

Criteria for Parameter Identification in Bayesian Lasso Methods for Covariance Analysis: Comparing Rules for Thresholding, p -value, and Credible Interval

Lijin Zhang, Junhao Pan & Edward Haksing Ip

To cite this article: Lijin Zhang, Junhao Pan & Edward Haksing Ip (2021): Criteria for Parameter Identification in Bayesian Lasso Methods for Covariance Analysis: Comparing Rules for Thresholding, p -value, and Credible Interval, Structural Equation Modeling: A Multidisciplinary Journal, DOI: [10.1080/10705511.2021.1945456](https://doi.org/10.1080/10705511.2021.1945456)

To link to this article: <https://doi.org/10.1080/10705511.2021.1945456>



Published online: 30 Aug 2021.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)



Criteria for Parameter Identification in Bayesian Lasso Methods for Covariance Analysis: Comparing Rules for Thresholding, p -value, and Credible Interval

Lijin Zhang ^a, Junhao Pan ^a, and Edward Haksing Ip ^b

^aSun Yat-sen University; ^bWake Forest School of Medicine

ABSTRACT

The lasso is a commonly used regularization method that is increasing used in structural equation models (SEMs). Under the Bayesian framework, lasso is rendered more flexible and readily produces estimates of standard errors and the penalty parameter. However, in practice, it remains unclear what decision rule is appropriate for parameter identification; in other words, determining what size estimate is large enough to be included into the model. The current study compared three decision rules for parameter identification – thresholding, p -value, and credible interval in confirmatory factor analysis. Specifically, two distinct parameter spaces were studied: cross-loadings and residual correlations. Results showed that the thresholding rule performed best in balancing power and Type I error rate. Different thresholds for standardized estimates were needed for different conditions. Guidelines for parameter identification and recommended thresholding values were also provided. Results of the current study have the potential to extend to a broad range of SEMs.

KEYWORDS

Confirmatory factor analysis; cross-loading; residual correlation; regularization

Introduction

Modern data analysis must often manage a large amount of data. When the number of variables is large and the sample size is insufficient, the problem of over-fitting occurs and leads to weakened model generalizability. To create a balance between model simplicity and model fit, regularization methods such as lasso (least absolute shrinkage and selection operator; Tibshirani, 1996) have been used in many fields and increasingly in social sciences (Lindström & Dahl, 2020). The idea behind lasso is to penalize models that are overly complex (i.e., that contain a large number of parameters) by using a more conservative criterion than traditional methods such as maximum likelihood or least square. To illustrate the principle of lasso, consider the following example for fitting a regression model with J predictors X to outcome data Y by minimizing the estimation function:

$$L^{\text{Lasso}}(\beta) = (|Y - X\beta|)^2 + \lambda \sum_{j=1}^J |\beta_j| \quad (1)$$

where $\beta(J \times 1)$ indicates the vector of regression coefficients for the J predictors, and $L^{\text{Lasso}}(\beta)$ and $(|Y - X\beta|)^2$ represent the loss function of lasso regression and least square difference, respectively. The term $\sum_{j=1}^J |\beta_j|$ is the lasso penalty function that consists of the sum of absolute values of regression coefficients. The tuning parameter, $\lambda(\geq 0)$, indicates the strength of the penalty. The addition of the penalty term to the least square criterion tends to result in simpler models that only include strong predictors.

Compared to models without a penalty function, models with lasso regularization have proved to have higher generalizability (Yarkoni & Westfall, 2017). Compared with the ridge regularization method, which uses a quadratic penalty function $\sqrt{\sum_{j=1}^J \beta_j^2}$ and enforces less shrinkage strength on smaller coefficients, lasso has been found to perform better in variable selection (Hesterberg et al., 2008).

The Bayesian lasso version has been increasingly used in the social sciences (Brandt et al., 2018; Feng et al., 2017; Pan et al., 2017), and this trend is expected to continue (McNeish, 2015) for several reasons. First, lasso can be readily applied in Bayesian analyses by using the double exponential priors (Park & Casella, 2008). Second, with advances in Bayesian computational tools, such as the Gibbs sampler, the Bayesian lasso method can provide estimates of standard errors that are difficult to obtain under a frequentist framework (Kyung et al., 2010). Third, the tuning parameters, which are traditionally chosen based on computationally intensive methods, such as cross-validation method and multiple-model comparison, can be more conveniently estimated with other coefficients simultaneously under the Bayesian lasso paradigm (Hans, 2009; Park & Casella, 2008).

Recently, novel methods that extend the lasso have been developed for structural equation modeling (SEM) and network analysis. Examples include the graphical lasso network models (Costantini et al., 2019), network analysis with an adaptive lasso method (Marcus et al., 2017), the Bayesian lasso confirmatory factor analysis (CFA; Chen et al., 2021; Pan et al., 2017), exploratory mediation analysis (Serang et al., 2017), the Bayesian adaptive lasso for ordinal regression with latent variables (Feng et al.,

2017), and regularized multiple-indicators and multiple-causes (MIMIC) models (Jacobucci et al., 2019). An R-package *regsem* was also developed by Jacobucci et al. (2020) to conduct regularized SEM. Lasso and adaptive lasso methods are available under the frequentist framework in *regsem*.

Although the application of lasso has sharply increased, decision rules for parameter identification (i.e., for determining whether a coefficient is non-zero and included in the model) tend to vary greatly across studies. At least three criteria for parameter identification – the thresholding rule, the p -value rule, and the credible interval rule – have been used in the literature. The lack of a standard for reporting of results for parameter identification creates confusion for result interpretation and problems such as difficulty in comparing results across studies.

In this paper, we compare the three criteria for identifying parameter for inclusion into a model using the Bayesian lasso.

Specifically, we focus on confirmatory factor analysis (CFA) and examine two parameter spaces that are amenable to regularization method: cross-loadings and residual correlations. To illustrate the research question, we use a data set regarding burnout in elementary school male teachers ($N = 372$) (Byrne, 1994, 2012). Participants were asked to respond to the 22-item Maslach Burnout Inventory (MBI; Maslach & Jackson, 1981, 1986). The MBI used a 7-point Likert scale (0 = feeling has never been experienced, 6 = feeling experienced daily) and includes three dimensions of burnout. Figure 1a shows the factor structure specified in original articles. While it is unlikely that each item is only loaded on one factor, as specified a priori, it is also unlikely that all items significantly load on all factors. In other words, we expect sparsity in the cross-loading pattern, which contains a total of 44 possible cross-factor loading.

By applying the Bayesian lasso to the data, we obtained estimates for all cross-loadings. Figure 1b shows the results of

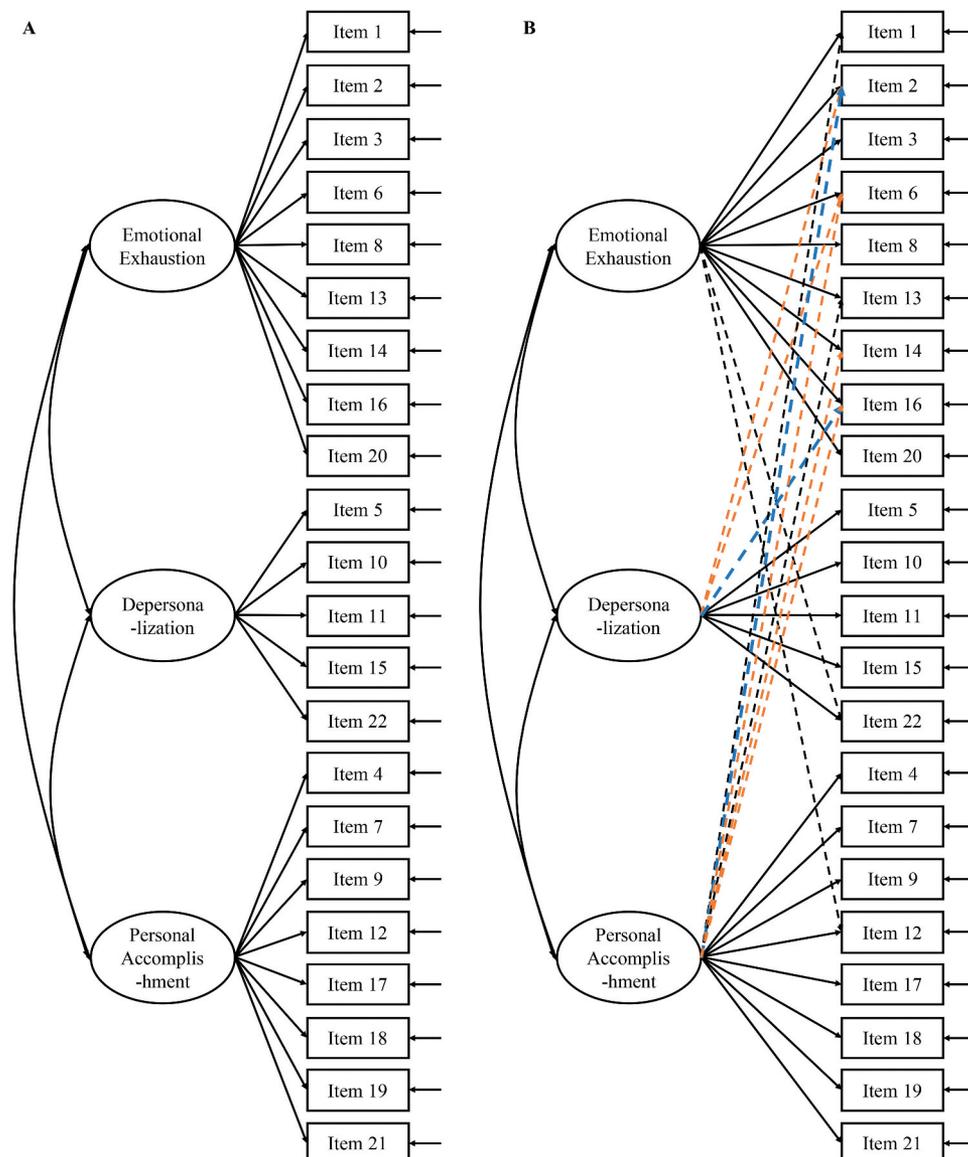


Figure 1. Measurement model of the Maslach Burnout Inventory. Figure 1a: Three dimensions of the Maslach Burnout Inventory. Figure 1b: Results of the Bayesian Lasso CFA. Black solid line: main loadings, black dotted line: cross-loadings identified by all three criteria, blue dotted line: cross-loadings identified by threshold and p -value, orange dotted line: cross-loadings identified only by threshold.

applying the three parameter identification criteria – threshold 0.1, p -value and 95% highest posterior density (HPD) interval – to the estimated values. Although there is some agreement of identified cross-loadings among the criteria, the numbers of identified significant cross-loadings were different – i.e., 11, 6, and 4 cross-loadings were, respectively, identified. Without knowing the true loading values, it is not possible to tell which method works best.

In this paper, our goal is to evaluate, through extensive simulation experiments, the performance of the three criteria under different conditions. We apply the Bayesian lasso to regularize both cross-loadings or residual correlations (both within- and across-factor). Based on the results of the evaluation, we provide further recommendations on the decision rules for using the Bayesian lasso CFA.

The remainder of the paper is organized as follows. First, we provide background for the Bayesian lasso for CFA and the three criteria and decision rules for parameter identification. We next describe the design for the simulation study. Results are reported by two sets of parameters – cross-loading and residual correlation, and three metrics – power, Type I error, and percentage of correct identification. Finally, we provide a discussion of the findings.

Bayesian lasso confirmatory factor analysis

The Bayesian lasso method has been used in CFA to detect possible cross-loadings of indicators (Chen et al., 2021) as well as to identify non-zero residual covariances (Pan et al., 2017). Suppose there are J items and N participants, the CFA is specified as follows:

$$\mathbf{y}_i = \boldsymbol{\mu} + \boldsymbol{\Lambda}\boldsymbol{\omega}_i + \boldsymbol{\varepsilon}_i, i = 1, 2, \dots, N \quad (2)$$

where \mathbf{y}_i ($J \times 1$) represents the observed values of J items for the i -th participant. The vector $\boldsymbol{\mu}$ represents the intercept, and $\boldsymbol{\Lambda}$ represents the loading matrix that encode the relationship between observed variables in \mathbf{y}_i and latent variables in $\boldsymbol{\omega}_i$. The error term $\boldsymbol{\varepsilon}_i$, which follows $N[\mathbf{0}, \boldsymbol{\Psi}_\varepsilon]$, represents residuals of observed variables. In traditional CFA models, the following two assumptions are made: (1) the structure of loadings matrix is specified a priori according to substantive theory such that each item is loaded on one and only one latent factor, and (2) the residual variance-covariance matrix $\boldsymbol{\Psi}_\varepsilon$ is diagonal. In other words, all cross-loadings and residual covariances are assumed to be zero. These two assumptions are rather restrictive and not likely to be entirely satisfied in real-world applications. If non-zero parameters exist and are not identified and properly handled, it could result in an inferior model fit (Muthén & Asparouhov, 2012), and the structural estimate could be biased. For example, Brandt et al. (2020) found that even missing a single cross-loading can cause an unacceptably large bias of interaction effect. Similarly, it has been shown that correctly identified non-zero residual covariances lead to improved precision in the estimates of structural parameters (Pan et al., 2017).

Under the Bayesian framework, the strict assumptions of no cross-loading and zero residual covariance are relaxed through the assignment of priors to the corresponding parameters.

Specifically, the Bayesian lasso assigns double exponential priors to the parameter, thus allowing cross-loadings and residual covariances to fluctuate around zero. The Bayesian lasso method simultaneously identifies non-negligible cross-loadings and/or residual covariances in a joint estimation procedure.

Because regularization methods work by shrinking many small parameters toward zero and only retain significantly large non-zero entries, the technique is especially useful when one expects the parameter space to be sparse (e.g., when only a few non-zero cross-loadings exist). Note that subsequent to using the Bayesian lasso or another regularized method, it is quite common to reanalyze the data with the identified freed (non-zero) parameters without regularization (Muthén & Asparouhov, 2013; Serang & Jacobucci, 2020; Serang et al., 2017).

The identification of “significantly large” non-zero entries require careful operationalization. While the Bayesian lasso can produce probabilistic measures such as a credible interval or p -value as evidence of a non-zero parameter value, a user is still required to determine whether a parameter should be included into the model, a process we call “parameter identification.” The process sets parameters that are deemed small to zero.

Different ways of parameter identification with the lasso method

In general, there are three common ways of parameter identification in lasso: (1) the thresholding rule; (2) the p -value rule; and (3) the credible interval rule.

When standard error is not available or difficult to derive (e.g., in a frequentist estimation), it is common to apply a simple thresholding rule to the (absolute) value of the estimated parameter for parameter identification. Apparently, an intuitive cutoff value is zero or a small value. In a frequentist setting, Serang et al. (2017) compared lasso with the zero-thresholding rule and other p -value based methods (such as multivariate delta method) in an exploratory mediation analysis. Results showed that lasso outperformed p -value-based methods in correctly identifying the mediators. Liang and Jacobucci (2019) also chose zero as the threshold in detecting measurement bias in maximum likelihood estimation. Yuan and Liu (2020) adopted the same threshold to select loadings in CFA and found that lasso did not perform well under this situation. Unlike the other studies, Serang and Jacobucci (2020) chose 0.001 instead of 0 as threshold when selecting mediators to exclude the negligible effect.

Although the zero thresholding rule has been widely used in the frequentist Lasso, it is not feasible in the Bayesian Lasso. In frequentist Lasso, parameters are viewed has a fixed value and can be shrunken exactly to zero. On the other hand, in the Bayesian Lasso, parameters are random variables with distributions and cannot be shrunken to zero. Within the Bayesian estimation, some researchers have opted to use the cutoff of 0.1 ($|\beta| > 0.1$). This cutoff value is especially popular with the Bayesian lasso method (Feng et al., 2017; Guo et al., 2012; Hoti & Sillanpää, 2006; Peterson et al., 2014). The cutoff value can also be justified from a substantive standpoint. Muthén and Asparouhov (2012) suggested that cross-loading of less than

0.1 can be considered to have little practical importance. Cohen (1988) also pointed out that the correlation coefficient as 0.1 is a typical value of low-effect size. Although substantively the cutoff value of 0.1 seems reasonable, the criterion may not have sufficient power to detect truly non-zero parameters. This is due partly to the shrinkage effect of regularization methods – that is, it tends to deflate parameters in achieving a high degree of parsimony.

In frequentist lasso applications when the standard error is

Table 1. Design of simulation study.

| Factors | Conditions |
|---------------------------|---|
| Sample Size | 200, 500, 1000 |
| Model Size | 2 factors and 10 items, 3 factors and 18 items |
| Effect Size of Parameters | 0, 0.1, 0.2, 0.3 for cross-loadings 0, 0.1, 0.3, 0.7 for residual correlations |

interval could maintain low Type I error rates. In this paper, the nominal α level at 0.05 is used throughout.

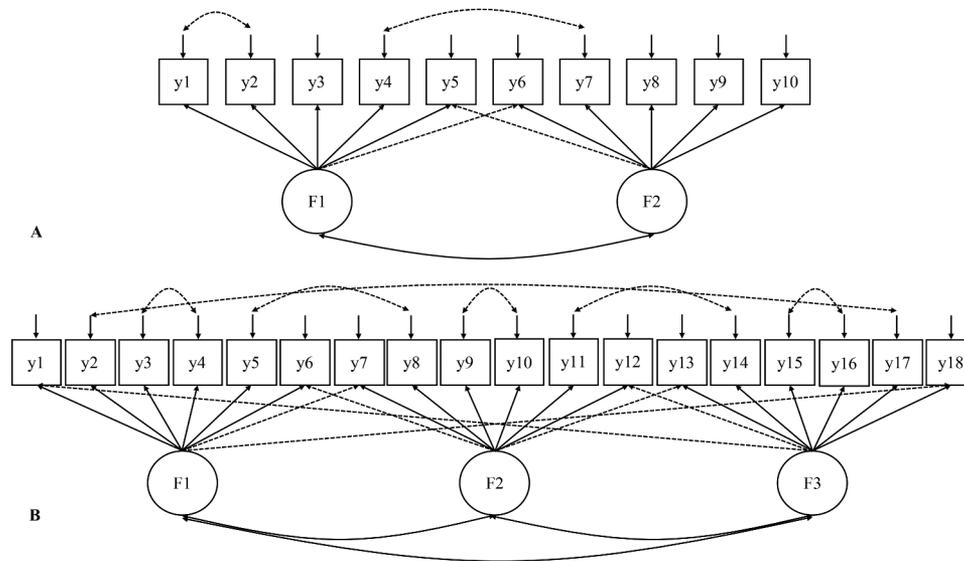


Figure 2. Measurement models in the simulation study. Figure 2a: measurement model with 2 factors and 10 items, Figure 2b: measurement model with 3 factors and 18 items. Dotted line with single arrow: cross-loadings, dotted line with double arrow: residual co-variances, solid line with single arrow: main loadings and residual variances, solid line with double arrow: factor co-variances.

available, the p -value (with cutoff at 0.05) is the conventional standard for parameter identification. For lasso regression, the p -value can be obtained using the R-package *covTest* (Lockhart et al., 2014). For the Bayesian lasso method, it is relatively straightforward to obtain p -values. Additionally, credible intervals such as the HPD interval can also be calculated using Markov Chain Monte Carlo (MCMC; Gilks et al., 1996) algorithms. In network analysis, Epskamp et al. (2018) demonstrated how frequentist-lasso regularization can provide confidence intervals using a bootstrap method. The authors further argued that because the distribution of Lasso regularized parameters is far from normal (Pötscher & Leeb, 2009), thus using the bootstrapped confidence interval is more appropriate than using the p -value.

When applying lasso to SEM, multiple testing of parameter significance is common. There does not appear to be consensus as to whether the α level should be corrected. Epskamp et al. (2018) pointed out the α level would be corrected to 0.000003 even with a small 20-node network if the Bonferroni correction is used. Such stringent criterion is not practical and thus not recommended by the authors. Pan et al. (2017) adopted the nominal α level at 0.05 when calculating the HPD intervals of residual covariances and demonstrated that a 95% HPD

Simulation study

The Monte Carlo simulation study was conducted to investigate the performance of three criteria – thresholding, p -value, and credible interval – as well as their associated decision rules for detecting significant cross-loading/residual correlation under a range of experimental conditions. For each condition, 100 replications were generated and each sample was analyzed by Bayesian lasso CFA.

Method

Data generation

The following factors were selected as conditions for the Monte Carlo study for cross-loading: sample size (3 levels), model size (2 levels), and the magnitude of cross-loading (4 levels). Similar conditions were used for residual correlations (Table 1). As such, we had a total of $3 \times 2 \times 4 = 24$ conditions for each set of parameters or a total of 48 conditions for cross-loading and residual correlation. For model size, two measurement models of different number of factors were considered: a two-factor model with five items per factor and a three-factor model with

six items per factor (Figure 2). These model sizes were commonly used in previous empirical and simulation studies (Chen et al., 2021; Khong et al., 2013; Muthén & Asparouhov, 2012). The variances of factors and items were set at one and the main loadings were set at 0.7.

Apart from the conditions of zero cross-loading and residual correlation, we considered three levels of non-zero effect size. Denote the factor loading between item a and factor b $\lambda_{a,b}$. Following Lu et al. (2016) and Muthén and Asparouhov (2012), two or 20% cross-loadings ($\lambda_{6,1}, \lambda_{5,2}$) in the two-factor model, and six or 16.67% cross-loadings ($\lambda_{7,1}, \lambda_{18,1}, \lambda_{6,2}, \lambda_{13,2}, \lambda_{1,3}, \lambda_{12,3}$) in the three-factor model (Figure 2) were set at values of 0.1, 0.2 or 0.3. Similarly, 2 or 3.64% out of 55 ($= C(10, 2)$) off-diagonal elements in the residual variance-covariance matrix were non-zero in the two-factor model. Six or 3.51% out of 171 ($= C(18, 2)$) residual covariances were non-zero in the three-factor model. The residual correlations were set at 0.1, 0.3 or 0.7 to represent, low, medium, and high correlation (Muthén & Asparouhov, 2012; Pan et al., 2017).

To avoid the possible confounding effect between non-trivial cross-loadings and residual correlations, the conditions of non-zero cross-loadings and non-zero residual correlations were separately analyzed. In other words, when the purpose was to detect significant cross-loadings, the non-diagonal parameters of residual variance-covariance matrix were set to zero in both data generation and model estimation and vice versa.

Model estimation

The specific models used for the simulation studies were: (1) M1: model with some non-zero cross-loadings and diagonal residual covariance matrix, and (2) M2: model with some non-zero, off-diagonal residual covariance entries but no cross-loading.

Using the MCMC algorithm, the posterior distribution of parameters can be estimated based on the specified priors and the data. Suppose there are J items and K factors for the confirmatory factor analysis, the priors for the Bayesian lasso confirmatory factor analysis are issued as follows:

$$\boldsymbol{\mu} \sim N(\boldsymbol{\mu}_0, \mathbf{H}_{\mu 0}), \boldsymbol{\Phi}^{-1} \sim \text{Wishart}(\mathbf{R}_0, \rho_0) \quad (3)$$

For M1, the loadings for the j -th item $\Lambda_j = \begin{pmatrix} \Lambda_j^m \\ \Lambda_j^c \end{pmatrix}$, where Λ_j^m and Λ_j^c , respectively represent the component of the loading on the designated factor and the component of cross-loading. The hierarchical priors for loadings are specified by:

Table 2. Three sets of values for hyperparameters.

| Set | μ_0 | $H_{\mu 0}$ | R_0 | ρ_0 | Λ_{0j} | H_{0j} | a_{0j} | b_{0j} | a_{lj} | b_{lj} | $a_{\lambda 0}$ | $b_{\lambda 0}$ |
|-----|---------|-------------|---------|----------|----------------|----------|----------|----------|----------|----------|-----------------|-----------------|
| 1 | 0 | 4I | I + 0.1 | K + 2 | 0 | 4I | 1 | 0.01 | 1 | 0.01 | 1 | 0.01 |
| 2 | 0 | 4I | I + 0.5 | K + 8 | 0 | I | 1 | 0.1 | 1 | 0.1 | 1 | 0.1 |
| 3 | 0 | 100I | I | K + 2 | 0 | 100I | 1 | 0.01 | 1 | 0.01 | 1 | 0.01 |

I: identity matrix, I + 0.1: diagonal elements = 1 and non-diagonal elements = 0.1, I + 0.5: diagonal elements = 1 and non-diagonal elements = 0.5, K: number of factors.

$$\begin{aligned} \Lambda_j^m &\sim N(\Lambda_{0j}, H_{0j}), \Lambda_j^c | \psi_{jj} \sim N(0, \psi_{jj} \mathbf{D}_{\tau_j}), \\ \psi_{jj}^{-1} &\sim \text{Gamma}(a_{0j}, b_{0j}), \mathbf{D}_{\tau_j} = \text{diag}(\tau_{j1}^2, \dots, \tau_{jk}^2), \\ \tau_{jk}^2 &\sim \text{Gamma}\left(1, \frac{\delta^2}{2}\right), \delta^2 \sim \text{Gamma}(a_{lj}, b_{lj}) \end{aligned} \quad (4)$$

For M2, the prior for the loadings is $\Lambda_j \sim N(\Lambda_{0j}, H_{0j})$. Following Khondker et al. (2013) and Wang (2012), graphical lasso priors are specified for the inverse of the residual variance-covariance matrix $\boldsymbol{\Psi}_\epsilon$:

$$\boldsymbol{\Psi}_\epsilon^{-1} = \boldsymbol{\Sigma} = (\sigma_{jj'})_{J \times J} \quad (5)$$

where independent exponential priors ($\frac{\lambda}{2} \exp(-\frac{\lambda}{2} \sigma_{jj'})$) and the double exponential priors ($\frac{\lambda}{2} \exp(-\lambda |\sigma_{jj'}|)$, $j < j'$) are assigned for the diagonal and the off-diagonal elements of $\boldsymbol{\Psi}_\epsilon^{-1}$, respectively. Moreover, $\lambda \sim \text{Gamma}(a_{\lambda 0}, b_{\lambda 0})$ where $a_{\lambda 0} = 1$ and $b_{\lambda 0}$ set at a small value.

For sensitivity analysis, we adopted three sets of values for hyperparameters $\boldsymbol{\mu}_0, \mathbf{H}_{\mu 0}, \mathbf{R}_0, \rho_0, \Lambda_{0j}, H_{0j}, a_{0j}, b_{0j}, a_{lj}, b_{lj}, a_{\lambda 0}, b_{\lambda 0}$ (Table 2).

The factor variances were set at 1.0 for model identification, and the MCMC algorithms proposed by Liu (2008) and Liu and Daniels (2006) was used to estimate the factor correlation matrix. Other parameters were estimated using the Gibbs sampler (Geman & Geman, 1984). We conducted several test runs to select the number of MCMC iterations and found that the algorithm converged in less than 8,000 iterations as judged by the value of the estimated potential-scale reduction (EPSR) values (< 1.2 ; Gelman, 1996). Based on the result from the test runs, all final estimates were obtained using 10,000 draws from the posterior distributions after a burn-in phase of 10,000 iterations. Data generation and estimation used R 3.5.3 (R Core Team, 2020).

Parameter identification

The following criteria for parameter identification were used in the current study:

- (1) Thresholds of magnitude 0, 0.05, 0.1, and 0.15 with the decision rule to include if the absolute value of the parameter estimate is larger than the cutoff. The thresholding rule applies to the estimates of residual correlations and standardized cross-loadings.
- (2) A p -value with $\alpha = 0.05$, with the decision rule to include if $p < .05$. The p -value can be different from the frequentist p -value, it is one-tailed and is based on MCMC samples rather than the z -test.¹ It has been used in the Bayesian analysis with *Mplus* software. As defined by Muthén (2010), the p -value is the proportion of the negative/positive posterior samples for or a positive/negative estimate, respectively.
- (3) A 95% HPD interval, with the decision rule to include if the point 0.0 is outside the 95% HPD interval.

¹We also checked the performance of the z -test, and results were similar to but slightly worse than the 95% HPD interval. Results based on the z -test are not reported in the current paper. Note that the worse performance of z -test may be due to the non-normal property of Lasso regularized parameters (Pötscher & Leeb, 2009).

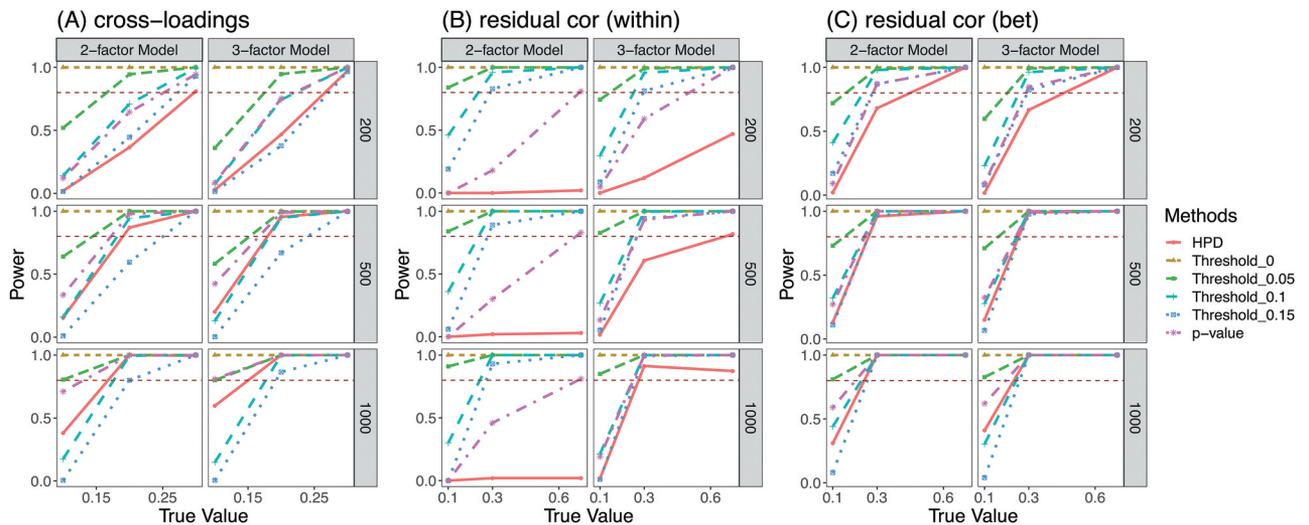


Figure 3. Power. Note: Sample size: 200, 500, and 1000. True value: 0.1, 0.2 and 0.3 for cross-loadings and 0.1, 0.3, 0.7 for residual correlations. Methods: methods of variable selection, thresholds include 0, 0.05, 0.1, and 0.15, HPD: 95% highest posterior density interval. Residual cor (with): within-factor residual correlations. residual cor (bet): between-factor residual correlations. The horizontal red dotted line indicates the threshold of acceptable power (0.8).

Note that for the thresholding criterion, the values of 0.1 were commonly used in previous Bayesian studies. We include the zero thresholding rule, which is commonly used in frequentist Lasso in this work and show that it is equivalent to no thresholding under the Bayesian framework. Here we also included two additional cutoff values of 0.05 and 0.15 to fully investigate the effect of different thresholds.

Evaluation criteria

Three metrics were considered for a performance comparison of the decision rules:

- (1) Power: the probability of correctly identifying the cross-loadings/residual correlations when the parameters are non-zero (Muthén & Asparouhov, 2012). Acceptable power is above 80% (L. K. Muthén & Muthén, 2002).
- (2) Type-I error rates: the probability of erroneously identifying the cross-loadings/residual correlations when the parameters are zero. The acceptable Type I error rate is within the 95% confidence interval of a binomial distribution [0.007, 0.093]

$$[0.05 \pm 1.96 \times \sqrt{0.05 \times (1 - 0.05) / n_{\text{rep}}}; \text{Cham et al., 2012}], \text{ where } n_{\text{rep}} \text{ is the number of replications.}$$

- (3) The ratio of correct identification to the total number of identified parameters (Yuan & Liu, 2020): There is no established criterion for this metric.

We focus on these three metrics in the current paper, the estimation results of Bayesian Lasso CFA were reported in the supplementary materials. The results of these metrics were similar in three sets of hyperparameters, the results of the first set of hyperparameters were reported in the paper and the others were also shown in the supplementary materials.

Results

Power

Figure 3 shows the performance of the three criteria for decision rules on power. Compared to the other two rules, the HPD interval rule performed worst in detecting non-zero cross-loadings when the sample size was small or when the target was within-factor residual correlations. For example, even when the sample size was 1,000 and the residual correlations were 0.7, the power of the 95% HPD interval rule in detecting within-factor residual correlations in the two-factor model was close to zero. The performance of the p -value decision rule was better than HPD interval rule in detecting cross-loadings and between-factor residual correlations, but was worse than the thresholding rule with threshold values of 0.05 and 0.1. However, the p -value rule had similar power problems as the HPD interval rule. For example, in the two-factor model, the p -value rule did not have sufficient power even when the sample size was 1,000 and the residual correlation was 0.3.

As expected, the zero-threshold produced 100% power in all the conditions. Under most conditions, power was maintained at approximately 80% by increasing the threshold to 0.05. The cutoff value of 0.1 only produced acceptable power when the estimated value was larger than 0.1. We also found that the power was higher for residual correlations than for cross-loadings with the same threshold and the same level of actual parameter values (0.1, 0.3). The threshold of 0.15 was acceptable in power for residual correlations but not for cross-loadings.

Additionally, we found an increasing sample size significantly improved the performance of the HPD interval and the p -value decision rules. However, the increasing sample size only led to a small improvement in the performance of the thresholding rules in general. Moreover, except for sample size, the performance of the HPD interval was sensitive to model size, especially when detecting within-factor residual correlations.

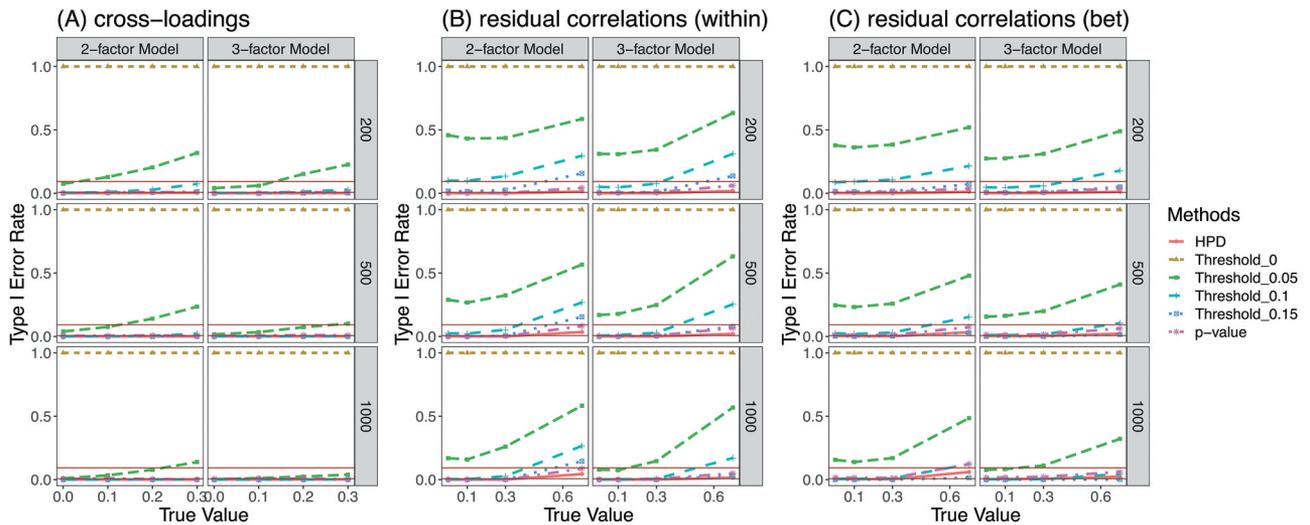


Figure 4. Type I error rates. *Note:* Type I error rates for zero cross-loadings and residual correlations under different modeling conditions: (1) Sample size: 200, 500, and 1000; (2) True value: 0 for all the cross-loadings and residual correlations, 0.1, 0.2, 0.3 for the pre-assigned non-zero cross-loadings, and 0.1, 0.3, 0.7 for non-zero residual correlations; (3) Methods: methods of variable selection, thresholds include 0, 0.05, 0.1, and 0.15, HPD: 95% highest posterior density interval. The two horizontal red lines indicate the bounds of acceptable Type I error rate (0.007–0.093).

Type I error rate

Figure 4 displays the performance of the three criteria for decision rules on Type I error rate. The general pattern was almost a mirror image of that of power. For example, the HPD interval rule produced the lowest Type I error rate and outperformed the thresholding rule in most conditions. The p -value rule also maintained the Type I error rate well in all the conditions. For the thresholding rule, Type I error rates were higher for residual correlations than for cross-loadings. As expected, the Type I error rate decreased with increasing cutoff value. The 0.05 cutoff produced Type I error rates higher than 0.093 for cross-loadings in some conditions and for residual correlations in almost all conditions. The cutoff of 0.1 maintained Type I error rates well at the nominal level for cross-loadings. However, for residual correlations conditions, the 0.1 cutoff

produced high Type I error rates (>0.3 in some conditions). Using the 0.15 cutoff value instead significantly improved performance.

Although the abovementioned results suggest using the thresholding rule at the 0.1 and 0.15 cutoffs, respectively, for cross-loading and residual correlation for maintaining the Type I error, using a lower cutoff may be appropriate for other purposes, such as exploring factor and residual covariance structures (Jacobucci et al., 2019; Serang et al., 2017). In the case in which power is emphasized, the two respective thresholds (0.05 for cross-loading and 0.1 for residual correlation) are favored. Results showed that for parameters with a true value of 0.1 or higher, these thresholds provided high power for detection than the 0.1-threshold for loadings and the 0.15-threshold for residual correlations (Figure 3).

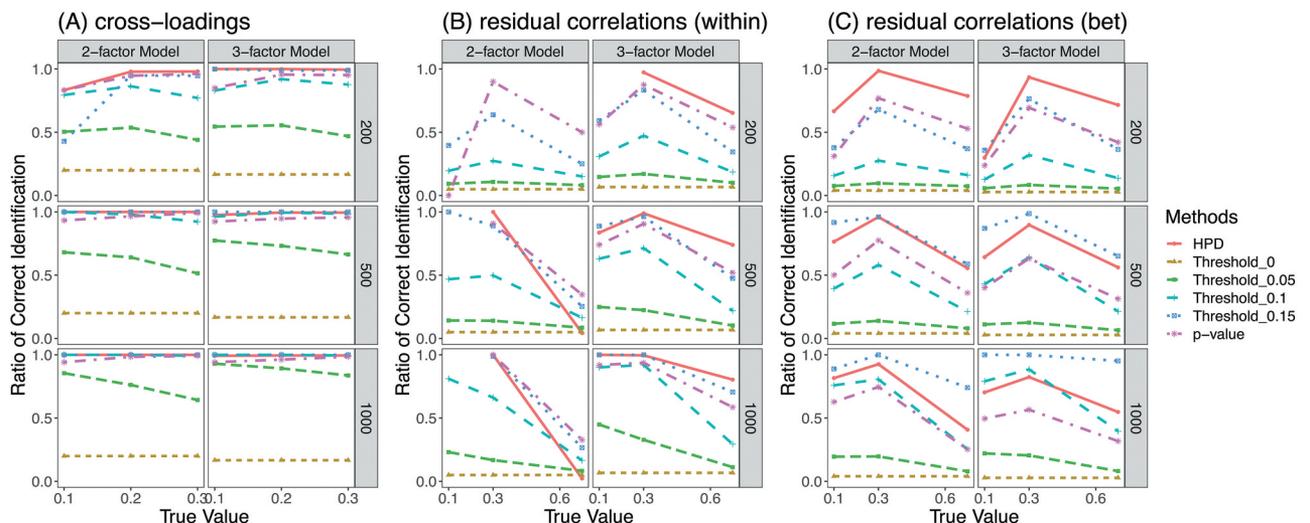


Figure 5. Ratio of correct identification. *Note:* Sample size: 200, 500, and 1000. True value: 0.1, 0.2 and 0.3 for cross-loadings and 0.1, 0.3, 0.7 for residual correlations. Methods: methods of variable selection, thresholds include 0, 0.05, 0.1, and 0.15, HPD: 95% highest posterior density interval. If a line is missing, it indicates that no cross-loading/residual correlation was identified by the method in that condition.

Table 3. Summary of recommendations.

| Sample Size | Parameters | Threshold 0.1 | | Threshold 0.15 | | HPD Interval | | <i>p</i> -value | |
|-------------|-----------------------------|---------------|--------|----------------|--------|--------------|--------|-----------------|--------|
| | | Power | Type I | Power | Type I | Power | Type I | Power | Type I |
| 200 | Cross-loadings ^a | × | √ | × | √ | × | √ | × | √ |
| | Residual correlations | √ | × | √ | √ | × | √ | × | √ |
| 500 | Cross-loadings | √ | √ | × | √ | √ | √ | √ | √ |
| | Residual correlations | √ | √ | √ | √ | × | √ | × | √ |
| 1000 | Cross-loadings | √ | √ | √ | √ | √ | √ | √ | √ |
| | Residual correlations | √ | √ | √ | √ | × | √ | × | √ |

Type I: Type I Error Rate; √: acceptable in most conditions, ×: unacceptable in many conditions, shaded: the best criterion in the corresponding condition.

^aNone of the criteria can provide sufficient power under this condition.

The ratio of correct identification

Figure 5 summarizes the results for the metric in the correct identification ratio. The ratio increased with an increasing sample and decreasing Type I error rate (Figure 5). The metric was higher for cross-loadings compared to residual correlations, which was partly caused by the relatively large proportion of zero residual correlations in the generative models.

Results showed the more conservative the rule, the higher the ratio of correct identification. The most conservative rule – a 95% HPD interval – performed best in terms of this metric in most conditions. The *p*-value rule also performed better than the thresholding rule for detecting within-factor residual correlations. Cutoff at zero performed worst, followed by the 0.05 cutoff. For correctly identifying cross-loading, cutoff values of 0.1 and 0.15 performed the best, and their performances were similar. The 0.15 cutoff outperformed the 0.1 cutoff in detecting residual correlations.

Discussion

The current study offers a comprehensive comparison between the thresholding rule, the *p*-value rule, and the HPD interval rule in parameter identification within the context of the Bayesian lasso CFA. In general, we found no universally superior criterion or decision rule. For example, criteria based on thresholding rules with cutoff value >0 showed a more robust performance than the HPD interval in balancing power and Type I error rate. However, the thresholding rules did not perform well in terms of percentage of correct identification. As different criteria suit different purposes (e.g., for exploratory purposes power is more important, whereas for formal hypothesis testing, Type I error rate is more important), we cannot recommend universal simple rules-of-thumb for parameter identification. Instead, we summarize our recommendations for which criterion to consider in a table that separates out major circumstances – metric to consider, sample size, and model parameter type (Table 3).

In general, power and Type I error rates for rules using thresholding criteria were higher for residual correlations than for cross-loadings. For cross-loading, the threshold of 0.1, which has been commonly used by previous studies (Feng et al., 2017; Guo et al., 2012; Hoti & Sillanpää, 2006; Peterson et al., 2014) provides a good balance – i.e., sufficient power and acceptable Type I error rates. The cutoff of 0.1, however, produced unacceptably high Type I error rates for residual correlations; as such, the cutoff value of 0.15 is instead recommended.

The current study identified cross-loadings and residual correlations separately to avoid possible-confounding effects. However, the

co-existence of cross-loadings and residual correlations may be expected in practice. We further conducted a small simulation study to preliminarily explore the performance of different criteria in detecting these two kinds of parameters simultaneously. Results were rather similar to the current study. We highlight key findings and include details in Supplementary Materials. For detecting the two kinds of parameters simultaneously, the phenomenon of low power of *p*-value and HPD interval methods was still present for within-factor residual correlations. Compared to performance when only one kind of parameter was present, we found the HPD interval and *p*-value methods were more sensitive to model size and provided lower power. Thus, we do not recommend to use *p*-value for detecting cross-loadings with small and medium sample size ($N < 500$). With small sample size, we recommend the use of the thresholding rule with different threshold values for different parameters.

Usage of the thresholding rule

The findings in this paper are based on the context of relaxed CFA assumptions and the Bayesian lasso tool for estimating the relaxed model. In the broader context of SEM, relaxing similar constraints using regularized methods has become an emerging and important trend (Jacobucci & Grimm, 2018; Lu et al., 2016; Muthén & Asparouhov, 2012). The 0.1 cutoff value, which we recommend for detecting cross-loading, can be a candidate for application to parameter identification for path coefficients in SEM, given the similar roles of path coefficient and factor loading as regression parameters. The simple 0.1 cutoff rule, pending further investigation, could be especially useful in helping researchers to specify complex structural models in which causal relationships are difficult to fully delineate from theoretical consideration. On the other hand, some SEM applications require conditional correlations to be properly modeled for improving model fit. For example, in a SEM with two time points, the same indicator variables of the same latent factor across the time points may be modeled as correlated. From the results in this paper, we conjecture that using a cutoff of 0.15 for such correlation would be appropriate if the Bayesian lasso model is adopted. As a caveat, the current study did find that the 0.1 and 0.15 cutoffs (respectively for cross-loading and residual correlation) were unable to provide sufficient power in some cases when the true value of the parameter is 0.1, although power is acceptable when the true value of the parameter is >0.1. In situations in which path coefficients (standardized) and correlations less than 0.1 are not deemed meaningful, the respective cutoffs at 0.1 and 0.15 represent reasonable choices. Our recommendation can serve as a preliminary guideline for future

investigation into the use of simple thresholding rules for broader use in SEM.

For exploratory analysis in SEM in which the purpose is to extract as many potentially important relationships as possible, the cutoff value can be lowered; we suggest using 0.05 for path coefficients and 0.1 for correlations. Further study is needed to fully assess the performance of these rules for exploratory SEM.

Limitations and future research

The current study has limitations. The number of factors investigated in the simulation study is relatively small, and the method is limited to the Bayesian lasso. For example, our simulation results show that the power of the HPD interval method in detecting within-factor residual correlations increases with the number of factors. However, the model sizes considered were limited. More work is needed to evaluate the performance of the HPD interval under various model sizes. Future studies should also investigate the criteria performance and decision rules in a broader setting, including CFA with a larger number of factors, as well as other extensions of the lasso method such as the adaptive lasso (Zou, 2006).

As pointed out by one of the reviewers, variable selection uncertainties are not considered in the commonly used two-step approach in Bayesian Lasso CFA. Lu et al. (2016) proposed a solution using a spike-and-slab prior (SSP). More work is needed to evaluate the performance of SSP in parameter identification especially with the three methods analyzed in the current paper.

Concluding remarks

In summary, the current paper provides guidelines on the use of different decision rules for parameter identification in a CFA when the Bayesian lasso method is applied. The guidelines are based on a comparison of several commonly used criteria. As pointed out by Hindman (2015), regularized approaches have proven to outperform traditional methods, and methods such as penalized regression (e.g., the Bayesian lasso) should be more widely applied for theory building and hypothesis testing in social sciences. The current study represents an important contribution that facilitates such application.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was funded by the National Natural Science Foundation of China under Grant [31871128]; the MOE (Ministry of Education) Project of Humanities and Social Science of China under Grant [18YJA190013], the NSF Grant (SES-1424875) and an NIH Grant (1UL1TR001420-01) from the United States.

ORCID

Lijin Zhang  <http://orcid.org/0000-0002-4222-8850>
Junhao Pan  <http://orcid.org/0000-0001-7156-6077>
Edward Haksing Ip  <http://orcid.org/0000-0002-4811-4205>

References

- Brandt, H., Cambria, J., & Kelava, A. (2018). An adaptive Bayesian lasso approach with spike- and-slab priors to identify multiple linear and nonlinear effects in structural equation models. *Structural Equation Modeling*, 25, 946–960. <https://doi.org/10.1080/10705511.2018.1474114>
- Brandt, H., Umbach, N., Kelava, A., & Bollen, K. A. (2020). Comparing estimators for latent interaction models under structural and distributional misspecifications. *Psychological Methods*, 25, 321–345. <https://doi.org/10.1037/met0000231>
- Byrne, B. M. (1994). Testing for the factorial validity, replication, and invariance of a measuring instrument: A paradigmatic application based on the Maslach Burnout Inventory. *Multivariate Behavioral Research*, 29, 289–311. https://doi.org/10.1207/s15327906mbr2903_5
- Byrne, B. M. (2012). *Structural equation modeling with Mplus: Basic concepts, applications, and programming*. Routledge.
- Cham, H., West, S. G., Ma, Y., & Aiken, L. S. (2012). Estimating latent variable interactions with nonnormal observed data: A comparison of four approaches. *Multivariate Behavioral Research*, 47, 840–876. <https://doi.org/10.1080/00273171.2012.732901>
- Chen, J., Guo, Z., Zhang, L., & Pan, J. (2021). A partially confirmatory approach to scale development with the Bayesian lasso. *Psychological Methods*, 26, 210–235. <https://doi.org/10.1037/met0000293>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Hillsdale.
- Costantini, G., Richetin, J., Preti, E., Casini, E., Epskamp, S., & Perugini, M. (2019). Stability and variability of personality networks. A tutorial on recent developments in network psychometrics. *Personality and Individual Differences*, 136, 68–78. <https://doi.org/10.1016/j.paid.2017.06.011>
- Epskamp, S., Borsboom, D., & Fried, E. I. (2018). Estimating psychological networks and their accuracy: A tutorial paper. *Behavior Research Methods*, 50, 195–212. <https://doi.org/10.3758/s13428-017-0862-1>
- Feng, X., Wu, H., & Song, X. (2017). Bayesian adaptive lasso for ordinal regression with latent variables. *Sociological Methods & Research*, 46, 926–953. <https://doi.org/10.1177/0049124115610349>
- Gelman, A. (1996). Inference and monitoring convergence. In W. R. Gilks, S. Richardson, & D. J. Spiegelhalter (Eds.), *Markov chain Monte Carlo in practice* (pp. 131–144). Chapman & Hall.
- Geman, S., & Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6, 721–741. <https://doi.org/10.1109/TPAMI.1984.4767596>
- Gilks, W. R., Richardson, S., & Spiegelhalter, D. J. (Eds.). (1996). *Markov chain Monte Carlo in practice*. Chapman & Hall.
- Guo, R., Zhu, H., Chow, S. M., & Ibrahim, J. G. (2012). Bayesian lasso for semi-parametric structural equation models. *Biometrics* 68, 567–577. <https://doi.org/10.1111/j.1541-0420.2012.01751.x>
- Hans, C. (2009). Bayesian Lasso regression. *Biometrika*, 96, 835–845. <https://doi.org/10.1093/biomet/asp047>
- Hesterberg, T., Choi, N. H., Meier, L., & Fraley, C. (2008). Least angle and l1 penalized regression: A review. *Statistics Surveys*, 2, 61–93. <https://doi.org/10.1214/08-SS035>
- Hindman, M. (2015). Building better models: Prediction, replication, and machine learning in the social sciences. *The Annals of the American Academy of Political and Social Science*, 659, 48–62. <https://doi.org/10.1177/0002716215570279>
- Hoti, F., & Sillanpää, M. J. (2006). Bayesian mapping of genotype expression interactions in quantitative and qualitative traits. *Heredity* 97, 4–18. <https://doi.org/10.1038/sj.hdy.6800817>
- Jacobucci, R., Brandmaier, A., & Kievit, R. (2019). A practical guide to variable selection in structural equation models with regularized MIMIC models. *Advances in Methods and Practices in Psychological Science*, 2, 55–76. <https://doi.org/10.1177/2515245919826527>
- Jacobucci, R., Grimm, K., Brandmaier, A., Serang, S., Kievit, R., Scharf, F., & Li, X. (2020). *regsem: Regularized structural equation models*. R package version 1.5.2. Retrieved from <https://cran.r-project.org/web/packages/regsem/index.html>

- Jacobucci, R., & Grimm, K. J. (2018). Comparison of frequentist and Bayesian regularization in structural equation modeling. *Structural Equation Modeling*, 25, 639–649. <https://doi.org/10.1080/10705511.2017.1410822>
- Khondker, Z. S., Zhu, H., Chu, H., Lin, W., & Ibrahim, J. G. (2013). The Bayesian covariance lasso. *Statistics and Its Interface*, 6, 243–259. <https://doi.org/10.4310/SII.2013.v6.n2.a8>
- Khong, K. W., Onyemeh, N. C., & Chong, Y. L. (2013). Bsem estimation of network effect and customer orientation empowerment on trust in social media and network environment. *Expert Systems with Applications*, 40, 4858–4870. <https://doi.org/10.1016/j.eswa.2013.02.020>
- Kyung, M., Gill, J., Ghosh, M., & Casella, G. (2010). Penalized regression, standard errors, and Bayesian lassos. *Bayesian Analysis*, 5, 369–411. <https://doi.org/10.1214/10-BA607>
- Liang, X., & Jacobucci, R. (2019). Regularized structural equation modeling to detect measurement bias: Evaluation of lasso, adaptive lasso, and elastic net. *Structural Equation Modeling*, 27, 722–734. <https://doi.org/10.1080/10705511.2019.1693273>
- Lindström, J. C., & Dahl, F. A. (2020). Model selection with lasso in multi-group structural equation models. *Structural Equation Modeling*, 27, 33–42. <https://doi.org/10.1080/10705511.2019.1638262>
- Liu, X. (2008). Parameter expansion for sampling a correlation matrix: An efficient GPX-RPMH algorithm. *Journal of Statistical Computation and Simulation*, 78, 1065–1076. <https://doi.org/10.1080/00949650701519635>
- Liu, X., & Daniels, M. J. (2006). A new efficient algorithm for sampling a correlation matrix based on parameter expansion and re-parameterization. *Journal of Computational and Graphical Statistics*, 15, 897–914. <https://doi.org/10.1198/106186006X160681>
- Lockhart, R., Taylor, J., Tibshirani, R. J., & Tibshirani, R. (2014). A significance test for the lasso. *The Annals of Statistics*, 42, 413–468. <https://doi.org/10.1214/13-AOS1175>
- Lu, Z. H., Chow, S. M., & Loken, E. (2016). Bayesian factor analysis as a variable-selection problem: Alternative priors and consequences. *Multivariate Behavioral Research*, 51, 519–539. <https://doi.org/10.1080/00273171.2016.1168279>
- Marcus, D. K., Preszler, J., & Zeigler-Hill, V. (2017). A network of dark personality traits: What lies at the heart of darkness? *Journal of Research in Personality*, 73, 56–62. <https://doi.org/10.1016/j.jrp.2017.11.003>
- Maslach, C., & Jackson, S. E. (1981). *Maslach Burnout Inventory manual*. Consulting Psychologists Press.
- Maslach, C., & Jackson, S. E. (1986). *Maslach Burnout Inventory manual* (2nd ed. ed.). Consulting Psychologists Press.
- McNeish, D. M. (2015). Using Lasso for predictor selection and to assuage overfitting: A method long overlooked in behavioral sciences. *Multivariate Behavioral Research*, 50, 471–484. <https://doi.org/10.1080/00273171.2015.1036965>
- Muthén, B., & Asparouhov, T. (2013). BSEM measurement invariance analysis. *Mplus Web Notes: No.17*. Mplus. Retrieved from www.statmodel.com
- Muthén, B. O., & Asparouhov, T. (2012). Bayesian structural equation modeling: A more flexible representation of substantive theory. *Psychological Methods*, 17, 313–335. <https://doi.org/10.1037/a0026802>
- Muthén, B. O. (2010). *Bayesian analysis in Mplus: A brief introduction*. <https://www.statmodel.com/download/IntroBayesVersion%203.pdf>
- Muthén, L. K., & Muthén, B. O. (2002). How to use a Monte Carlo study to decide on sample size and determine power. *Structural Equation Modeling*, 9, 599–620. https://doi.org/10.1207/S15328007SEM0904_8
- Pan, J. H., Ip, E. H., & Dubé, L. (2017). An alternative to post hoc model modification in confirmatory factor analysis: The Bayesian lasso. *Psychological Methods*, 22, 687–704. <https://doi.org/10.1037/met0000112>
- Park, T., & Casella, G. (2008). The Bayesian Lasso. *Journal of the American Statistical Association*, 103, 681–686. <https://doi.org/10.1198/016214508000000337>
- Peterson, C., Vannucci, M., Karakas, C., Choi, W., Ma, L., & Maletic-savatic, M. (2014). Inferring metabolic networks using the Bayesian adaptive graphical lasso with informative priors. *Statistical and Its Interface*, 6, 547–558. <https://doi.org/10.4310/SII.2013.v6.n4.a12>
- Pötscher, B. M., & Leeb, H. (2009). On the distribution of penalized maximum likelihood estimators: The lasso, scad, and thresholding. *Journal of Multivariate Analysis*, 100, 2065–2082. <https://doi.org/10.1016/j.jmva.2009.06.010>
- R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Serang, S., & Jacobucci, R. (2020). Exploratory mediation analysis of dichotomous outcomes via regularization. *Multivariate Behavioral Research*, 55, 69–86. <https://doi.org/10.1080/00273171.2019.1608145>
- Serang, S., Jacobucci, R., Brimhall, K. C., & Grimm, K. J. (2017). Exploratory mediation analysis via regularization. *Structural Equation Modeling*, 24, 733–744. <https://doi.org/10.1080/10705511.2017.1311775>
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58, 267–288. <https://doi.org/10.1111/j.1467-9868.2011.00771.x>
- Wang, H. (2012). Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis*, 7, 867–886. <https://doi.org/10.1214/12-BA729>
- Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 12, 1100–1122. <https://doi.org/10.1177/1745691617693393>
- Yuan, K.-H., & Liu, F. (2020). Which method is more reliable in performing model modification: Lasso regularization or lagrange multiplier test? *Structural Equation Modeling*, 28, 69–81. <https://doi.org/10.1080/10705511.2020.1768858>
- Zhang, L., Pan, J., Dube, L., & Ip, E. H. (2021). blcfa: An R package for Bayesian model modification in confirmatory factor analysis. *Structural Equation Modeling*, 28, 649–658. <https://doi.org/10.1080/10705511.2020.1867862>
- Zou, H. (2006). The adaptive Lasso and its oracle properties. *Journal of the American Statistical Association*, 101, 1418–1429. <https://doi.org/10.1198/016214506000000735>