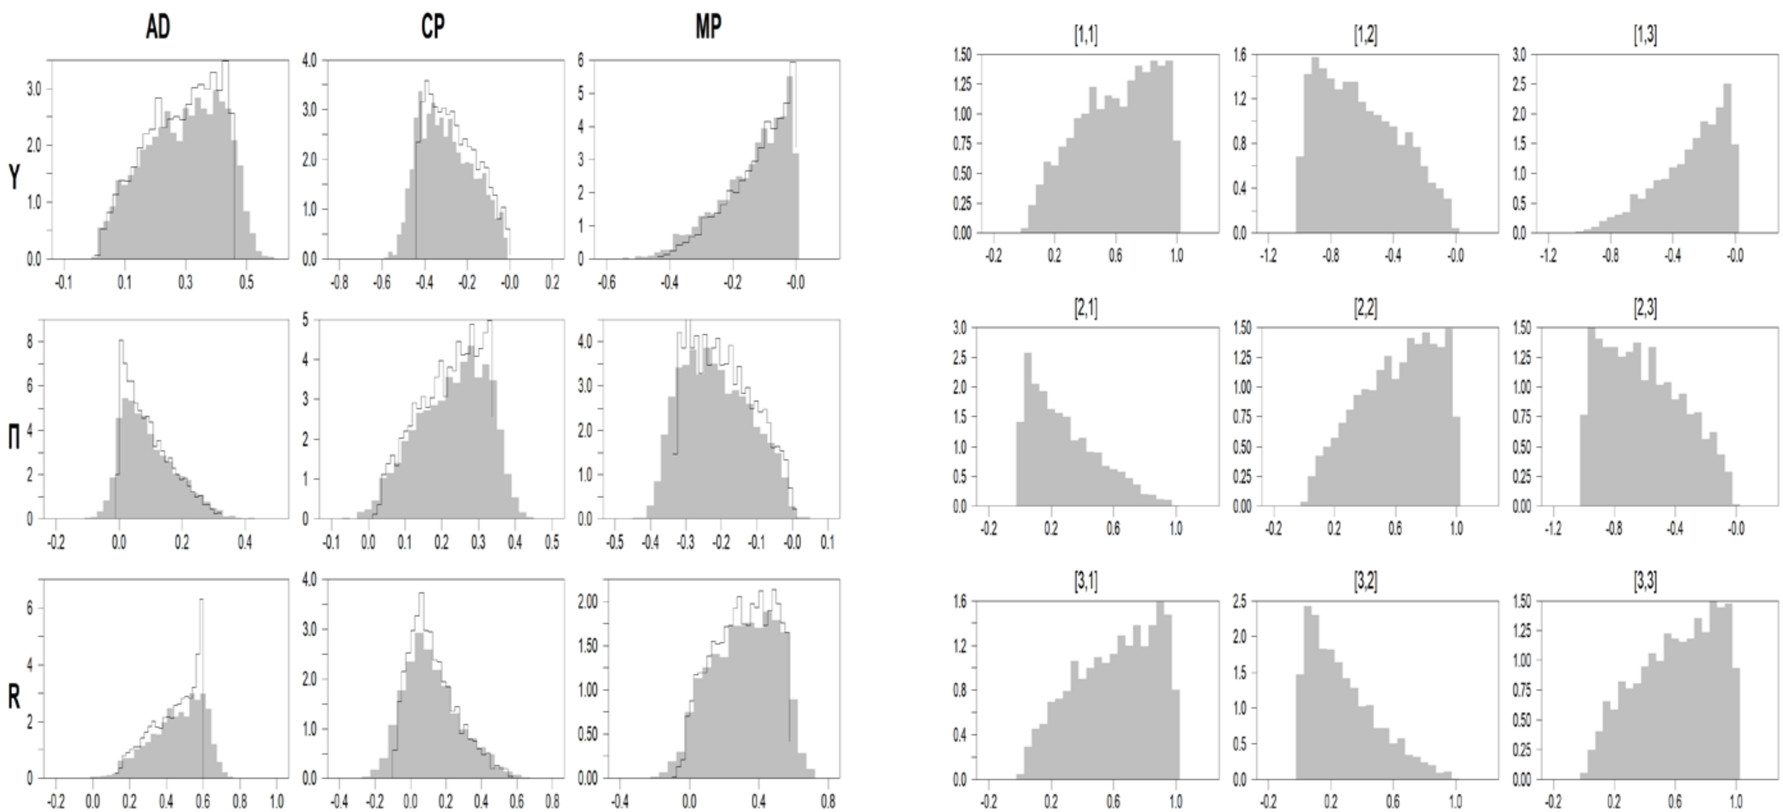


Notes: The shading (solid line) refers to histograms generated under the Bayesian (non-Bayesian) setup. In the right panel, element $(i, j)$ refers to the element in the *ith* row and *jth* column of a matrix.

Figure 5: Histograms of the responses at the horizon of one for the BH model

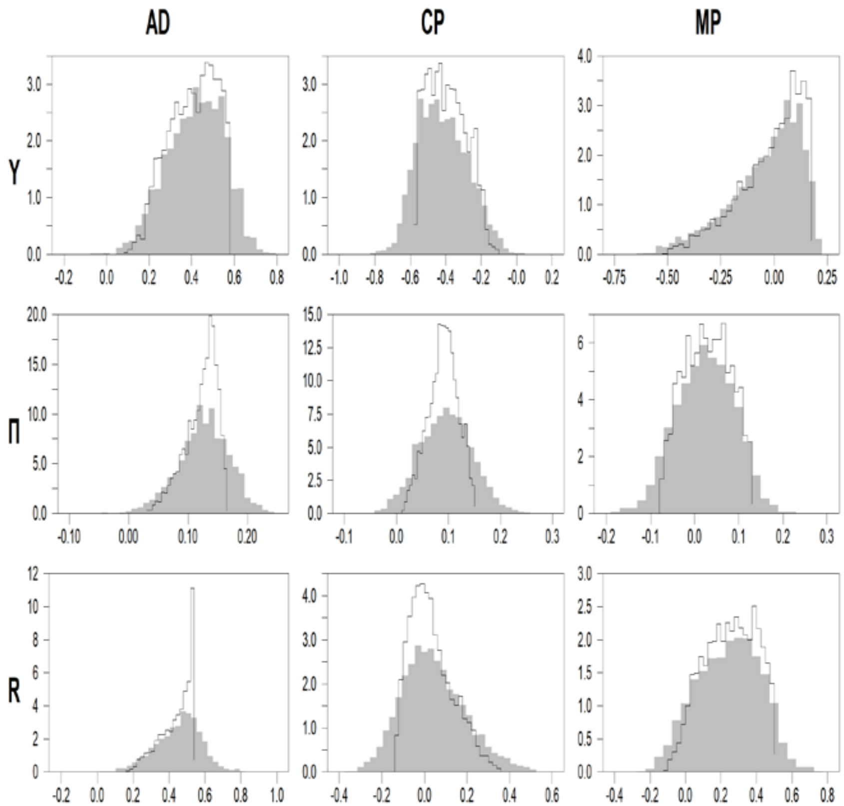


Note: The shading (solid line) refers to histograms generated under the Bayesian (non-Bayesian) setup.

Figure 6: Responses of the variables to the shocks for the PS model

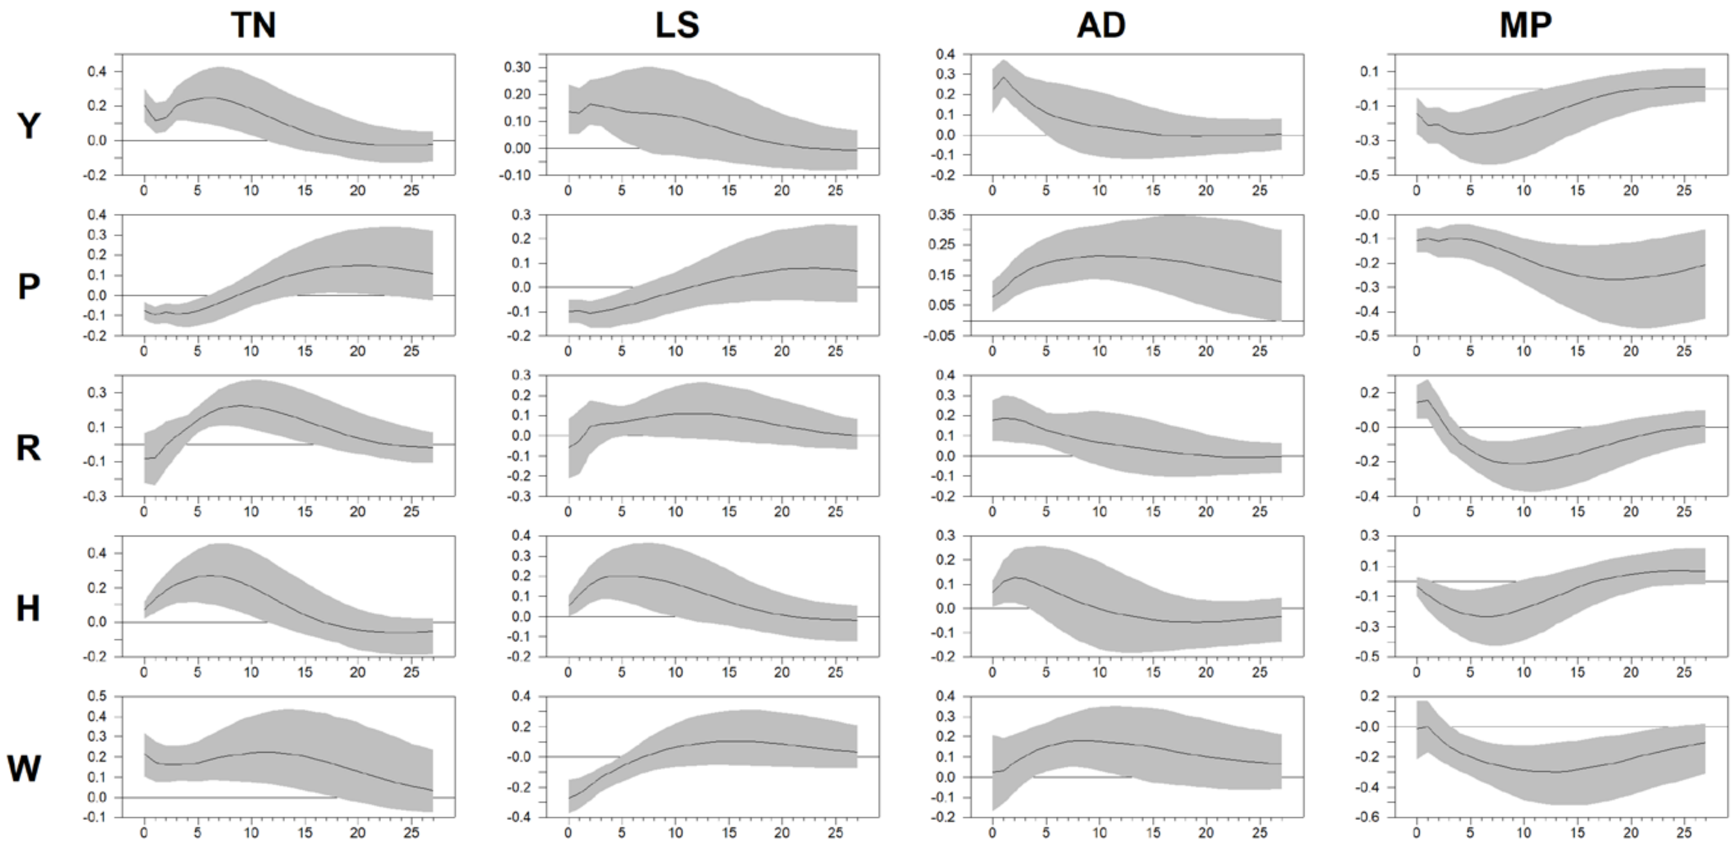


Note: The solid line and the shading are the median responses and the 16%–84% band responses, respectively.

Figure 7: Histograms of the impact responses and elements of accepted rotation matrices for the PS model

(A) Impact response matrix

(B) Accepted rotation matrix

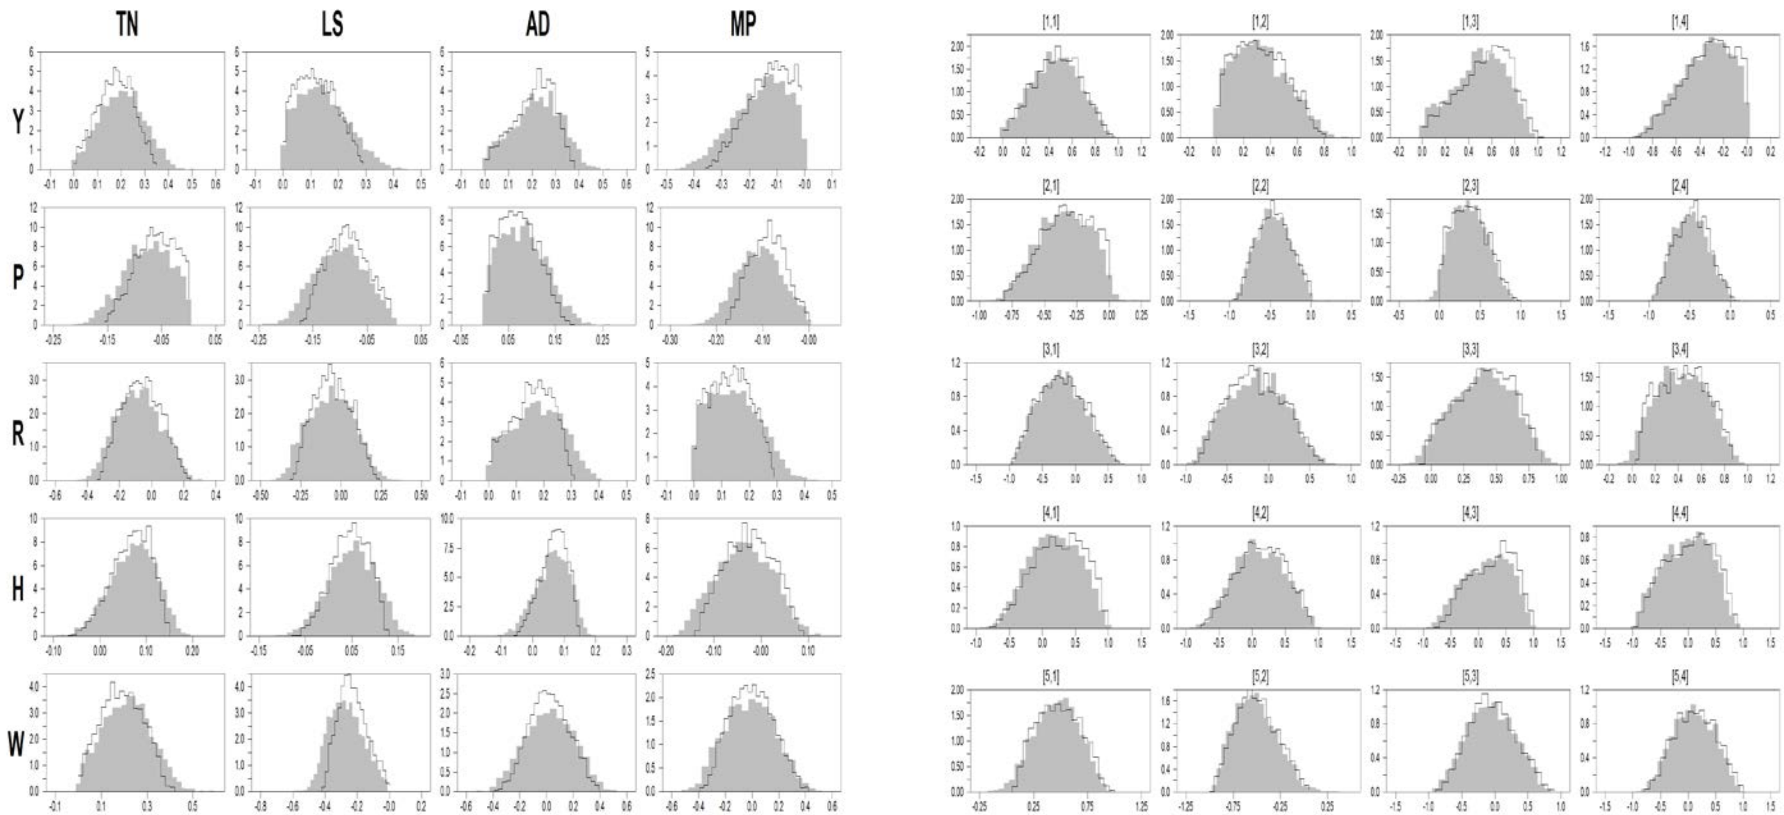


Notes: The shading (solid line) refers to histograms generated under the Bayesian (non-Bayesian) setup. In the right panel, element $(i, j)$ refers to the element in the *ith* row and *jth* column of a matrix.

Figure 8: Histograms of the impact responses and elements of accepted rotation matrices for the PS model when the sign restrictions are imposed on the rotation matrix

(A) Impact response matrix

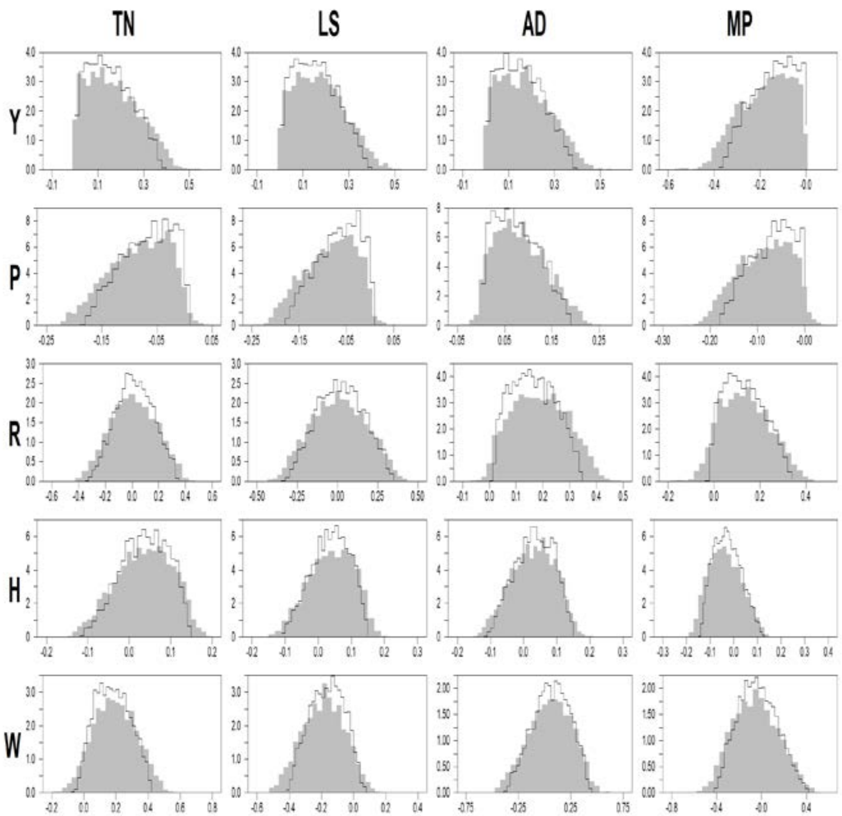


(B) Accepted rotation matrix

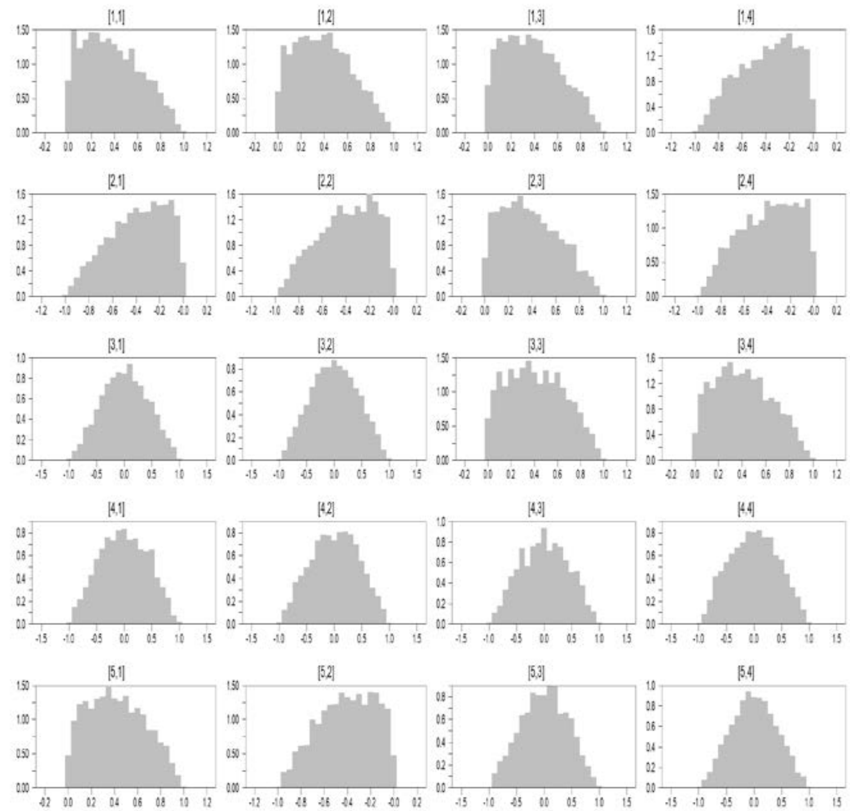


Notes: The shading (solid line) refers to histograms generated under the Bayesian (non-Bayesian) setup. In the right panel, element $(i, j)$ refers to the element in the *ith* row and *jth* column of a matrix.

Figure 9: Histograms of the responses at the horizon of one for the PS model

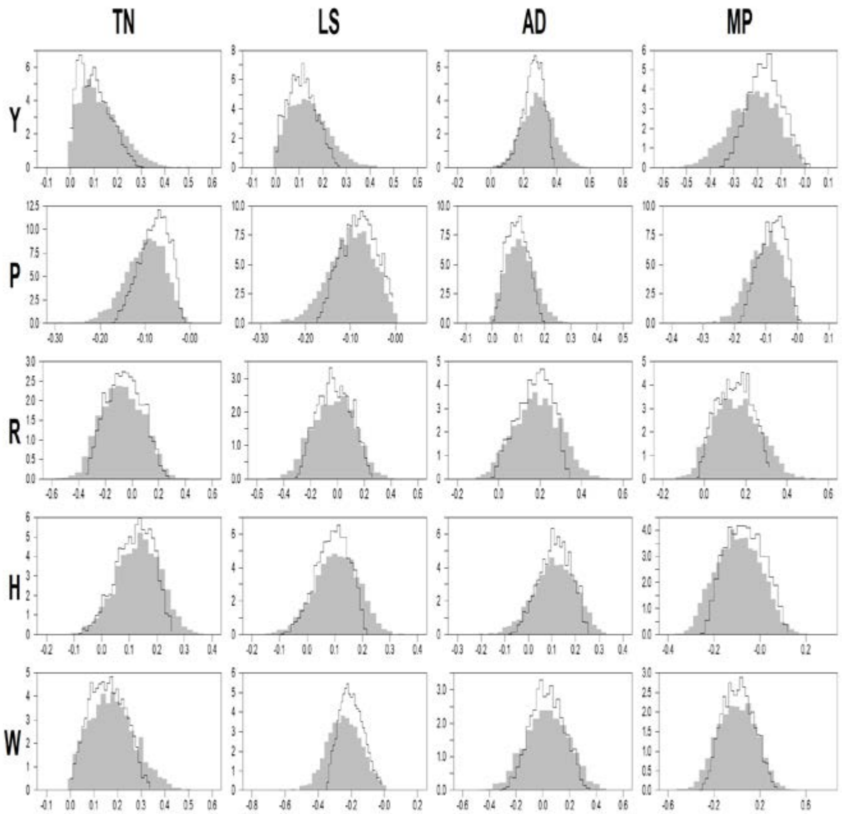


Notes: The shading (solid line) refers to histograms generated under the Bayesian (non-Bayesian) setup.