Useful Prior Information in Sign-Identified Structural Vector Autoregression: Replication of Baumeister and Hamilton (2015)*

Markku Lanne
Faculty of Social Sciences, University of Helsinki

and

Jani Luoto†
Faculty of Social Sciences, University of Helsinki

Summary

Baumeister and Hamilton (2015) encouraged researchers to explicitly defend prior information and examine its effect on the posterior results in sign-identified structural vector autoregressions. In an application to the U.S. labor market, they followed this principle by carefully motivating the priors for short-run labor elasticities by previous empirical evidence. However, we show that their findings are strongly driven by an arbitrary restriction on the long-run response of employment to the labor-demand shock. As an alternative, we propose an approach to identification that makes efficient use of non-Gaussianity in the data to facilitate point instead of set identification.

Keywords: non-Gaussian vector autoregressions, sign restrictions, Bayesian inference, set identification, GMM estimation.

*Financial support from the Academy of Finland (grant 308628) is gratefully acknowledged.

†Correspondence to: Jani Luoto, Economics, Faculty of Social Sciences, P.O. Box 17 (Arkadiankatu 7), FI-00014 University of Helsinki, Finland. E-mail: jani.luoto@helsinki.fi. Tel. +358294128718. Fax +358294128736.

1 Introduction

Baumeister and Hamilton (2015; hereafter BH) showed that impulse response functions obtained by the traditional sign-restriction approach to structural vector autoregressions (SVARs) involve implicit informative prior distributions whose influence does not vanish even asymptotically. This result follows from their Proposition 2 according to which, with non-informative prior beliefs on the other parameters of their SVAR model, the parameters capturing the simultaneities between the variables are only set identified, and hence, within the set, their joint marginal posterior remains proportional to their joint marginal prior. Therefore, in impulse response analysis based on the traditional sign-restriction approach, little is learned from the data, but the results are dictated by the implicit prior.

Instead of relying on the implicit prior, BH recommended researchers to openly defend their prior beliefs and to identify their role in influencing posterior conclusions. In practice, they accomplished the latter by comparing the marginal prior and posterior distributions of the parameters when illustrating their recommended methods. While this may, in principle, be helpful in disentangling the roles of prior and likelihood information in drawing structural conclusions, we argue that such comparisons do not necessarily provide information on how prior beliefs influence posterior conclusions. This follows from the fact that the parameters may be dependent even if their priors are independent, as pointed out by Müller (2012) and Koop et al. (2013) in the context of dynamic stochastic general equilibrium models.

Following their own recommendation, in the empirical application to the U.S. labor market, BH carefully defended priors on the short-run labor-supply and labor-demand elasticities by referring to estimates obtained in the previous literature. In this application, such information is widely available, but in most cases defending priors by previous results may not be straightforward. Moreover, with only two variables and two parameters of interest, BH were able to use likelihood contour plots in the analysis, while this is not feasible in a higher-dimensional parameter space. These plots showed that the parameters of interest are dependent, as the maximum likelihood estimate of one parameter is a function of the other, and thus, comparisons of marginal priors and posteriors of the parameters are not likely to yield useful information on the effects of prior beliefs on posterior conclusions.

For instance, a tight marginal prior for either parameter would result in a tight posterior for the other parameter irrespective of its own marginal prior.

In addition to the carefully defended priors on the short-run elasticities, BH introduced a long-run restriction on the effect of the labor-demand shock on employment, with no economic justification. The latter restriction received little support from the data, but because of it, comparison of the marginal prior and posterior distributions of the parameters of interest fails to provide information about the relative importance of the prior and observed data. As a matter of fact, this long-run restriction was actually driving BH's results, which we demonstrate in Section 2.1 by comparing the marginal posterior distributions of the short-run labor-supply and labor-demand elasticities obtained under tight and loose priors on this restriction but keeping the marginal priors intact. The observation that they are quite different indicates that BH's analysis failed to identify the relative importance of the short-run and long-run priors for posterior results, and hence their role in influencing the structural conclusions remains far from obvious. Moreover, as we also show in Section 2.1, quite different results are obtained if the corresponding long-run restriction is (equally arbitrarily) imposed on the effect of the labor-supply shock on employment instead.

BH advocated openly acknowledging the prior information used in the analysis of sign-identified SVAR models and assessing its role in obtaining the posterior results. However, as we demonstrate, their advice is in practice difficult to follow, and even their own empirical illustration failed to convey the role of all prior beliefs in influencing the conclusions in a transparent manner. The main problem is that even with informative priors, sign restrictions only yield set identification, when the structural errors are assumed normal. In contrast, by making use of potential non-Gaussianity in the data, along the lines of the recent advances in the statistical identification literature,[1] point identification can be reached, and thus, many of the problems discussed above be solved. In particular, Lanne et al. (2017) show that the parameters of a SVAR model are uniquely identified (up to permutation of the structural shocks) if at most one of them is marginally normally distributed. While Lanne et al. make quite stringent assumptions, including independence

[1] See Kilian and Lütkepohl (2017, Chapter 14) for a survey of the burgeoning literature on statistical identification of SVAR models.

and parametric marginal distributions for each of the structural shocks, we follow Lanne and Luoto (2018), and consider the generalized method of moments (GMM) framework which requires no explicit distributional assumptions and allows for various forms of conditional heteroskedasticity. Once the statistically identified shocks have been obtained, they can be labelled by the theory-implied signs. In this way, we can sidestep the ambiguity related to the examination of the role of prior information in drawing posterior conclusions,

This paper is organized as follows. In Section 2.1, we show that BH's results are actually driven by their arbitrary long-run restriction, whereas the marginal priors only play a minor role. We also demonstrate how an alternative long-run restriction yields quite different conclusions. Section 2.2 contains the results of the point-identified SVAR model. We resoundingly reject BH's long-run restriction, and obtain impulse responses quite different from theirs. In particular, a positive labor-demand shock is found to have a significantly positive and persistent effect on employment. Finally, Section 3 concludes.

2 Model of U.S. labor market

We demonstrate the issues brought up in the Introduction by means of BH's model of labor supply and labor demand:

$$\begin{aligned}
\Delta n_t &= k^d + \beta^d \Delta w_t + b_{11}^d \Delta w_{t-1} + b_{12}^d \Delta n_{t-1} + b_{21}^d \Delta w_{t-2} \\
&\quad + b_{22}^d \Delta n_{t-2} + \cdots + b_{m1}^d \Delta w_{t-m} + b_{m2}^d \Delta n_{t-m} + u_t^d, \\
\Delta n_t &= k^s + \alpha^s \Delta w_t + b_{11}^s \Delta w_{t-1} + b_{12}^s \Delta n_{t-1} + b_{21}^s \Delta w_{t-2} \\
&\quad + b_{22}^s \Delta n_{t-2} + \cdots + b_{m1}^s \Delta w_{t-m} + b_{m2}^s \Delta n_{t-m} + u_t^s, \qquad (1)
\end{aligned}$$

where Δn_t is the growth rate of employment, Δw_t is the growth rate of real compensation per hour, β^d is the short-run wage elasticity of demand, and α^s is the short-run wage elasticity of supply. The vector of structural errors $\mathbf{u}_t = (u_t^d, u_t^s)'$ is independently and identically normally distributed with diagonal covariance matrix $\mathbf{D}$.[2]

[2] We use the same data as BH and in the Bayesian analyses their Matlab code that were kindly provided by Christiane Baumeister on her website at `https://sites.google.com/site/cjsbaumeister/research`.

2.1 Bayesian analysis under sign restrictions

BH show that if $\mathbf{u}_t$ is Gaussian, the maximum likelihood (ML) estimates of β^d and α^s are dependent such that there is an infinite number of pairs (β^d, α^s) that maximize the likelihood function. Therefore, to narrow down the range of plausible parameter values, they consulted the previous literature and consequently used truncated Student t priors with means –0.6 and 0.6 for β^d and α^s, respectively. In addition, they imposed a tight prior (with variance $V = 0.1$) on a restriction on the long-run response of employment to the labor-demand shock u_t^d. BH's Figure 8 reveals that the response of employment to the demand shock and the posterior distribution of α^s are strongly driven by the long-run restriction. However, their Figure 6 (panel C) indicates that there is very little support to this restriction (with the small prior variance that most of their results are based on).

Figure 1 depicts the marginal prior and posterior distributions of the parameters of interest assuming a very loose prior ($V = 1000$) on the long-run restriction. Comparison of it to BH's Figure 6 (Panels A and B) based on a tight prior of the long-run restriction ($V = 0.1$) illustrates the limitations of inspecting marginal prior and posterior distributions in figuring out how observed data causes prior beliefs to be revised. Despite being based on the same marginal priors of β^d and α^s, the marginal posteriors in these two cases are quite different, and the differences must arise from the tightness of the restriction on the long-run response of employment on the labor-demand shock. It is noteworthy that the marginal distributions in Figure 1 very closely resemble those obtained by the traditional approach to sign restrictions depicted in BH's Figure 2. This indicates that BH's carefully defended marginal priors indeed play a minor role. Moreover, the impulse responses of employment to the labor-supply and labor-demand shocks in the case of the loose prior for the long-run restriction (not shown to save space, but available upon request) closely resemble each other (in contrast to BH's impulse responses obtained under the tight prior for the long-run restriction in their Figure 7). These findings suggest that, assuming normality, little information can be extracted from the data alone about the parameters of interest and the effects of the labor-supply and labor-demand shocks. It seems that BH's results were actually driven by the tight prior for their (implausible) long-run restriction. In other words, even though BH carefully defended the marginal priors of the key parameters, they failed to report how little the observed data alone cause these prior beliefs to be revised.

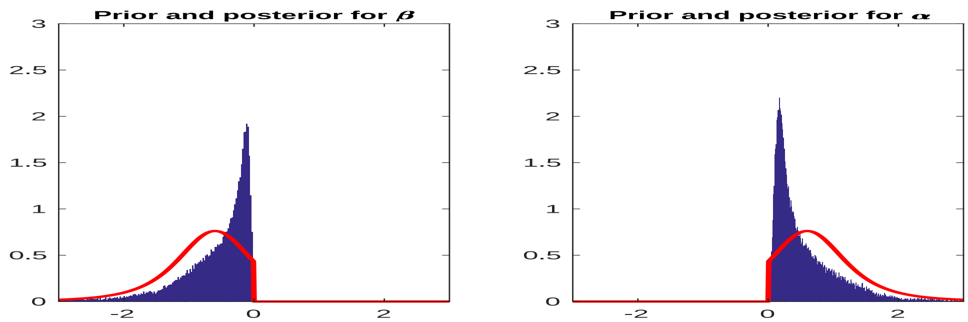

Figure 1: Posterior (histograms) and prior (solid lines) distributions of the short-run elasticities of labor demand β^d and labor supply α^s based on a loose prior on the restriction $b^s_{11} + \cdots + b^s_{m1} + \alpha^s = 0$ (with variance $V = 10000$), and BH's truncated Student t priors.

In order to highlight the arbitrariness of BH's long-run restriction, we re-estimated the model using the same short-run priors, but restricting the long-run effect of the labor-supply (instead of labor-demand) shock on employment. The resulting marginal prior and posterior distributions of β^d and α^s and the impulse responses are depicted in Figures 2 and 3, respectively. Despite the same marginal priors, the posterior distributions in Figure 2 are quite different from those based on BH's long-run restriction in their Figure 6. It is now the posterior of α^s that closely resembles its prior, while the posterior of β^d is concentrated on values just below zero. Under the alternative long-run restriction, the impulse responses of the shocks are also reversed compared to those in BH's Figure 7: the supply shock has virtually no effect on employment, and the demand shock has only a minor effect on wages.

2.2 Point identification by non-Gaussianity

Instead of carefully specified priors for the key parameters and examination of their effects on the posterior results advocated by BH, we consider a frequentist approach that makes efficient use of non-Gaussianity in the structural errors. We start out by estimating a reduced-form eighth-order vector autoregressive (VAR(8)) model for $(\Delta n_t, \Delta w_t)'$ by ordinary least squares. The residuals exhibit clear non-normality, indicating that at least one

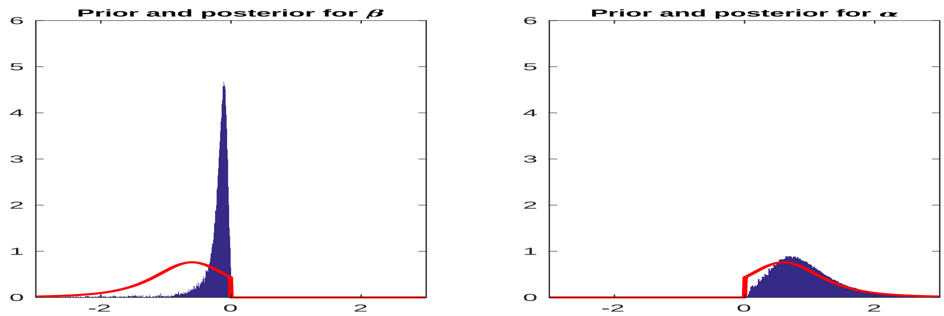

Figure 2: Posterior (histograms) and prior (solid lines) distributions of the short-run elasticities of labor demand β^d and labor supply α^s based on a tight prior on the restriction $b_{11}^d + \cdots + b_{m1}^d + \beta^d = 0$ (with variance $V = 0.1$), and BH's truncated Student t priors.

of the structural errors must be non-Gaussian, which guarantees point identification.[3]

We do not entertain any specific non-Gaussian distributions for the structural errors, but follow Lanne and Luoto (2018) and estimate the SVAR model by the efficient GMM based on the following moment conditions:

$$E\left(\mathbf{u}_t \otimes \mathbf{x}_{t-1}\right) = 0_{2k\times 1} \tag{2a}$$

$$E\left(u_t^d\right)^2 - d_{11} = 0 \tag{2b}$$

$$E\left(u_t^s\right)^2 - d_{22} = 0 \tag{2c}$$

$$E(u_t^d u_t^s) = 0 \tag{2d}$$

where $\otimes$ denotes the Kronecker product, $\mathbf{x}_{t-1}$ is a $((2m+1)\times 1)$ vector containing the m lags of $\mathbf{y}_t$, and d_{ii} $(i = 1, 2)$ are the diagonal elements of $\mathbf{D}$. Conditions (2a), implicitly assume that the lag length 8 is sufficient to make the components of the error term $\mathbf{u}_t$ serially uncorrelated, while (2b) and (2c) concern the error variances, and condition (2d) makes the errors orthogonal. In addition, we impose either of the following co-kurtosis

[3] The values of the Jarque-Bera test statistic for non-normality are 11.192 and 11.852, with asymptotic p values 0.015 and 0.011, respectively, indicating non-Gaussianity at the 5% significance level. Following, Kilian and Demiroglu (2000), we reconfirmed this conclusion by bootstrap. The corresponding bootstrapped (based on 10,000 replications) 5% (1%) critical values equal equal 8.880 and 5.941 (18.626 and 13.028), respectively, likewise indicating rejection of normality at the 5%, but not at the 1% significance level.

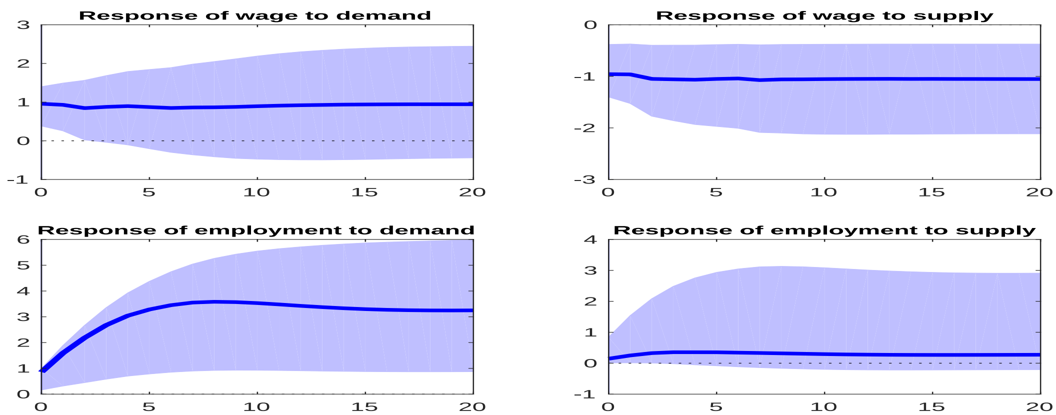

Figure 3: Posterior medians and 95% credibility sets of the impulse responses of the labor-demand and labor-supply shocks based on the tight prior on the restriction $b_{11}^d + \cdots + b_{m1}^d + \beta^d = 0$ (with variance $V = 0.1$) and BH's truncated Student t priors.

conditions:

$$E(\left(u_t^d\right)^3 u_t^s) = 0 \tag{3a}$$

$$E(\left(u_t^s\right)^3 u_t^d) = 0 \tag{3b}$$

The goal of the latter conditions is to obtain shocks that are close to being independent without actually imposing independence. They provide useful information for estimation in the presence of non-Gaussianity, whereas they are redundant under Gaussianity.

As shown by Lanne and Luoto (2018), when standard GMM regularity conditions hold, moment conditions (2a)–(2d) and either (3a) or (3b) yield local and global point identification. However, identification is lost if both of conditions (3a) and (3b) are included. The relevant moment selection condition (RMSC) of Hall et al. (2007) selects conditions (2a)–(2d) and (3b).[4] With 37 moment conditions and 36 parameters, the model is over-identified, and the p value of Hansen's (1982) J test of over-identification restrictions equals 0.705, indicating validity of the moment conditions.

Additional information is needed to label the estimated equations (or, equivalently, the shocks). To that end, we use the theoretically implied signs of the short-run labor-supply

[4]The values of the RMSC equal –220.7 and –228.3, when estimation is based on conditions (3a) and (3b) in addition to conditions (2a)–(2d), respectively.

and labor-demand elasticities, and label the equation with a positive (negative) coefficient estimate of Δw_t the supply (demand) equation. The resulting estimates of the labor-supply elasticity α^s and labor-demand elasticity β^d equal 0.765 and -0.197 with asymptotic standard errors 0.196 and 0.057, respectively. According to asymptotic t tests, both parameters are significantly different from zero at conventional significance levels. The point estimate of β^d is in accordance with the range based on microeconometric studies (from –0.75 to –0.15) referred to by BH, but not with BH's empirical result (their posterior estimate of β^d lies close to unity in absolute value). Also, while our 95% confidence interval for α^s has a lot of overlap with the estimates reported in microeconometric and macroeconometric studies referred to by BH, the most of the probability mass of BH's posterior distribution of α^s lies even below the minimum of the estimates in the microeconometric literature (0.15).

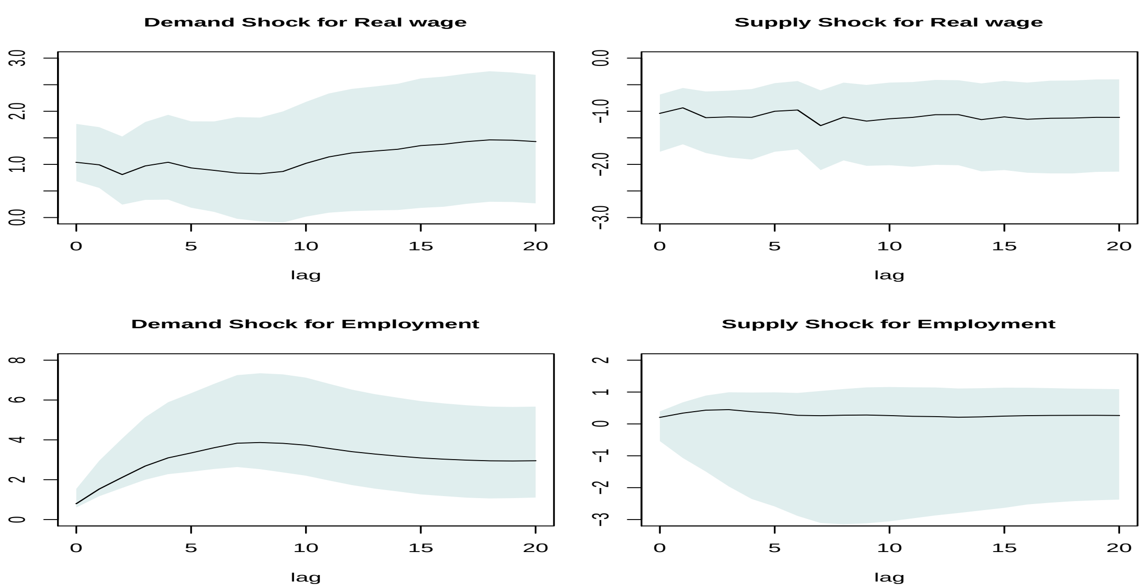

Figure 4: Impulse responses of the labor-supply and labor-demand shocks, and their pointwise 95% Hall's percentile confidence bands obtained by bootstrap with 10,000 replications.

The statistically significant effect of both shocks on the real wage on impact at the 5% level is reconfirmed by the bootstrapped 95% confidence bands of the impulse responses of the two shocks displayed in the upper panel of Figure 4. These responses are also in accordance with BH's posterior results, whereas we obtain quite different effects of both shocks on employment, as seen in the lower panel of Figure 4. Specifically, we find the effect of the labor-demand shock much stronger than BH, while the effect of the labor-supply shock turns out insignificant. The former finding is not surprising, however, as BH's results were strongly driven by a relatively tight prior on the restriction on the long-run labor-demand elasticity, as already discussed. One advantage of our approach is that any restrictions on the parameters can be tested in a straightforward manner, and BH's long-run restriction turns out to be clearly rejected with a p value 0.0003 in a Wald test.

3 Conclusion

In this note, we have questioned the general usefulness of BH's recommendation of "explicitly defending the prior information used in the analysis and reporting the way in which the observed data causes these prior beliefs to be revised" in the empirical analysis of sign-identified SVARs. While it may provide an improvement over the traditional way of inference in sign-identified SVARs, especially under Gaussianity, both approaches fail to make efficient use of potential non-Gaussianity. Moreover, in addition to being infeasible in large models in practice, comparisons of the marginal prior and posterior distributions of the parameters proposed by BH are not, in general, useful in disentangling the roles of prior and observed information on the posterior if the parameters are dependent.

We have illustrated the above points in BH's empirical application to the U.S. labor market. Their carefully defended priors on the short-run labor elasticities had in fact essentially no effect on the posteriors, which were dominated by the restriction on the long-run response of employment to the labor-demand shock. When the latter restriction was virtually absent, BH's priors of the short-run elasticities resulted in more or less the same posteriors as the traditional sign-restricted approach. Moreover, completely different conclusion were reached when BH's arbitrary long-run restriction was replaced by an alternative one.

As a potentially preferable approach, we have suggested exploiting non-Gaussianity in the structural errors in a way that yields point identification instead of set identification. In particular, we estimated the bivariate SVAR model of BH by the GMM with moment conditions informative under non-Gaussian structural errors. In this framework, we resoundingly rejected BH's long-run restriction, and, in contrast to BH, obtained estimates of the labor-supply and labor-demand elasticities that conform to estimates obtained in the previous literature. Moreover, we found the labor-demand shock to have a significant and persistent effect on employment. This goes contrary to BH, who, presumably because of the (incorrect) long-run restriction, concluded this effect to be unimportant.

References

Baumeister, C. & Hamilton, J. D. (2015). Sign restrictions, structural vector autoregressions, and useful prior information. *Econometrica*, 83, 1963–1999.

Hall, A.R., Inoue, A., Jana, K. & Sin, C. (2007). Information in generalized method of mometns estimator in misspecified models. *Journal of Econometrics*, 138, 488–512.

Kilian, L. & Demiroglu, U. (2000). Residual-based tests for normality in autoregressions: Asymptotic theory and simulation evidence. *Journal of Business and Economic Statistics*, 18, 40–50.

Kilian, L., & Lütkepohl, H. (2017). *Structural vector autoregressive analysis*. Oxford: Oxford University Press.

Koop, G., Pesaran, M. H. & Smith, R. P. (2013). On identification of Bayesian DSGE models. *Journal of Business and Economic Statistics*, 31, 300–314.

Lanne, M. & Luoto, J. (2018). GMM estimation of non-Gaussian structural vector autoregression (HECER Working Paper 423). University of Helsinki.

Lanne, M. Meitz, M. & Saikkonen, P. (2017). Identification and estimation of non-Gaussian structural vector autoregressions. *Journal of Econometrics*, 196, 288–304.

Müller, U. (2012). Measuring prior sensitivity and prior informativeness in large Bayesian models. *Journal of Monetary Economics*, 59, 581–597.