


Opaque Prior Distributions in Bayesian Latent Variable Models

Edgar C. Merkle¹ , Oludare Ariyo² , Sonja D. Winter³ ,

Mauricio Garnier-Villarreal⁴ 

[1] *Department of Psychological Sciences, University of Missouri, Columbia, MO, USA.* [2] *School of Mathematics, Statistics and Actuarial Science, University of Essex, United Kingdom.* [3] *College of Education & Human Development, University of Missouri, Columbia, MO, USA.* [4] *Department of Sociology, Vrije Universiteit Amsterdam, Amsterdam, The Netherlands.*

Methodology, 2023, Vol. 19(3), 228–255, <https://doi.org/10.5964/meth.11167>

Received: 2023-01-20 • **Accepted:** 2023-08-26 • **Published (VoR):** 2023-09-29

Handling Editor: Francisco J. Abad, Universidad Autónoma de Madrid, Madrid, Spain

Corresponding Author: Edgar C. Merkle, 28A McAlester Hall, University of Missouri, Columbia, MO 65203, USA.
E-mail: merklee@missouri.edu

Abstract

We review common situations in Bayesian latent variable models where the prior distribution that a researcher specifies differs from the prior distribution used during estimation. These situations can arise from the positive definite requirement on correlation matrices, from sign indeterminacy of factor loadings, and from order constraints on threshold parameters. The issue is especially problematic for reproducibility and for model checks that involve prior distributions, including prior predictive assessment and Bayes factors. In these cases, one might be assessing the wrong model, casting doubt on the relevance of the results. The most straightforward solution to the issue sometimes involves use of informative prior distributions. We explore other solutions and make recommendations for practice.

Keywords

Bayesian SEM, prior distributions, Bayesian psychometrics, Stan, blavaan

Bayesian latent variable models, including structural equation models and time series models, are becoming increasingly complex to accommodate the datasets that modern technology permits (e.g., Asparouhov et al., 2018; Bürkner, 2021; Depaoli et al., 2021; dos Santos et al., 2022; Driver et al., 2017; Enders et al., 2018; Garnier-Villarreal et al., 2021; Haaf et al., 2020; Haaf & Rouder, 2023; Hoijtink & van de Schoot, 2018; Kaplan



et al., 2023; Keller & Enders, 2023; Levy & Enders, 2021; Magnus & Garnier-Villarreal, 2022; Man & Culpepper, 2022; Miočević et al., 2021; Paganin et al., 2023; Rast et al., 2022; Uanhoro, 2023a, 2023b; Ulitzsch et al., 2022; Van De Schoot et al., 2013; Van Erp & Browne, 2021; Wilcox et al., 2021; Zyphur et al., 2021). To ensure that these complex models are working correctly, it is important to fully understand the problematic issues that can arise in simpler Bayesian models with latent variables. We think that this understanding is not yet fully developed. For example, researchers often overlook the fact that multiple WAIC metrics are available for the same model, depending on whether or not one includes latent variables in the likelihood (Merkle et al., 2019). And recent evidence indicates that small-variance priors can mask misfit in other parts of the model (Jorgensen & Garnier-Villarreal, 2023). Continuing with the theme of problematic issues, the current paper discusses problems with prior distributions that are easy to overlook and that can lead to incorrect results and model summaries.

The issues that we consider here can generally be called *opaque prior distributions*. They involve the fact that, for some Bayesian models, the prior distribution that a researcher specifies is not the prior distribution that the estimation method actually uses. This happens because prior distributions are influenced by various details surrounding MCMC implementation, beyond the researcher's specified prior. These issues are already known to many statisticians, but they have some unique manifestations in Bayesian psychometric models with latent variables. Further, these issues have not received much attention in the context of Bayesian psychometric modeling.

In psychometric models with latent variables, opaque prior distributions can arise from positive definite constraints associated with model covariance matrices, from sign indeterminacies in factor loadings, and from order constraints associated with threshold parameters. One example was considered by Ghosh and Dunson (2009), who describe a parameter expansion algorithm for estimating Bayesian factor analysis models. They developed a Gibbs sampler for an overparameterized model in which the factor loadings were not identified, then translated the unidentified parameters to identified parameters via postprocessing. Of interest here, they showed that the priors for the identified loadings (obtained via postprocessing) differed from the priors for the unidentified loadings. We consider this example in more detail later.

Opaque prior distributions can cause a variety of problems. One way in which these problems arise is through recent studies where researchers seek to recommend the "best" prior distributions for various models. For example, specific prior distributions have been recommended for factor loadings and correlation parameters in confirmatory factor analysis (Lüdtke et al., 2021; Ulitzsch et al., 2023), factor loadings and intercept parameters in item factor analysis with dichotomous indicators (Bainter, 2017), regression parameters in mediation models with latent variables (Miočević et al., 2021), covariance matrices in SEM (Liu et al., 2022; Van Zundert et al., 2022), random effect parameters in multilevel SEMs (Van Erp & Browne, 2021; Zitzmann et al., 2021), and variance parameters in

approximate measurement invariance modeling (Pokropek et al., 2020). However, as we will show, opaque priors imply that such recommendations can be specific to the software used. In other words, alternative software packages may implement the same model in a different way, leading to a different implied prior distribution even though the user always specified the same prior distribution.

We could also experience problems with metrics that explicitly involve evaluation of the prior, including Bayes factors (e.g., Heck et al., 2023; Kass & Raftery, 1995) and prior predictive checks (e.g., Vanpaemel, 2020). In both cases, if a researcher specifies a model with opaque prior distributions, then it is easy to use the wrong prior distributions to compute the metrics. Relatedly, some methods for verifying the accuracy of one's MCMC algorithm involve generating data from prior distributions. If researchers use the wrong priors, then they may conclude that their MCMC algorithm is problematic even though it runs correctly, or vice versa.

The intent of this paper is to illustrate opaque prior distributions and to provide solutions for avoiding them. This can allow for reproducible results across software implementations, and it can lead to improved prior predictive assessments and other metrics. In the pages below, we show how opaque priors can arise from positive definite constraints, from sign indeterminacies, and from parameter order constraints. For each of these topics, we provide examples of the problem, discuss how the problem can lead to compounding problems in model assessment, and provide recommendations for avoiding the problem. Finally, we summarize and make general recommendations for practice.

Positive Definite Constraints

In Bayesian modeling, prior distributions for the covariance matrix often involve the inverse-Wishart (IW) distribution due to its conditional conjugacy. However, the IW distribution can be problematic because it assumes the same amount of prior information for the entire covariance matrix. To overcome this challenge, various strategies have been suggested for separately specifying priors on the variance and correlation parameters underlying the covariance matrix (Barnard et al., 2000). These priors have been shown to perform better than the classical IW (Alvarez et al., 2014; Ariyo et al., 2022; Huang et al., 2013). However, things become more complicated for matrix dimensions of three or more, because certain restrictions must be imposed to ensure positive definiteness of the matrix (Barnard et al., 2000; Daniels & Kass, 1999; Huang et al., 2013; Hurtado Rúa et al., 2015; Wei & Higgins, 2013).

The situation becomes even more complicated when the covariance matrix has model-imposed constraints, which can arise in SEMs with correlated residuals or with across-group equality constraints. In these models, we cannot impose an IW (or other prior) on the full covariance matrix because those priors will not respect the constraints imposed by the model. An easily-implemented approach is to place priors on individual

parameters within the covariance matrix (and we consider other approaches later). The problem is that, when we build a covariance matrix using these parameters, the resulting matrix will sometimes be non-positive definite. We elaborate below.


Illustration

An example comes from the popular “Political Democracy” model originally described by Bollen (1989), shown below in *lavaan* syntax (Rosseel, 2012). Intended to describe countries’ relationships between their levels of industrialization and democracy, the model contains four observed variables that are measured once during 1960 (y_1, \dots, y_4) and again during 1965 (y_5, \dots, y_8). The residuals of the 1960 variables are allowed to correlate with the residuals of the corresponding 1965 variables. Additionally, there exists a pair of similar variables collected during 1960 and again during 1965, leading to two more residual correlations. The full structure of the residual covariance matrix is shown in Figure 1.

As described earlier, we could elect to put a univariate prior distribution on each variance (or standard deviation or precision) in this matrix, and also on the six residual correlations that are not fixed to zero. But this is problematic because the univariate priors on correlations can yield non-positive definite correlation matrices. The specific problem that occurs depends on the software package. For example, JAGS will stop as soon as it encounters a correlation matrix that is not positive definite. This means that univariate priors on correlations cannot often be used. On the other hand, Stan will report that a non-positive definite matrix was encountered, reject it, and continue sampling. We are left with univariate priors that are collectively constrained to be positive definite. If we consider only the space of positive definite correlation matrices under these priors, then the priors are usually more informative than we originally specified. That is, the prior distributions are opaque: the analyst specifies a set of priors that are different from the implied prior distributions, which must obey the constraint of positive definiteness.

Figure 1

Political Democracy Model, Structure of Residual Covariance Matrix With Free Parameters Marked by ×.




```
model <- '  
  # measurement model  
  ind60 =~ x1 + x2 + x3  
  dem60 =~ y1 + a*y2 + b*y3 + c*y4  
  dem65 =~ y5 + a*y6 + b*y7 + c*y8  
  # regressions  
  dem60 ~ ind60  
  dem65 ~ ind60 + dem60  
  # residual correlations  
  y1 ~~ y5  
  y2 ~~ y4 + y6  
  y3 ~~ y7  
  y4 ~~ y8  
  y6 ~~ y8  
'
```

How can we characterize the implied prior distributions? We could simply generate thousands of correlation matrices from the prior, then discard matrices that are not positive definite, then visualize what is left. As an example of this, we specified a $\text{Uniform}(-1, 1)$ prior for each free correlation in [Figure 1](#). We then generated 100,000 correlation matrices with the desired structure, discarding the 57,818 matrices that were not positive definite. Finally, we examined the distributions of the remaining matrices, with a histogram for a single correlation parameter appearing in [Figure 2](#). We see that the resulting distribution is no longer uniform; it is approximately symmetric around 0, with more density near 0 than near -1 or 1 . This is because our uniform priors did not account for the fact that correlation matrices must be positive definite.

Figure 2

Bollen Model, Implied Prior Distribution of a Single Correlation Parameter After Accounting for Positive Definite Constraints



Positive Definiteness

To technically describe the positive definite constraint here, we can simultaneously permute the rows and columns of the correlation matrix to obtain a block diagonal matrix. A block diagonal matrix is useful because the determinant of the full matrix is the product of determinants of each individual block within the matrix. This allows us to examine whether or not the full matrix is positive definite, by working with submatrices of smaller dimension.

The Cuthill-McKee algorithm (Cuthill & McKee, 1969) allows us to automatically find an appropriate, block-diagonal permutation. After carrying out this algorithm (via the *netprioR* package; Schmich, 2022) and permuting the rows and columns, we arrive at the matrix in Figure 3. This shows that, to keep the full matrix positive definite, we only need to worry about the 4×4 block in the lower right that involves y_2 , y_4 , y_6 , and y_8 . The priors on the (y_3, y_7) and (y_1, y_5) correlations have no influence on the positive definiteness of the full matrix, because a 2×2 matrix is positive definite for any correlation in $(-1, 1)$, so that we can safely put a uniform prior on each of those correlations.


After applying traditional rules for computing matrix determinants, we can express the determinant of the 4×4 block of correlations as

$$\det(\mathbf{R}_{(4 \times 4)}) = 1 + (r_1 r_4 - r_2 r_3)^2 - \sum_{i=1}^4 r_i^2. \quad (1)$$

This shows analytically how the positive definite constraint of our correlation matrix influences our univariate priors. The univariate priors are collectively constrained by the requirement that Equation (1) be greater than 0. For MCMC algorithms that sample

Figure 3

Political Democracy Model, Structure of Permuted Correlation Matrix With Free Parameters Marked by ×.



the correlations individually, each correlation is impacted by the current values of the other correlations. For example, if the current values of r_1 , r_2 , and r_3 are .9, .1, and .1, respectively, then r_4 must be less than .9 in order to maintain a positive definite matrix. But if the first three correlations assume different values, then it would be possible to observe values of r_4 above .9.

Implications for Bayes Factors

To show how this issue becomes problematic for applied work, we consider the calculation of Bayes factors using the Bollen model. Imagine that we wish to know whether the two residual correlations involving y_2 are necessary. To address this question, we could compute a Bayes factor comparing a model that includes those two residual correlations, to a model without those residual correlations. A computationally-cheap way to do this is the Savage-Dickey method (Dickey & Lientz, 1970; Wagenmakers et al., 2010). This involves comparing the prior distributions of the two residual covariances to their posterior distributions, focusing on the point at which those covariances equal 0.

The Savage-Dickey method is slightly more complicated than usual because the two residual correlations in question influence the possible values that other residual correlations can take (due to positive definiteness). This changes the prior distributions on other residual correlations, as we move from a model with the two focal correlations freed to a model with the two focal correlations fixed. Heck (2019) discusses that, in such a situation, we need to supply a correction term to the usual Savage-Dickey calculation that accounts for the change in prior distributions (also see Verdinelli & Wasserman, 1995).

If we ignore (or do not realize) all of the above and set Uniform priors on each individual correlation, then the prior density of the two focal correlations at 0 equals 0.25. This corresponds to a log-density of -1.39 . In contrast, the true joint prior density

on the two correlations of the full model (with the two focal correlations free, respecting the positive definite constraint) can be approximated using the positive definite correlation matrices that we randomly generated earlier. This approximation involves a density estimation method from the *ks* package in R (Duong, 2022). Our approximate joint density at 0 is now 0.46, corresponding to a log-density of -0.78 . And we have an additional correction term to account for the fact that the priors for the γ_4 – γ_8 and for the γ_6 – γ_8 correlations are impacted by whether or not the γ_2 residual correlations are fixed to 0.

These evaluations lead to two separate Bayes factors for the model with correlations, relative to the model without correlations. Using our incorrect priors that do not account for positive definite constraints, we obtain a log-Bayes factor of 4.66 in favor of the model with correlations. Using our priors that do account for positive definite constraints, along with the correction from Heck (2019), we obtain a log-Bayes factor of 5.54 in favor of the model with correlations. Both of these Bayes factors provide support for the model with correlations, but possibly at different levels of evidence. For example, if we subscribe to the rules of thumb offered by Kass and Raftery (1995), then these two Bayes factors lead us to conclude “strong” evidence using the incorrect prior calculation, and “very strong” evidence using the correct calculation.

We agree with you, the reader, that these cutoffs are arbitrary and that the Bayes factors do not differ by very much. But the point is that the Bayes factor systematically differ depending on how we compute prior densities. These differences will sometimes lead to different substantive conclusions in practice, with the easier computation (i.e., ignoring the positive definite constraint) being incorrect.

Solutions

While it was straightforward to visualize implied prior distributions in the Bollen model, the process becomes inefficient for correlation matrices whose dimension is larger than 3 or 4. For such correlation matrices, few of the randomly-generated matrices will be positive definite, and it will take a long time to obtain a sufficient number of positive definite matrices to describe the implied prior. Additionally, every unique structure of correlation matrix will have a unique positive definite constraint, similar to the one from Equation (1). So we desire more general solutions that do not involve random generation of correlation matrices. We describe three solutions below that differ in complexity and in the extent to which they fully solve the problem.

Informative Priors

The simplest, partial solution is to maintain the univariate priors on individual correlations, but make those priors informative around 0. For example, instead of placing Uniform priors on the correlations, we might use Beta distributed priors. The Beta distribution is typically defined for the interval (0, 1), but we can transform it so that

the interval $(-1, 1)$ is supported (and note that the Uniform is a special case of the Beta). To make this distribution more informative around 0, we can increase the value of both shape hyperparameters, for example Beta(5, 5). The implied prior distributions will then be closer to what the user specifies, because these informative priors will more often lead to positive definite correlation matrices. But this is not a full solution because, even with informative priors, we may still encounter non-positive definite correlation matrices. And it is not straightforward to predict when or how often this will happen. Additionally, depending on one’s application, certain informative priors may be inappropriate.

Priors on Cholesky Decomposition

A more general solution to this problem comes from putting priors on the Cholesky decomposition of the correlation matrix, which is related to the Ghosh et al. (2021) approach for time series models. The Cholesky decomposition is advantageous because, to ensure that the correlation matrix stays positive definite, we only have to ensure that the diagonal elements of the Cholesky decomposition are positive.

For the Political Democracy model, the 4×4 covariance matrix from the bottom right corner of Figure 3 has a Cholesky decomposition with structure

$$\begin{pmatrix} c_{11} & & & & \\ c_{21} & c_{22} & & & \\ c_{31} & -c_{21}c_{31}/c_{22} & c_{33} & & \\ 0 & c_{42} & c_{43} & c_{44} & \end{pmatrix}, \tag{2}$$

where the entries $\{c_{11}, c_{22}, c_{33}, c_{44}\}$ are constrained to be positive, while $\{c_{21}, c_{31}, c_{42}, c_{43}\}$ are unconstrained. Additionally, the entry in row 3, column 2 is fully determined by other entries. If we place gamma priors (say) on the diagonal entries and normal priors on the remaining c variates, we can maintain the desired structure of the covariance matrix while also maintaining positive definiteness.

A disadvantage of this approach is that the entries of the Cholesky decomposition do not necessarily have intuitive interpretations, so that it is difficult to set informative priors. Each diagonal entry is related to the portion of the corresponding variable’s variance that cannot be accounted for by variables that occur further to the left of the matrix. Each off-diagonal entry is related to a partial correlation conditioned on variables further to the left of the matrix (see Joe, 2006; Lewandowski et al., 2009; Pourahmadi et al., 2007). This implies that the order of the variables matters. A further difficulty is that there does not appear to be an automatic way to obtain a Cholesky structure, like that of Equation (2), for arbitrarily-structured covariance matrices.

Combining LKJ with Informative Priors

A final solution was recently described in a blog post by [Martin \(2021\)](#) and can be implemented in Stan. This solution involves use of a prior distribution that is the product of (i) a Lewandowski-Kurowicka-Joe (LKJ) prior ([Lewandowski et al., 2009](#)) on the full correlation matrix, and (ii) informative priors on individual entries of the correlation matrix. The resulting prior inherits the positive definiteness from (i) while also inheriting the informativeness from (ii). At the moment, it is not clear that this approach can be used to fix individual entries of the correlation matrix; instead, [Martin \(2021\)](#) recommends highly-informative priors around the fixed value that is desired (similar to the notion of “approximate zeros” in a factor loading matrix). The highly-informative priors may be sufficient to replace some hard constraints, but they may also sometimes cause problems with convergence of the MCMC chains.

Sign Indeterminacies

We now turn to a problem that is more specific to SEM: sign indeterminacies of loading parameters. It is well known that, if we change the signs of all loadings, the SEM likelihood (usually) stays the same. To avoid this issue, SEM software typically “prefers” positive loadings through various aspects of implementation. First, for both Bayesian and frequentist models, the loadings’ starting values are often set to positive numbers. Additionally, if a single loading is fixed for identification, it is almost always fixed to +1. This often leads other loadings towards positive values.

Especially when using software like JAGS or Stan, researchers commonly fix the latent variance to 1 and place truncated normal priors on the factor loadings, where the distributions are truncated from below at 0 (e.g., [Curtis, 2010](#)). This forces all loadings to be positive and resolves sign indeterminacies in the model. But this solution is problematic because it does not allow for indicators with “bad” loadings (whose posterior distributions overlap with zero), and it does not allow for reversed indicators (whose valence is opposite that of other indicators). [Peeters \(2012\)](#) shows that, to achieve parameter identification of the likelihood, only one loading per latent variable must be sign restricted (with the latent variance being fixed to 1). Thus, fixing the signs of all loadings is overly restrictive from a parameter identification standpoint.

When a single loading per factor is fixed to 1 for identification, we should not need to fix the signs of any other loadings. If we instead fix the latent variance to 1 for identification, then an improved solution (over fixing all signs to positive) is to employ relabeling algorithms (e.g., [Erosheva & Curtis, 2017](#)). Under this approach, we allow factor loadings to flip between positive and negative values during MCMC estimation. Then, after model estimation, we change the signs of loadings depending upon the signs of some focal loading parameters (and, if the model includes factor correlations or factor

regressions, we also may need to change those signs). This strategy leads to positive loading values, while allowing for the possibility that some loadings are negative.

The relabeling algorithms' preferences for positive loadings can conflict with researchers' desires to use noninformative prior distributions for the loadings (say, Normal with a mean of 0 and a large variance). That is, the software's preference for positive loadings conflicts with the noninformative prior distributions, which state that both positive and negative loadings are equally likely. More generally, factor loadings are influenced by the model identification constraints (e.g., [Bollen et al., 2022](#)). To set meaningful prior distributions on factor loadings, researchers need to consider the specific identification constraints that will be used.

Illustration

To illustrate the interaction between sign indeterminacy and prior distributions, it is sufficient to consider the usual confirmatory factor model that is fit to the [Holzinger and Swineford \(1939\)](#) data. We suspect that most people reading this far know the dataset, which contains scores on various tests of mental ability. We are focusing on the 3-factor model that is traditionally fit to the version of the data from *lavaan* ([Rosseel, 2012](#)), where each factor is associated with 3 observed variables.

The Holzinger-Swineford factor model has five types of model parameters: intercepts, loadings, factor standard deviations, factor correlations, and residual standard deviations. We assign true ("population") values to all these parameters. Intercepts receive true values of 0, factor standard deviations receive true values of 1, factor correlations receive true values of 0, and residual standard deviations receive true values of 1. Finally and importantly, loadings receive true values of -1 .

Using these true values, we generated a dataset of 1,000 observations and re-fit the 3-factor model back to the data (where true values were treated as unknown). We used common, non-informative priors for the parameters of the estimated model, which are currently the defaults in *blavaan* ([Merkle et al., 2021](#); [Merkle & Rosseel, 2018](#)).


$$\begin{aligned} \text{Intercept} &\sim N(0, 1000) \\ \text{Loading} &\sim N(0, 100) \\ \text{Latent covariance matrix} &\sim \text{LKJ}(1) \\ \text{Residual SD} &\sim \text{Gamma}(1, .5), \end{aligned}$$

where the Normal distributions are parameterized with variances, and where the LKJ prior is placed on the entire latent covariance matrix at once and respects the positive definite constraints described earlier. To identify the model, we fixed each latent variance to 1. In this case, *blavaan* uses a relabeling algorithm to handle sign indeterminacy. The estimation had three chains, 500 warmup samples per chain, and 1,000 posterior samples per chain.

The resulting posterior distributions of the loadings appear in Figure 4. These distributions are centered near +1, with the distributions being fully on the positive side of the space. So we have a situation where the true loading values were -1 , we used noninformative priors that exhibit little influence on the posterior, and the posterior distributions of loading values are nowhere near the true values.

Figure 4

Estimated Posterior Distributions of Loadings, Holzinger-Swineford 3-Factor Model



At this point, readers might object that this just illustrates sign indeterminacy. And posterior inferences about factor loadings are not generally impacted here. But the example highlights that, for loadings, a noninformative prior centered at 0 usually ignores the identification constraint that was chosen. That is, when employing noninformative priors, we are usually attempting to say that we have no idea about the loadings' values, or to avoid influencing the results of the model estimation. But the prior ignores the fact that (i) we typically fix a loading to be positive for identification, and (ii) observed variables are usually positively correlated, so that we can expect other loadings to be positive. So our original priors, which were intended to be *noninformative*, may actually state that an implausible part of the parameter space is plausible (also see Gelman et al., 2017; Seaman et al., 2012). Further, as mentioned earlier, the prior distribution of relabelled factor loadings generally differs from the prior distribution of un-relabelled factor loadings, which can impact post-estimation computations that rely on evaluation of the prior.

Implications for MCMC Validation

Sign indeterminacy also complicates MCMC algorithm validation, which is used to ensure that MCMC samplers are working correctly. The MCMC validation process is difficult even without sign indeterminacy, because randomness is inherent in MCMC sampling. This means that we cannot simply examine whether the posterior means and standard deviations match other samplers to many decimal places. While it is possible to obtain analytic posterior distributions for certain models, analytic results are the exception instead of the rule for models estimated via MCMC.

Simulation-Based Calibration (SBC; [Modrák et al., 2022](#); [Talts et al., 2018](#)) is an MCMC validation method that has received recent attention and implementation. Given a model of interest (including likelihood and priors), simulation-based calibration can be described in four steps:

1. Generate many sets of parameter values from the prior distribution.
2. For each set of parameters from Step 1, generate an artificial dataset.
3. Fit the model of interest to each artificial dataset from Step 2.
4. Examine whether the resulting posterior distributions look like the prior distribution.

If the MCMC sampler is working correctly, then the posteriors from Step 3 should look like the priors from which we started. We can graphically examine this idea by comparing the posterior means from Step 3 to the parameter values from Step 1; we should see an identity line when plotting the parameter values against the posterior means.


Using the same model from the previous section, we used the *SBC* package ([Kim et al., 2022](#)) to conduct simulation-based calibration under two sets of priors, with 1000 simulated data sets each. Set 1 was exactly the same as the noninformative priors from the previous section. Set 2 was also similar to the previous section, differing only in the priors for the loadings: instead of $\text{Normal}(0, 100)$, the priors for loadings were $\text{Normal}(1, 1/16)$ (where the second number is a variance). This is an informative prior reflecting the belief that all loadings should be similar to one another (recall that a single loading is being fixed to 1 for identification).

Results for the Set 1 and Set 2 priors are shown in [Figure 5](#) and [Figure 6](#), respectively. In both figures, the parameter values simulated from the prior are on the x-axis, and the posterior means estimated from the artificial data are on the y-axis. Each panel represents a factor loading (there are 6 free loadings in the model). Each point represents a replication, and “correct” MCMC algorithms should lead to points that are crowded around the blue diagonal line.

[Figure 5](#) shows that, for noninformative priors on factor loadings, the points follow a V or an X pattern. The V pattern occurs for loadings that are constrained to be positive during estimation, while the X pattern occurs for the remaining loadings. This means

Figure 5

Prior Parameter Values Versus Posterior Means for $N(0,100)$ Priors on Loadings




that the estimation algorithm often recovers the true parameter value, but it also often flips the sign of the parameter value. While this is not a problem if we realize the sign indeterminacy issue, it is easy to overlook in the context of simulation-based calibration. For example, typical SBC summaries involve cumulative distributions, which would lead us to conclude that there are problems with the MCMC algorithm, and which would not provide clear clues about sign indeterminacy. See [Merkle et al. \(2021\)](#) for some discussion of similar analyses.

[Figure 6](#), on the other hand, is closer to what we would hope to see from a correct MCMC algorithm. In this figure, the points generally form a cloud around the blue diagonal line, implying that the posterior means match the parameter values that generated the data (note that the axis limits have changed, as compared to the previous figure). The informative priors on factor loadings work similarly to use of truncated normal priors, where the signs of the loadings were generally restricted to be positive. But informative priors differ from the truncated normals in that they allow for the possibility of negative loadings.

Figure 6

Prior Parameter Values Versus Posterior Means for $N(1, \frac{1}{16})$ Priors on Loadings



Solutions

If researchers are aware of how their software handles sign indeterminacy, then they can potentially avoid the issues described here and successfully employ noninformative prior distributions. Barring that, we advise researchers to explicitly consider the loadings’ expected signs when setting prior distributions. In many models, we expect that all the observed variables corresponding to a factor will have the same direction of relationship with that factor. Additionally, a single loading is often set to 1 for identification. In a situation like this, it is often reasonable to place priors on the free loadings that have a mean of 1 and a standard deviation of, say, .5. These priors look very informative at first glance, as compared to, say, a Normal prior with a mean of 0 and a variance of 10,000. But the suggested priors better represent what the researcher knows about the signs of the loadings, combined with the fact that some loadings are being fixed to 1.

Instead of fixing a single loading to 1 for identification, researchers may fix the latent variance to 1. In this case, as described by Peeters (2012), one loading per latent variable must be constrained to be positive (or negative) in order to achieve identification. If that loading has a Normal prior whose support includes the negative (positive) reals, then we

again face a conflict where our stated prior is not the implied prior. The implied prior for the sign-constrained loading may become truncated normal or truncated t , depending on other details of the MCMC implementation (see Ghosh & Dunson, 2009), while the priors on remaining loadings are as stated. While a truncated distribution arises from the identification constraint here, we think that truncated priors on all loadings should be avoided, because they prevent the loadings' posterior distributions from overlapping with zero. The relabeling strategies described by Erosheva and Curtis (2017) should be preferred.

There are alternative prior distributions that avoid the issue and/or that make it easier to specify informative prior distributions, though they are not readily available in popular software. Graves and Merkle (2022) studied prior distributions on ratios of factor loadings, as well as prior specification under effect coding (Little et al., 2006). Ratios of loadings avoid the sign-switching issue (e.g., the ratio of two negative loadings remains positive), and effect coding can make it easier to specify informative prior distributions. Additionally, in the case of exploratory factor analysis, Heaps (2022) recently proposed a prior distribution for the common factor covariance matrix. That is, factor analysis models typically imply a covariance matrix of the form $\Lambda\Lambda' + \Psi$, where Λ is the factor loading matrix with many fewer columns than rows. Heaps (2022) proposes to place a matrix normal prior on $\Lambda\Lambda'$, which encodes researcher knowledge about shared variation in the observed variables. The term $\Lambda\Lambda'$ remains invariant across factor loadings' signs and rotations, so it avoids the issues described here. The proposed priors can also shrink sets of loadings towards zero, which is similar to confirmatory factor analysis. It remains to be seen how this prior could be translated to more general SEMs, where the model-implied covariance matrix becomes more complex.

Order Constraints

Finally, we describe prior distributions for order-constrained parameters, which are commonly seen in SEMs for ordinal variables with more than two categories. For these models, there exist threshold parameters that chop each underlying continuous variable into observed, ordered categories. The threshold parameters must be ordered so that they correspond to the ordering of the observed variables. For example, the lowest threshold chops off the lowest category, the second threshold chops off the bottom two categories, and so on.

The prior distributions for threshold parameters are often opaque, because the priors that researchers specify often have no order constraints. This is commonly done to improve the software's ease of use: researchers are accustomed to setting univariate Normal priors on individual parameters, and the priors with order constraints typically do not have simple forms. But the software always imposes order constraints here, which changes the prior distribution in various manners. Researchers often do not realize

that anything happened, which may be especially problematic when setting informative priors.

Example

Say that a researcher fits a factor analysis model to a set of 4-category ordinal variables, and that she specifies a Normal(0,5) prior on all threshold parameters in the model. Because there are four categories per variable, we require three order-constrained thresholds per variable. We wish to know what these priors look like, after accounting for the order constraints.

We consider two ways that we could translate Normal(0,5) priors to three ordered parameters (also see Padgett et al., 2023). First, we could imagine drawing three separate variates from the normal distribution, then ordering the variates to obtain ordered thresholds. Second, we could imagine drawing the first (lowest) threshold from a Normal(0,5), then adding a Lognormal(0,5) variate to that threshold in order to obtain the second threshold. Once we have obtained the second threshold, we could add another Lognormal variate to obtain the third threshold. Lognormal variables can only take positive values, so we are guaranteed to have an ordered set of thresholds under this approach.

For both of these translations, the threshold parameters' prior distributions differ from the Normal(0,5) distribution that the researcher originally declared. We expand on this point below, separately for the two methods.

Reordering

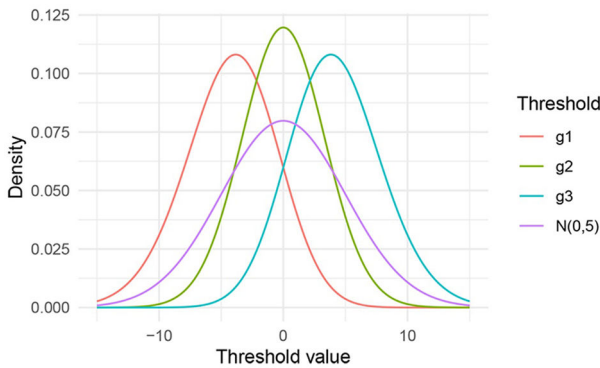
When we draw three values from the Normal distribution and then order them, the act of ordering influences the resulting prior distributions. The specific distributions can be described via statistical theory on order statistics. For our example, the Normal(0,5) priors translate into the following probability density functions (pdfs) for individual thresholds:

$$\begin{aligned} p(g_1) &= 3 \times \phi(g_1/5) \times [1 - \Phi(g_1/5)]^2 \\ p(g_2) &= 6 \times \phi(g_2/5) \times \Phi(g_2/5) \times [1 - \Phi(g_1/5)] \\ p(g_3) &= 3 \times \phi(g_2/5) \times \Phi(g_2/5)^2, \end{aligned}$$

where $\phi()$ is the standard normal pdf and $\Phi()$ is the standard normal cdf. These distributions are visualized in Figure 7, with the Normal(0,5) distribution overlaid for comparison. The figure shows that the prior for the lower threshold (g_1) is centered below 0, while the prior for the upper threshold (g_3) is centered above 0. None of the three distributions matches the Normal(0,5), despite the fact that the researcher declared a Normal(0,5) prior for all three parameters.

Figure 7

Priors of Threshold Parameters Whose Stated Priors are Normal(0,5), as Implied by Order Constraints



Lognormals

The *blavaan* package uses the lognormal approach mentioned previously. The Normal(0,5) prior goes on the lowest threshold parameter per item. Then, Normal(0,5) priors are placed on log-differences between subsequent parameters. For example, considering the three thresholds from the example, we would have

$$\begin{aligned} g_1 &\sim \text{Normal}(0, 5) \\ \log(g_2 - g_1) &\sim \text{Normal}(0, 5) \\ \log(g_3 - g_2) &\sim \text{Normal}(0, 5), \end{aligned}$$

where we could alternatively say that the differences between thresholds follow Lognormal priors.

Just like the previous section, the above priors can be translated to priors on individual thresholds. It is obvious that the prior for g_1 continues to be the stated prior, which is Normal(0,5). But the priors for g_2 and g_3 involve the sum of a normal distribution and Lognormal distribution(s). The resulting distributions can be written as

$$\begin{aligned} p(g_2) &= \int_{-\infty}^{g_2} \phi((g_1 - \mu)/\sigma) \times \log N(g_2 - g_1, \mu, \sigma) \partial g_1 \\ p(g_3) &= \int_{-\infty}^{g_3} \int_{-\infty}^{g_2} \phi((g_1 - \mu)/\sigma) \times \log N(g_2 - g_1, \mu, \sigma) \times \log N(g_3 - g_2, \mu, \sigma) \partial g_1 \partial g_2, \end{aligned}$$


where $\log N(x, \mu, \sigma)$ is the density function of the Lognormal distribution with mean μ and standard deviation σ (both on the log scale), evaluated at x .

While the priors for g_2 and g_3 do not have nice forms, we can numerically approximate the integrals to visualize them. These distributions are shown in Figure 8, which is

arranged similarly to Figure 7. As stated before, the lowest threshold, g_1 , has the stated Normal(0,5) distribution. The peaks of the distributions of the remaining thresholds are closer to that of g_1 , as compared to the distributions in Figure 7. The distributions of g_2 and g_3 are also skewed to the right, reflecting the fact that we are adding a positive variate to the distribution of g_1 .

Figure 8

Priors of Threshold Parameters Whose Stated Priors are Normal(0,5), as Implied by Placing Priors on Log-Differences



Solution

Unlike the previous issues with positive definite constraints and sign constraints, researchers do not have to consider changing their priors in order to address order constraints. The main solution is to be aware of the fact that, if one places univariate prior distributions on a set of order-constrained parameters, then some translation will take place to ensure that the parameters are ordered correctly. And this translation will influence the implied prior distribution of each parameter. It is worthwhile to understand how this is handled by one’s software, especially for prior predictive assessments, Bayes factor calculation, and simulation-based calibration methods.

General Discussion

In this paper, we considered the idea of *opaque* prior distributions in Bayesian models with latent variables, where model estimation and prior specification interact to potentially yield unexpected results. The problem arises from the fact that software implementations must respect various model constraints, while researchers' prior specifications do not always respect the same constraints. This leads researchers to declare one set of prior distributions, which can imply a different set of prior distributions depending on the software implementation. The issue is particularly problematic for software verification, for comparing model results across different pieces of software, and for computing metrics that explicitly involve evaluation of the prior distribution.

The three issues that we considered were (i) positive definite constraints on model covariance matrices; (ii) sign indeterminacy and constraints used to identify model parameters (typically factor loadings); and (iii) order constraints on subsets of model parameters (typically thresholds/intercepts). These issues occur with different frequencies, with (ii) and (iii) occurring more often than (i). To expand on this, issue (iii) occurs for most models that have ordinal variables with more than two categories, issue (ii) occurs for most measurement models (with free loadings), and issue (i) occurs for models with residual covariances, or other combinations of fixed and free covariances. We could have a worst-case scenario, such as a multiple group model with ordinal variables and across-group parameter constraints, where all three issues occur at once.

To avoid problems associated with opaque priors, we offer the following recommendations for practice:

1. If one's model involves covariance matrices without parameter constraints, use a single prior for the full covariance (or correlation) matrix (LKJ, inverse Wishart, etc).
2. If one's model involves a covariance matrix with parameter constraints, consider putting a prior on the Cholesky decomposition, or use matrix identities to see whether the full matrix can be broken into blocks that are easier to handle. If these are unavailable, use informative priors on the correlations that place more density close to 0.
3. For factor loadings, consider the expected direction of the relationship between each observed variable and the corresponding latent variable(s), along with the loading identification constraints. Use priors that place most density in this expected direction.
4. Be aware of how order constraints influence priors for thresholds, especially if one is doing model assessments that directly involve prior evaluation.

Out of these recommendations, the priors on constrained covariance matrices are most difficult to handle. Future work could make it easier for researchers to place reasonable priors on constrained covariance matrices.

Importantly, the issue of opaque priors does *not* mean that all results associated with affected models are wrong. In the presence of opaque priors, we can still obtain accurate posterior summaries, including model information criteria (like WAIC and LOO) and some fit indices (e.g., [Garnier-Villarreal & Jorgensen, 2020](#)). On the other hand, the prior distributions that researchers describe in their papers may be incorrect, as will prior predictive checks and other model summaries that directly rely on prior distribution evaluation, such as Bayes factors. Researchers should especially be careful about applying Bayes factor computation strategies (e.g., [Gronau et al., 2020](#)) to latent variable models, to ensure that they are evaluating the correct priors.

Opaque priors are vaguely similar to applied modeling of ordinal variables (e.g., [Bürkner & Vuorre, 2019](#); [Liddell & Kruschke, 2018](#)), where researchers ignore the fact that they have ordinal variables, treat them as continuous, and sometimes obtain reasonable results. Similarly, researchers can ignore the fact that they have opaque priors, estimate their model, and sometimes obtain reasonable results. In both cases, it is difficult to predict exactly when the results will be reasonable and when they will not. And ignoring the issues do not make them disappear.

We conclude by considering that some non-Bayesian researchers may find this paper appealing, because they can use it to justify phrases like “Bayesian methods are difficult to use.” We agree that priors present extra complications that do not exist for other methods, but we find the extra complications to be worthwhile. In our experience, wrestling with prior distributions can lead to a deeper, more sober understanding of one’s model and how it interacts with data. This understanding might be achieved via other, non-Bayesian routes, but it will require the time and effort that Bayesians devote to prior distributions.

Funding: This work was supported by the Institute of Education Sciences, United States Department of Education, Grant R305D210044 to the University of Missouri.

Acknowledgments: We thank Bob Carpenter and two reviewers for comments that improved the paper. All remaining errors are due to the authors.

Competing Interests: The authors have declared that no competing interests exist.

Computational Details: All results were obtained using the R system for statistical computing ([R Core Team, 2022](#)), version 4.3.1, with major reliance on packages *blavaan* ([Merkle et al., 2021](#)), *rstan* ([Stan Development Team, 2022](#)), and *SBC* ([Kim et al., 2022](#)).

Data Availability: The R codes to replicate the results of the current study are freely available and can be found in the [Supplementary Materials](#).

Supplementary Materials

The supplementary materials provided are the R code scripts to replicate the results of the current study (see [Merkle et al., 2023](#)).

Index of Supplementary Materials

Merkle, E. C., Ariyo, O., Winter, S. D., & Garnier-Villarreal, M. (2023). *Supplementary materials to "Opaque prior distributions in Bayesian latent variable models"* [R code scripts]. PsychOpen GOLD. <https://doi.org/10.23668/psycharchives.13256>

References

- Alvarez, I., Niemi, J., & Simpson, M. (2014). *Bayesian inference for a covariance matrix*. arXiv. <https://doi.org/10.48550/arXiv.1408.4050>
- Ariyo, O., Lesaffre, E., Verbeke, G., & Quintero, A. (2022). Model selection for Bayesian linear mixed models with longitudinal data: Sensitivity to the choice of priors. *Communications in Statistics – Simulation and Computation*, 51(4), 1591–1615. <https://doi.org/10.1080/03610918.2019.1676439>
- Asparouhov, T., Hamaker, E. L., & Muthén, B. (2018). Dynamic structural equation models. *Structural Equation Modeling: A Multidisciplinary Journal*, 25(3), 359–388. <https://doi.org/10.1080/10705511.2017.1406803>
- Bainter, S. A. (2017). Bayesian estimation for item factor analysis models with sparse categorical indicators. *Multivariate Behavioral Research*, 52(5), 593–615. <https://doi.org/10.1080/00273171.2017.1342203>
- Barnard, J., McCulloch, R., & Meng, X.-L. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica*, 10(4), 1281–1311.
- Bollen, K. A. (1989). *Structural equations with latent variables*. John Wiley & Sons.
- Bollen, K. A., Lilly, A. G., & Luo, L. (2022). Selecting scaling indicators in structural equation models (SEMs). *Psychological Methods*. Advance online publication. <https://doi.org/10.1037/met0000530>
- Bürkner, P.-C., & Vuorre, M. (2019). Ordinal regression models in psychology: A tutorial. *Advances in Methods and Practices in Psychological Science*, 2(1), 77–101. <https://doi.org/10.1177/2515245918823199>
- Bürkner, P.-C. (2021). Bayesian item response modeling in R with brms and Stan. *Journal of Statistical Software*, 100(5), 1–54. <https://doi.org/10.18637/jss.v100.i05>
- Curtis, S. M. (2010). BUGS code for item response theory. *Journal of Statistical Software*, 36(1), 1–34. <https://doi.org/10.18637/jss.v036.c01>
- Cuthill, E., & McKee, J. (1969). Reducing the bandwidth of sparse symmetric matrices. In *Proceedings of the 24th Association for Computing Machinery National Conference* (pp. 157–172). Association for Computing Machinery. <https://doi.org/10.1145/800195.805928>

- Daniels, M. J., & Kass, R. E. (1999). Nonconjugate Bayesian estimation of covariance matrices and its use in hierarchical models. *Journal of the American Statistical Association*, 94(448), 1254–1263. <https://doi.org/10.2307/2669939>
- Depaoli, S., Lai, K., & Yang, Y. (2021). Bayesian model averaging as an alternative to model selection for multilevel models. *Multivariate Behavioral Research*, 56(6), 920–940. <https://doi.org/10.1080/00273171.2020.1778439>
- Dickey, J. M., & Lientz, B. P. (1970). The weighted likelihood ratio, sharp hypotheses about chances, the order of a Markov chain. *Annals of Mathematical Statistics*, 41(1), 214–226. <https://doi.org/10.1214/aoms/1177697203>
- dos Santos, J. R. S., Azevedo, C. L. N., & Fox, J.-P. (2022). Bayesian longitudinal item response modeling with multivariate asymmetric serial dependencies. *Journal of Statistical Computation and Simulation*, 92(3), 488–523. <https://doi.org/10.1080/00949655.2021.1965604>
- Driver, C. C., Oud, J. H. L., & Voelkle, M. C. (2017). Continuous time structural equation modeling with R package ctsem. *Journal of Statistical Software*, 77(5), 1–35. <https://doi.org/10.18637/jss.v077.i05>
- Duong, T. (2022). ks: Kernel smoothing [Computer software manual]. R package version 1.14.0. <https://CRAN.R-project.org/package=ks>.
- Enders, C. K., Keller, B. T., & Levy, R. (2018). A fully conditional specification approach to multilevel imputation of categorical and continuous variables. *Psychological Methods*, 23(2), 298–317. <https://doi.org/10.1037/met0000148>
- Erosheva, E. A., & Curtis, S. M. (2017). Dealing with reflection invariance in Bayesian factor analysis. *Psychometrika*, 82(2), 295–307. <https://doi.org/10.1007/s11336-017-9564-y>
- Garnier-Villarreál, M., & Jorgensen, T. D. (2020). Adapting fit indices for Bayesian structural equation modeling: Comparison to maximum likelihood. *Psychological Methods*, 25(1), 46–70. <https://doi.org/10.1037/met0000224>
- Garnier-Villarreál, M., Merkle, E. C., & Magnus, B. E. (2021). Between-item multidimensional IRT: How far can the estimation methods go? *Psych*, 3(3), 404–421. <https://doi.org/10.3390/psych3030029>
- Gelman, A., Simpson, D., & Betancourt, M. (2017). The prior can often only be understood in the context of the likelihood. *Entropy*, 19(10), Article 555. <https://doi.org/10.3390/e19100555>
- Ghosh, J., & Dunson, D. B. (2009). Default prior distributions and efficient posterior computation in Bayesian factor analysis. *Journal of Computational and Graphical Statistics*, 18(2), 306–320. <https://doi.org/10.1198/jcgs.2009.07145>
- Ghosh, R. P., Mallick, B., & Pourahmadi, M. (2021). Bayesian estimation of correlation matrices of longitudinal data. *Bayesian Analysis*, 16(3), 1039–1058. <https://doi.org/10.1214/20-BA1237>
- Graves, B., & Merkle, E. C. (2022). A note on identification constraints and information criteria in Bayesian latent variable models. *Behavior Research Methods*, 54(2), 795–804. <https://doi.org/10.3758/s13428-021-01649-8>

- Gronau, Q. F., Singmann, H., & Wagenmakers, E.-J. (2020). bridgesampling: An R package for estimating normalizing constants. *Journal of Statistical Software*, *92*(10), 1–29. <https://doi.org/10.18637/jss.v092.i10>
- Haaf, J. M., Merkle, E. C., & Rouder, J. N. (2020). Do items order? The psychology in IRT models. *Journal of Mathematical Psychology*, *98*, Article 102398. <https://doi.org/10.1016/j.jmp.2020.102398>
- Haaf, J. M., & Rouder, J. N. (2023). Does every study? Implementing ordinal constraint in meta-analysis. *Psychological Methods*, *28*(2), 472–487. <https://doi.org/10.1037/met0000428>
- Heaps, S. E. (2022). *Structured prior distributions for the covariance matrix in latent factor models*. ArXiv. <https://doi.org/10.48550/arXiv.2208.07831>
- Heck, D. W. (2019). A caveat on the Savage–Dickey density ratio: The case of computing Bayes factors for regression parameters. *British Journal of Mathematical and Statistical Psychology*, *72*(2), 316–333. <https://doi.org/10.1111/bmsp.12150>
- Heck, D. W., Boehm, U., Böing-Messing, F., Bürkner, P.-C., Derks, K., Dienes, Z., Dienes, Z., Fu, Q., Gu, XZ., Karimova, D., Kiers, H. A. L., Klugkist, I., Kuiper, R. M., Lee, M. D., Leenders, R., Lepplaa, H. J., Linde, M., Ly, A., Meijerink-Bosman, M., . . . Hoijsink, H. (2023). A review of applications of the Bayes factor in psychological research. *Psychological Methods*, *28*(3), 558–579. <https://doi.org/10.1037/met0000454>
- Hoijsink, H., & van de Schoot, R. (2018). Testing small variance priors using prior-posterior predictive p values. *Psychological Methods*, *23*(3), 561–569. <https://doi.org/10.1037/met0000131>
- Holzinger, K. J., & Swineford, F. A. (1939). *A study of factor analysis: The stability of a bi-factor solution* (No. 48). University of Chicago Press.
- Huang, A., & Wand, M. P., (2013). Simple marginally noninformative prior distributions for covariance matrices. *Bayesian Analysis*, *8*(2), 439–452. <https://doi.org/10.1214/13-BA815>
- Hurtado Rúa, S. M., Mazumdar, M., & Strawderman, R. L. (2015). The choice of prior distribution for a covariance matrix in multivariate meta-analysis: A simulation study. *Statistics in Medicine*, *34*(30), 4083–4104. <https://doi.org/10.1002/sim.6631>
- Joe, H. (2006). Generating random correlation matrices based on partial correlations. *Journal of Multivariate Analysis*, *97*(10), 2177–2189. <https://doi.org/10.1016/j.jmva.2005.05.010>
- Jorgensen, T. D., & Garnier-Villarreal, M. (2023). Limited utility of small-variance priors to detect local misspecification in Bayesian structural equation models. In M. Wiberg, D. Molenaar, J. González, J.-S. Kim, & H. Hwang (Eds.), *Quantitative psychology* (Vol. 422, pp. 85–95). Springer Nature.
- Kaplan, D., Chen, J., Yavuz, S., & Lyu, W. (2023). Bayesian dynamic borrowing of historical information with applications to the analysis of large-scale assessments. *Psychometrika*, *88*(1), 1–30. <https://doi.org/10.1007/s11336-022-09869-3>
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773–795. <https://doi.org/10.1080/01621459.1995.10476572>

- Keller, B. T., & Enders, C. K. (2023). An investigation of factored regression missing data methods for multilevel models with cross-level interactions. *Multivariate Behavioral Research*. <https://doi.org/10.1080/00273171.2022.2147049>
- Kim, S., Moon, H., Modrák, M., & Säilynoja, T. (2022). *SBC: Simulation based calibration for rstan/cmdstanr models [Computer software manual]*. GitHub. <https://hyunjimoon.github.io/SBC/>
- Levy, R., & Enders, C. K. (2021). Full conditional distributions for Bayesian multilevel models with additive or interactive effects and missing data on covariates. *Communications in Statistics – Simulation and Computation*, 52(7), 2899–2923. <https://doi.org/10.1080/03610918.2021.1921799>
- Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9), 1989–2001. <https://doi.org/10.1016/j.jmva.2009.04.008>
- Liddell, T. M., & Kruschke, J. K. (2018). Analyzing ordinal data with metric models: What could possibly go wrong? *Journal of Experimental Social Psychology*, 79, 328–348. <https://doi.org/10.1016/j.jesp.2018.08.009>
- Little, T. D., Slegers, D. W., & Card, N. A. (2006). A non-arbitrary method of identifying and scaling latent variables in SEM and MACS models. *Structural Equation Modeling: A Multidisciplinary Journal*, 13(1), 59–72. https://doi.org/10.1207/s15328007sem1301_3
- Liu, H., Depaoli, S., & Marvin, L. (2022). Understanding the deviance information criterion for SEM: Cautions in prior specification. *Structural Equation Modeling: A Multidisciplinary Journal*, 29(2), 278–294. <https://doi.org/10.1080/10705511.2021.1994407>
- Lüdtke, O., Ulitzsch, E., & Robitzsch, A. (2021). A comparison of penalized maximum likelihood estimation and markov chain Monte Carlo techniques for estimating confirmatory factor analysis models with small sample sizes. *Frontiers in Psychology*, 12, Article 615162. <https://doi.org/10.3389/fpsyg.2021.615162>
- Magnus, B. E., & Garnier-Villarre, M. (2022). A multidimensional zero-inflated graded response model for ordinal symptom data. *Psychological Methods*, 27(2), 261–279. <https://doi.org/10.1037/met0000395>
- Man, A. X., & Culpepper, S. A. (2022). A mode-jumping algorithm for Bayesian factor analysis. *Journal of the American Statistical Association*, 117(537), 277–290. <https://doi.org/10.1080/01621459.2020.1773833>
- Martin, S. R. (2021). *Informative priors for correlation matrices: An easy approach*. Stephen R. Martin, PhD. <http://smart.in/informative-priors-for-correlation-matrices-an-easy-approach/>
- Merkle, E. C., Fitzsimmons, E., Uanhoro, J., & Goodrich, B. (2021). Efficient Bayesian structural equation modeling in Stan. *Journal of Statistical Software*, 100(6), 1–22. <https://doi.org/10.18637/jss.v100.i06>
- Merkle, E. C., Furr, D., & Rabe-Hesketh, S. (2019). Bayesian model assessment: Use of conditional vs marginal likelihoods. *Psychometrika*, 84(3), 802–829. <https://doi.org/10.1007/s11336-019-09679-0>
- Merkle, E. C., & Rosseel, Y. (2018). *blavaan*: Bayesian structural equation models via parameter expansion. *Journal of Statistical Software*, 85(4), 1–30. <https://doi.org/10.18637/jss.v085.i04>

- Miočević, M., Levy, R., & MacKinnon, D. P. (2021). Different roles of prior distributions in the single mediator model with latent variables. *Multivariate Behavioral Research*, 56(1), 20–40. <https://doi.org/10.1080/00273171.2019.1709405>
- Modrák, M., Moon, A. H., Kim, S., Bürkner, P., Huurre, N., Faltejsková, K., Gelman, A. & Vehtari, A. (2022). *Simulation-based calibration checking for Bayesian computation: The choice of test quantities shapes sensitivity*. ArXiv. <https://doi.org/10.48550/arXiv.2211.02383>
- Padgett, R. N., Morgan, G. B., & Lomas, T. (2023). *Regularizing threshold priors with sparse response patterns in Bayesian factor analysis with categorical indicators*. ArXiv. <https://doi.org/10.48550/arXiv.2307.10503>
- Paganin, S., Paciorek, C. J., Wehrhahn, C., Rodríguez, A., Rabe-Hesketh, S., & de Valpine, P. (2023). Computational strategies and estimation performance with Bayesian semiparametric item response theory models. *Journal of Educational and Behavioral Statistics*, 48(2), 147–188. <https://doi.org/10.3102/10769986221136105>
- Peeters, C. F. W. (2012). Rotational uniqueness conditions under oblique factor correlation metric. *Psychometrika*, 77(2), 288–292. <https://doi.org/10.1007/s11336-012-9259-3>
- Pokropek, A., Schmidt, P., & Davidov, E. (2020). Choosing priors in Bayesian measurement invariance modeling: A Monte Carlo simulation study. *Structural Equation Modeling: A Multidisciplinary Journal*, 27(5), 750–764. <https://doi.org/10.1080/10705511.2019.1703708>
- Pourahmadi, M., Daniels, M. J., & Park, T. (2007). Simultaneous modelling of the Cholesky decomposition of several covariance matrices. *Journal of Multivariate Analysis*, 98(3), 568–587. <https://doi.org/10.1016/j.jmva.2005.11.002>
- Rast, P., Martin, S., Liu, S., & Williams, D. (2022). A new frontier for studying within-person variability: Bayesian multivariate generalized autoregressive conditional heteroskedasticity models. *Psychological Methods*, 27(5), 856–873. <https://doi.org/10.1037/met0000357>
- R Core Team. (2022). *R: A language and environment for statistical computing* [Computer software manual]. R Foundation for Statistical Computing. <https://www.r-project.org/>
- Rosseel, Y. (2012). *lavaan: An R package for structural equation modeling*. *Journal of Statistical Software*, 48(2), 1–36. <http://www.jstatsoft.org/v48/i02/>
- Schmich, F. (2022). *netprioR: A model for network-based prioritisation of genes* [Computer software manual]. R package version 1.24.0. <http://bioconductor.org/packages/netprioR>
- Seaman, J. W., Seaman, J. W., & Stamey, J. D. (2012). Hidden dangers of specifying noninformative priors. *American Statistician*, 66(2), 77–84. <https://doi.org/10.1080/00031305.2012.695938>
- Stan Development Team. (2022). *RStan: The R interface to Stan*. R package version 2.21.7. <https://mc-stan.org/>
- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., & Gelman, A. (2018). *Validating Bayesian inference algorithms with simulation-based calibration*. ArXiv. <https://doi.org/10.48550/arXiv.1804.06788>
- Uanhoro, J. O. (2023a). Hierarchical covariance estimation approach to meta-analytic structural equation modeling. *Structural Equation Modeling: A Multidisciplinary Journal*, 30(4), 532–546. <https://doi.org/10.1080/10705511.2022.2142128>

- Uanhoro, J. O. (2023b). Modeling misspecification as a parameter in Bayesian structural equation models. *Educational and Psychological Measurement*. <https://doi.org/10.1177/00131644231165306>
- Ullrich, E., Lüdtke, O., & Robitzsch, A. (2023). Alleviating estimation problems in small sample structural equation modeling: A comparison of constrained maximum likelihood, Bayesian estimation, and fixed reliability approaches. *Psychological Methods*, 28(3), 527–557. <https://doi.org/10.1037/met0000435>
- Ullrich, E., Pohl, S., Khorramdel, L., Kroehne, U., & von Davier, M. (2022). A response-time-based latent response mixture model for identifying and modeling careless and insufficient effort responding in survey data. *Psychometrika*, 87(2), 593–619. <https://doi.org/10.1007/s11336-021-09817-7>
- Van De Schoot, R., Kluytmans, A., Tummers, L., Lugtig, P., Hox, J., & Muthén, B. (2013). Facing off with Scylla and Charybdis: A comparison of scalar, partial, and the novel possibility of approximate measurement invariance. *Frontiers in Psychology*, 4, Article 770. <https://doi.org/10.3389/fpsyg.2013.00770>
- Van Erp, S., & Browne, W. J. (2021). Bayesian multilevel structural equation modeling: An investigation into robust prior distributions for the doubly latent categorical model. *Structural Equation Modeling: A Multidisciplinary Journal*, 28(6), 875–893. <https://doi.org/10.1080/10705511.2021.1915146>
- Vanpaemel, W. (2020). Strong theory testing using the prior predictive and the data prior. *Psychological Review*, 127(1), 136–145. <https://doi.org/10.1037/rev0000167>
- Van Zundert, C., Somer, E., & Miočević, M. (2022). Prior predictive checks for the method of covariances in Bayesian mediation analysis. *Structural Equation Modeling: A Multidisciplinary Journal*, 29(3), 428–437. <https://doi.org/10.1080/10705511.2021.1977648>
- Verdinelli, I., & Wasserman, L. (1995). Computing Bayes factors using a generalization of the Savage–Dickey density ratio. *Journal of the American Statistical Association*, 90(430), 614–618. <https://doi.org/10.1080/01621459.1995.10476554>
- Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive Psychology*, 60(3), 158–189. <https://doi.org/10.1016/j.cogpsych.2009.12.001>
- Wei, J., & Higgins, J. (2013). Bayesian multivariate meta-analysis with multiple outcomes. *Statistics in Medicine*, 32(17), 2911–2934. <https://doi.org/10.1002/sim.5745>
- Wilcox, K. T., Jacobucci, R., Zhang, Z., & Ammerman, B. A. (2021). *Supervised latent Dirichlet allocation with covariates: A Bayesian structural and measurement model of text and covariates*. PsyArXiv. <https://doi.org/10.31234/osf.io/62tc3>
- Zitzmann, S., Lüdtke, O., Robitzsch, A., & Hecht, M. (2021). On the performance of Bayesian approaches in small samples: A comment on Smid, McNeish, Miočević, and van de Schoot (2020). *Structural Equation Modeling: A Multidisciplinary Journal*, 28(1), 40–50. <https://doi.org/10.1080/10705511.2020.1752216>

Zyphur, M. J., Hamaker, E. L., Tay, L., Voelke, M., Preacher, K. J., Zhang, Z., Allison, P. D., Pierides, D. C., Koval, P., & Diener, E. F. (2021). From data to causes III: Bayesian priors for general cross-lagged panel models (gclm). *Frontiers in Psychology, 12*.

<https://doi.org/10.3389/fpsyg.2021.612251>



Methodology is the official journal of the European Association of Methodology (EAM).



leibniz-psychology.org

PsychOpen GOLD is a publishing service by Leibniz Institute for Psychology (ZPID), Germany.