A Comparison of Penalized Maximum Likelihood Estimation and Markov Chain Monte Carlo
Techniques for Estimating Confirmatory Factor Analysis Models with Small Sample Sizes

Oliver Lüdtke^{1,2}, Esther Ulitzsch¹, and Alexander Robitzsch^{1,2}

¹ Leibniz Institute for Science and Mathematics Education, Kiel, Germany

² Centre for International Student Assessment, Germany

Author Note

Correspondence concerning this article should be sent to Oliver Lüdtke, Leibniz
Institute for Science and Mathematics Education (IPN), Olshausenstr. 62, 24118 Kiel,
Germany. Email: oluedtke@ipn.uni-kiel.de

Abstract

With small to modest sample sizes and complex models, maximum likelihood (ML) estimation of confirmatory factor analysis (CFA) models can show serious estimation problems such as nonconvergence or parameter estimates that are outside the admissible parameter space. In the present article, we discuss two Bayesian estimation methods for stabilizing parameter estimates of a CFA: Penalized maximum likelihood (PML) estimation and Markov Chain Monte Carlo (MCMC) methods. We clarify that these use different Bayesian point estimates from the joint posterior distribution—the mode (PML) of the joint posterior distribution, and the mean (EAP) or mode (MAP) of the marginal posterior distribution—and discuss under which conditions the two methods produce different results. In a simulation study, we show that the MCMC method clearly outperforms PML and that these performance gains can be explained by the fact that MCMC uses the EAP as a point estimate. We also argue that it is often advantageous to choose a parameterization in which the main parameters of interest are bounded and suggest the four-parameter beta distribution as a prior distribution for loadings and correlations. Using simulated data, we show that selecting weakly informative four-parameter beta priors can further stabilize parameter estimates, even in cases when the priors were mildly misspecified. Finally, we derive recommendations and propose directions for further research.

Keywords: measurement error; latent variable models; Bayesian methods; prior distribution; Markov Chain Monte Carlo; penalized maximum likelihood estimation; constrained maximum likelihood estimation; confirmatory factor analysis

1. Introduction

In the social and behavioral sciences, constructs (e.g., intelligence, extraversion) are often conceptualized as latent variables that are measured by error-prone observed indicators (e.g., items). Structural equation modeling (SEM) is a very prominent approach to correct for measurement error when assessing multivariate relationships among latent constructs (Bollen, 1989; Hoyle, 2012). In the SEM approach, a measurement part is distinguished from a structural part. In the measurement part, measurement models are specified to allow for an error-free estimation of the relations in the structural model. In research practice, maximum likelihood (ML) estimation is routinely used to obtain parameter estimates for structural equation models. However, one major limitation of ML estimation is that it needs large sample sizes to reveal its optimal properties (e.g., unbiasedness, efficiency). With small to modest sample sizes and complex models, ML estimation can show serious estimation problems such as nonconvergence or parameter estimates that are outside the admissible parameter space (e.g., negative variances; see Anderson & Gerbing, 1984; Boomsma, 1985; Chen et al., 2001; Gagné & Hancock, 2006; Hoogland & Boomsma, 1998; Smid & Rosseel, 2020; Wolf et al., 2013).

In the recent decade, several researchers have shown that the Bayesian approach has the potential to solve some of the estimation problems that occur in small sample applications of SEM (e.g., Kaplan & Depaoli, 2012; Lee, 2007; Muthén & Asparouhov, 2012; Song & Lee, 2009). First, if appropriate prior distributions are used, the Bayesian approach guarantees that parameter estimates will be within the admissible range, and estimation problems can usually be avoided. Second, Bayesian methods allow for the stabilization of parameter estimates by specifying weakly informative prior distributions for the SEM parameters (Can et al., 2015; Chen et al., 2014; Depaoli & Clifton, 2015; Lee & Song, 2004; Lüdtke et al., 2018; McNeish, 2016; Miocevic et al., 2020; van Erp et al., 2018). The basic idea is that incorporating even a small amount of information in the prior distribution of the SEM

parameters provides some direction for their estimation. However, at the same time, inferences can still be driven by the data (Gelman et al., 2014).

In the present article, we focus on the estimation of confirmatory factor analysis (CFA) models in which a number of latent factors are measured by a set of observed variables (Bollen, 1989). We investigate two critical issues in the Bayesian estimation of CFA models with small sample sizes. First, we discuss two methods for obtaining Bayesian parameter estimates: penalized maximum likelihood (PML) estimation and Markov Chain Monte Carlo (MCMC) methods. We clarify under which conditions the two methods produce different results and compare their performance to the traditional ML estimation of CFA models. Second, we discuss the specification of prior distributions in the Bayesian approach and argue that it can be advantageous to choose a parameterization in which the model parameters (i.e., standardized loadings, latent correlations) are bounded. More specifically, we suggest the four-parameter beta distribution as a prior distribution for bounded parameters and investigate in a simulation study how the specification of thoughtful prior distributions can help to stabilize parameter estimates in small sample size conditions (Smid et al., 2020; Zitzmann et al., 2020).

The article is organized as follows. We start by describing how the basic CFA is estimated with traditional ML estimation. We then discuss the specification of CFA models in the Bayesian approach and describe two different estimation methods (PML and MCMC) for obtaining Bayesian estimates. In the context of a CFA with two latent factors, we discuss issues of parameterization, the specification of prior distributions, and illustrate conditions under which PML and MCMC produce different estimates. We then present the results of a simulation study in which we compare traditional ML estimation and the two Bayesian estimation approaches. In this simulation, we evaluate the influence of correctly and misspecified prior distributions on the quality of parameter estimates in small sample size conditions. Finally, we derive recommendations and propose directions for further research.

2. Confirmatory Factor Analysis

Let $\mathbf{x}$ denote a vector of p observed variables. Then a CFA model with m latent factors is represented as follows

$$\mathbf{x} = \mathbf{v} + \mathbf{\Lambda}\boldsymbol{\eta} + \boldsymbol{\varepsilon} \quad (1)$$

where $\mathbf{v}$ is $p \times 1$ vector containing intercepts, $\mathbf{\Lambda}$ is a $p \times m$ matrix of factor loadings, $\boldsymbol{\eta}$ is a $m \times 1$ vector of latent factors, and $\boldsymbol{\varepsilon}$ denotes the vector of multivariate normally distributed residuals with zero mean vector and covariance matrix $\boldsymbol{\Omega}$. In the following, we assume that the mean structure is saturated and completely reflected in the intercepts, that is $E(\mathbf{x}) = \mathbf{v}$.

Thus, the focus is on modeling the covariance structure $\boldsymbol{\Sigma}$ of the observed variables.

The covariance matrix of the observed variables $\boldsymbol{\Sigma}$ can be written as a function of the model parameters of the CFA

$$\boldsymbol{\Sigma}(\boldsymbol{\theta}) = \mathbf{\Lambda}\boldsymbol{\Phi}\mathbf{\Lambda}^T + \boldsymbol{\Omega} \quad (2)$$

where $\boldsymbol{\Phi}$ is the $m \times m$ covariance matrix of the latent factors, $\boldsymbol{\theta} = (\theta_1, \dots, \theta_q)$ is a $q \times 1$ vector that contains the q nonredundant parameters in $\mathbf{\Lambda}$, $\boldsymbol{\Phi}$, and $\boldsymbol{\Omega}$ that are estimated; and $\boldsymbol{\Sigma}(\boldsymbol{\theta})$ is the model-implied covariance matrix. Thus, the covariance of the observed variables can be decomposed into a part due to the covariance structure of the latent factors, and a part that is due to measurement error.

2.1. Maximum Likelihood Estimation

ML estimation has been routinely used to obtain parameter estimates of CFA models (Jackson et al., 2009). Let $\mathbf{x}_i$ denote a vector of p independently and identically distributed observed variables ($i = 1, \dots, n$) that is multivariate normally distributed. Then for an observed data set $\mathbf{X} = \{\mathbf{x}_i\}_{i=1, \dots, n}$, the likelihood function is written as

$$L(\boldsymbol{\theta}|\mathbf{X}) = \prod_{i=1}^n f(\mathbf{x}_i; \mathbf{v}, \boldsymbol{\Sigma}(\boldsymbol{\theta})) \quad (3)$$

where $f(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the multivariate normal density with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. It is known that the sample-based covariance matrix $\mathbf{S} = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})'$ is

a sufficient statistic for Σ , and hence for $\Sigma(\theta)$, which also implies sufficiency for θ . Thus, the likelihood can be written as $L(\theta|\mathbf{X}) = L(\theta|\mathbf{S})$, and the sample covariance matrix $\mathbf{S}$ of the p observed variables can be used as input in the SEM framework. The log-likelihood can be simplified as (Bollen, 1989)

$$l(\theta) = \log L(\theta|\mathbf{S}) = \frac{n}{2} [-p \cdot \log(2\pi) - \log|\Sigma(\theta)| - \text{tr}(\Sigma^{-1}(\theta)\mathbf{S})] \quad (4)$$

where tr is the trace operator, that is, the sum of the diagonal elements of a square matrix.

Note that $n\mathbf{S}$ follows a Wishart distribution (Hoff, 2009). The value $\hat{\theta}_{\text{ML}} = \arg \max_{\theta} l(\theta)$ that maximizes $l(\theta)$ is the ML estimate. It should be emphasized that the latent variables η do not appear in the likelihood in Equation 4. Therefore, it has also been referred to as the marginal likelihood where the latent variables are integrated out (Fox et al., 2017; Merkle et al., 2019).

The desirable properties of ML (e.g., most efficient estimates) are based on asymptotic theory and are only guaranteed to hold with large sample sizes (Yuan & Bentler, 2007). In small samples and complex models, ML is prone to serious estimation problems such as failure to converge or inadmissible solutions (e.g., negative variance estimates or correlations that are larger than 1; Anderson & Gerbing, 1984; Chen et al., 2001; Wothke, 1993; Yuan & Chan, 2008). Furthermore, in small to medium samples, SEMs that correct for measurement error, even though approximately unbiased, can produce much more variable estimates of structural relationships (i.e., larger empirical sampling variance) than biased manifest approaches that ignore measurement error and use manifest scale scores (Ledgerwood & Shrout, 2011; Savalei, 2019; see also Li & Beretvas, 2013; Zitzmann et al., 2016). This bias-variance trade-off can result in the counterintuitive finding that with small sample sizes, ML estimation of SEMs can provide less accurate estimates of structural relationships in terms of root mean squared error (RMSE) than a manifest approach which uses the error-prone scale scores (Carroll et al., 2006; Fuller, 1987).

2.2. Constrained Maximum Likelihood Estimation

As pointed out, in standard ML estimation, parameter estimates are not constrained to any specific interval, and, for example, nothing prevents variance estimates from becoming negative (Held & Bové, 2014; Savalei & Kolenikov, 2008). In order to mitigate estimation problems and avoid parameter estimates that are outside the admissible parameter space, researchers sometimes use *constrained* ML estimation for obtaining parameter estimates in SEM. For example, Lüdtke et al. (2018) showed in simulation studies that constrained ML estimation of the trait-state-error model for multiwave data (Kenny & Zautra, 1995) outperformed unconstrained ML estimation in terms of the frequency of estimation problems and the accuracy of the parameter estimates (see also Chen et al., 2001; Gerbing & Anderson, 1987).

In constrained estimation, the parameter space over which optimization is performed is restricted to admissible values (e.g., variances are constrained to be positive; Schoenberg, 1997). To this end, inequality constraints that restrict parameter estimates to lower and upper bounds must be specified (see Savalei & Kolenikov, 2008). More specifically, in the constraint estimation approach, a multivalued function $\mathbf{h}$ is specified on the vector of SEM parameters, that is $\mathbf{h}(\boldsymbol{\theta}) > \mathbf{0}$. For example, if a parameter θ (e.g., correlation) is supposed to be bounded by a lower bound l and an upper bound u , that is $l < \theta < u$, the constraints would be given as follows: $\mathbf{h}(\theta) = (\theta - l, u - \theta) > (0, 0)$. Further possible constraints include restricting factor loadings or residual variances to positive values. Note that the constrained ML estimator is the parameter vector $\boldsymbol{\theta}$ that maximizes the log-likelihood in Equation 4 and fulfills the constraints that are imposed in $\mathbf{h}$. In most SEM programs such as *Mplus* (Muthén & Muthén, 2012) and *lavaan* (Rosseel, 2012), unconstrained ML estimation, that does not impose any restrictions on the admissible parameter space, is used as the default (see Kline, 2016). In the present article, we compare the performance of constrained and unconstrained

ML estimation of CFA models with Bayesian estimation techniques. These are discussed in the next section.

2.3. Bayesian Approach to Confirmatory Factor Analysis

In the Bayesian approach, statistical inference is based on the posterior distribution, which is determined by the likelihood function and the prior distribution $\pi(\boldsymbol{\theta})$ of model parameters (for a general introduction to the Bayesian approach, see Gelman et al., 2014; Jackman, 2009; Lynch, 2007). Using the observed data $\mathbf{X}$ (or the sufficient statistic $\mathbf{S}$) and the prior distributions, the joint posterior distribution $p(\boldsymbol{\theta}|\mathbf{X})$ of the parameters is determined by multiplying the likelihood with the prior:

$$p(\boldsymbol{\theta}|\mathbf{X}) = \frac{L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})}{\int L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}} = L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})C \propto L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta}) \quad (5)$$

where $C = 1/\int L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}$ is a normalizing constant. As can be seen, the posterior distribution is proportional to the product of the likelihood and the prior. If a researcher is not willing to make assumptions about a parameter, noninformative prior distributions that have only a minimal influence on the results are selected. Besides, the Bayesian approach offers the opportunity to stabilize parameter estimates by specifying a weakly informative prior distribution $\pi(\boldsymbol{\theta})$ “which contains some information—enough to ‘regularize’ the posterior distribution, that is, to keep it roughly within reasonable bounds—but without attempting to fully capture one’s scientific knowledge about the underlying parameter” (Gelman et al., 2014, pp. 51). Thus, the idea is to incorporate a small amount of information in $\pi(\boldsymbol{\theta})$ that provides some direction for the estimation of model parameters but, at the same time, still allows the inferences to be driven by the likelihood (Baldwin & Fellingham, 2013; Chung et al., 2013; Depaoli & Clifton, 2015; Lüdtke et al., 2013). In the following, we distinguish two methods for obtaining Bayesian estimates from the posterior distribution: PML and MCMC methods.

2.3.1. Penalized ML Estimation

The close relation of the Bayesian approach to ML estimation becomes evident by considering the log-posterior function, which is given by applying the logarithm to Equation 5 (and omitting the normalizing constant C)

$$w(\boldsymbol{\theta}) = l(\boldsymbol{\theta}) + \log \pi(\boldsymbol{\theta}) \quad (6)$$

The log-posterior $w(\boldsymbol{\theta}) = \log p(\boldsymbol{\theta}|\mathbf{X})$ is a function of the log-likelihood $l(\boldsymbol{\theta}) = \log L(\boldsymbol{\theta}|\mathbf{X})$ and additional information that is given by the prior $\log \pi(\boldsymbol{\theta})$. To obtain Bayesian estimates of the SEM parameters, the log-posterior function (i.e., the sum of the log-likelihood and the logarithm of the prior density) could be directly maximized to calculate the multivariate mode of the posterior distribution. This is also referred to as the maximum a posteriori (MAP) estimator (Gelman et al., 2014). Alternatively, the logarithm of the prior distribution $\log \pi(\boldsymbol{\theta})$ can be interpreted as a penalty term that is added to the log-likelihood function, which motivates the label “penalized” ML (Cole et al., 2014; see also Cousineau & Helie, 2013). The value $\hat{\boldsymbol{\theta}}_{\text{PML}} = \arg \max_{\boldsymbol{\theta}} w(\boldsymbol{\theta})$ that maximizes w is the PML estimate. Note that with a noninformative prior (i.e., π is a constant function with respect to $\boldsymbol{\theta}$), the $\hat{\boldsymbol{\theta}}_{\text{PML}}$ coincides with the ML estimate. In addition, it should be emphasized that constrained ML can be regarded as a variant of PML when uniform prior distributions are imposed on the admissible parameter space (e.g., Rindskopf, 2012).

PML estimation has been shown to circumvent estimation problems, particularly when the likelihood is flat, and has been successfully applied in a wide range of models to stabilize parameter estimates such as logistic regression models (Firth, 1993; Heinze & Schemper, 2002), latent class models (DeCarlo, Kim, & Johnson, 2011; Galindo-Garre & Vermunt, 2006), item response theory models (Harwell & Baker, 1991; Mislevy, 1986), or multilevel models (Chung et al., 2013). PML estimation can also be considered as a regularization

technique that has been proposed as a broad class of methods for improving the stability of parameter estimation (Jacobucci & Grimm, 2018; Fan et al., 2020).

2.3.2. *MCMC Estimation*

Another strategy for obtaining Bayesian estimates is to apply simulation-based techniques. This is motivated by the fact that in practice, the joint posterior distribution of the parameters is often difficult to evaluate because high-dimensional integration is required for computing the normalization constant C (see Equation 5; Held & Bové, 2014, Ch. 8). Simulation-based techniques are implemented in most general Bayesian software such as WinBUGS (Spiegelhalter et al., 2003), JAGS (Plummer, 2003), NIMBLE (de Valpine et al., 2017), or Stan (Carpenter et al., 2017) and use MCMC algorithms to approximate the posterior distribution by iteratively sampling from conditional distributions. The most prominent MCMC methods are Gibbs sampling, Metropolis-Hastings sampling, and the no-U-turn sampler (Gelman et al., 2014; Junker, Patz, & VanHoudnos, 2016).

In these methods, point and interval estimators are constructed from the posterior distribution by iteratively sampling from conditional distributions, constituting an MCMC chain ($t = 1, \dots, T$):

$$\boldsymbol{\theta}^{(t)} = \left(\theta_1^{(t)}, \dots, \theta_d^{(t)}, \dots, \theta_q^{(t)} \right) \sim p(\boldsymbol{\theta}|\mathbf{X}) \propto L(\boldsymbol{\theta}|\mathbf{X}) \cdot \pi(\boldsymbol{\theta}) \quad (7)$$

Usually, the initial samples are discarded (burn-in period) because the initial values are affected by the starting values (for a detailed discussion of assessing convergence in MCMC, see Cowles & Carlin, 1996; Gelman et al., 2014; Gill, 2007). In MCMC, the mean and the mode are often used as Bayesian estimates of the SEM parameters. More specifically, the expected a posteriori (EAP) estimator for a single SEM parameter θ_d is obtained by averaging across the T iterations, that is $\hat{\theta}_{d,\text{EAP}} = T^{-1} \sum_{t=1}^T \theta_d^{(t)}$. The mode $\hat{\theta}_{d,\text{MAP}}$ can be defined as the univariate mode of the kernel density estimate (Silverman, 1998) of the univariate density for the sample of $\theta_d^{(t)}$ ($t = 1, \dots, T$). It is important to understand that both estimates, $\hat{\theta}_{d,\text{EAP}}$ and

$\hat{\theta}_{d,\text{MAP}}$, target location estimators of the marginal univariate density p_d of the d th component of $\boldsymbol{\theta}$ in which the other components are integrated out. More formally, let $\boldsymbol{\theta}_{(-d)} = (\theta_1, \dots, \theta_{d-1}, \theta_{d+1}, \dots, \theta_q)$ denote the vector of parameters in which the d th entry of $\boldsymbol{\theta}$ has been omitted, then the univariate marginal posterior distribution of θ_d in which all other components of $\boldsymbol{\theta}$ are integrated out, is given by

$$p_d(\theta_d|\mathbf{X}) = \int p(\boldsymbol{\theta}|\mathbf{X})d\boldsymbol{\theta}_{(-d)} = C \int L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}_{(-d)} \quad (8)$$

and the corresponding location parameters are defined as

$$\hat{\theta}_{d,\text{EAP}} = \int \theta_d p_d(\theta_d|\mathbf{X})d\theta_d = \frac{\int \theta_d L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}}{\int L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}} \quad \text{and} \quad (9)$$

$$\hat{\theta}_{d,\text{MAP}} = \arg \max_{\theta_d} p_d(\theta_d|\mathbf{X}) = \arg \max_{\theta_d} \int L(\boldsymbol{\theta}|\mathbf{X})\pi(\boldsymbol{\theta})d\boldsymbol{\theta}_{(-d)} \quad (10)$$

Note that both $\hat{\theta}_{d,\text{EAP}}$ and $\hat{\theta}_{d,\text{MAP}}$ are functionals of the joint posterior distribution. For a sufficiently large number of MCMC samples T , these quantities are approximated by the empirical posterior distribution of the samples. The EAP $\hat{\theta}_{d,\text{EAP}}$ obtained from the univariate marginal posterior distribution p_d coincides with the EAP obtained from the multivariate posterior distribution p . However, it is crucial that the univariate modes $\hat{\theta}_{d,\text{MAP}}$ ($d = 1, \dots, q$) of the marginal posterior distributions may not equal the components of the mode $\hat{\boldsymbol{\theta}}_{\text{PML}}$ of the joint posterior distribution (Held & Bové, 2014). Thus, Bayesian estimation obtained with MCMC that utilizes the marginal posteriors (i.e., $\hat{\theta}_{d,\text{EAP}}$, and $\hat{\theta}_{d,\text{MAP}}$) will, in general, deviate from the PML estimates that are obtained from the joint posterior (i.e., $\hat{\boldsymbol{\theta}}_{\text{PML}}$), particularly in conditions with small sample sizes and skewed marginal posterior distributions. The consequences of obtaining Bayesian estimates from either the joint posterior or the marginal posterior will be further illustrated for a CFA with two latent factors in the next section.

Finally, it should be emphasized that in the presented MCMC approach, the parameter estimates are based on the marginal likelihood $L(\boldsymbol{\theta}|\mathbf{X})$ —or $L(\boldsymbol{\theta}|\mathbf{S})$ if the sufficient statistic $\mathbf{S}$ is used—in which the latent variables $\boldsymbol{\eta}$ are integrated out (in Equation 4). However, in many

applications of MCMC-based SEM, a joint estimation approach is used that relies on the joint likelihood $L(\boldsymbol{\theta}, \boldsymbol{\eta}|\mathbf{X})$, which also includes the latent variables $\boldsymbol{\eta}$ in the likelihood (Lee, 2007; Muthén & Asparouhov, 2012). In the joint estimation approach, the MCMC method¹ generates samples for $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$. When individual factor scores are not of interest, however, the latent variables $\boldsymbol{\eta}$ are nuisance parameters that may deteriorate the computational efficiency of the MCMC algorithm (Choi & Levy, 2017; Hayashi & Arav, 2006; Hecht et al., 2020; Lüdtke et al., 2018; Merkle et al., 2020).

3. CFA with Two Factors: Parameterization, Prior Distributions, and Estimation

Methods

In the following, we consider a CFA model with two latent variables, which are each measured by three observed indicator variables (see Figure 1). We use this model to discuss three issues that are of relevance in the practical application of the Bayesian approach: model parameterization, specification of prior distributions, and different estimation methods (PML and MCMC). Although this is a very simple model, it is a building block for many more complicated SEM models, such as latent mediation models or multilevel SEMs (Hoyle, 2012; Kline, 2016).

3.1. Parameterization

For mitigating estimation problems, it can be advantageous to choose a parameterization of the CFA that transforms the optimization problem of estimating unbounded parameters into an optimization involving bounded parameters. Furthermore, a parameterization with bounded or standardized parameters has the advantage that it is often more convenient for applied researchers for specifying thoughtful prior distributions (Smid et al., 2020; Zitzmann et al., 2020). Thus, we suggest employing a parameterization of the CFA

¹ In the joint estimation approach, the joint likelihood $L(\boldsymbol{\theta}, \boldsymbol{\eta}|\mathbf{X}) = \prod_{i=1}^n f(\mathbf{x}_i; \boldsymbol{\nu} + \boldsymbol{\Lambda}\boldsymbol{\eta}_i, \boldsymbol{\Omega}) \cdot f(\boldsymbol{\eta}_i; \boldsymbol{\theta}, \boldsymbol{\Phi})$ is considered, and the MCMC method generates samples for $\boldsymbol{\theta}$ and $\boldsymbol{\eta}$ ($t = 1, \dots, T$): $\boldsymbol{\theta}^{(t)}, \boldsymbol{\eta}^{(t)} \sim p(\boldsymbol{\theta}, \boldsymbol{\eta}|\mathbf{X}) \propto L(\boldsymbol{\theta}, \boldsymbol{\eta}|\mathbf{X}) \cdot \pi(\boldsymbol{\theta})$. Note that the joint estimation approach needs the raw data $\mathbf{X}$ instead of the sufficient statistic $\mathbf{S}$ (see Choi & Levy, 2017).

which directly estimates standardized loadings and the latent correlation. Let us assume that the two latent variables η_1 and η_2 are measured by indicators x_j ($j = 1, \dots, 6$)

$$x_j = \sigma_j(\lambda_j \eta_{m[j]} + \varepsilon_j) \quad (11)$$

where $m[\cdot]$ is a function that maps the item index j to the corresponding index $m[j]$ of the latent variable and σ_j is the standard deviation of the observed indicator x_j . The variances of the latent variables are set to one, i.e., $\text{Var}(\eta_1) = \text{Var}(\eta_2) = 1$, and the parameters λ_j are the standardized loadings. Note that the variance of an observed variable j is given by $\sigma_j^2 = \text{Var}(x_j) = \lambda_j^2 + \omega_{jj}$. Thus, Equation 11 can be written as $x_j = \sigma_j[\underbrace{(\tilde{\lambda}_j/\sigma_j)}_{\lambda_j} \eta_{m[j]} + \underbrace{(\tilde{\varepsilon}_j/\sigma_j)}_{\varepsilon_j}]$,

where $\tilde{\lambda}_j$ are the unstandardized loadings, and $\tilde{\varepsilon}_j$ the residuals of the unstandardized solution.

Technically, this parameterization can be implemented in the SEM framework by introducing an intermediate layer of latent variables (phantom variables; see Little, 2013). The error variances for the indicators are given by $\omega_{jj} = \text{Var}(\varepsilon_j) = 1 - \lambda_j^2$. Our main focus is on estimating the correlation between the latent variables, i.e., $\text{Cov}(\eta_1, \eta_2) = \rho$.²

As we are interested in estimating the latent correlation that corrects for the unreliability of the scale scores, it is instructive to see how the reliability of the manifest sum score is related to the standardized loadings of the measurement model. We assume that the standardized loadings of the indicators for a latent factor are equal and set to λ (tau-congeneric measurement model; Traub, 1994). Then, the indicator-specific reliability is given by $\text{Rel}_1 = \lambda^2$, and the reliability of the sum score of I items is

$$\text{Rel}_I = \frac{\lambda^2}{\lambda^2 + (1 - \lambda^2)/I} = \frac{\text{Rel}_1}{\text{Rel}_1 + (1 - \text{Rel}_1)/I} \quad (12)$$

² For constrained ML estimation, it can be shown that this parameterization of the CFA with standard deviations that are constrained to be positive and loadings that are restricted to the interval $[0, 1]$ is (analytically) equivalent to a parameterization in which the loadings and error variances of the unstandardized solution are constrained to be positive, and the variances of the latent factors are set to one. Thus, using the parameterization of the CFA model in Equation 11 for constrained ML estimation should—in theory—provide the same results as a CFA model with the corresponding inequality constraints on the loadings and error variances.

As can be seen, the reliability of the sum score Rel_I is a function of the indicator-specific reliability Rel_1 and the number of items. Thus, by rearranging terms, the reliability of an indicator can be written as

$$Rel_1 = \frac{Rel_I}{1 + I(1 - Rel_I)} \quad (13)$$

For example, with $I = 3$, a standardized loading of $\lambda = .58$ translates into a reliability of .60 for the sum score (see Table 1). This relationship between the standardized loading and the reliability of the sum score will be helpful when specifying the prior distributions for the loadings because, in most cases, it is easier to make plausible assumptions about the overall reliability of a scale than for every single item (Smid et al., 2020).

3.2. Specification of Prior Distributions

In the CFA model, the standardized loadings and latent correlations are bounded parameters. For bounded parameters, the beta distribution is a natural choice. The density of the beta distribution $X \sim \text{Beta}(a, b)$ on the interval $[0, 1]$ is given as:

$$f(x) = B_{a,b}^{-1} x^{a-1} (1-x)^{b-1}, \quad x \in [0, 1] \quad (14)$$

where $B_{a,b}$ is a normalizing constant. The mean and the variance can be computed as

$$E(X) = \frac{a}{a+b} \quad \text{and} \quad \text{Var}(X) = \frac{ab}{(a+b)^2(a+b+1)} \quad (15)$$

Alternatively, the beta distribution can be parameterized as a function of a mean μ and a prior sample size v , that is $X \sim \text{Beta}(\mu, v)$, where $\mu = a(a+b)^{-1}$, and $v = a+b-2$ (Hoff, 2009). The prior sample size is motivated by the fact that the uniform distribution, which reflects complete ignorance about a parameter, is given by setting $a = b = 1$. Thus a prior sample size of $v = 1 + 1 - 2 = 0$ corresponds to the uniform prior on $[0, 1]$. The variance of the beta distribution is given as $\text{Var}(X) = \mu(1-\mu)(v+3)^{-1}$. For given μ and v , the original a and b parameters are determined by $a = (v+2)\mu$, and $b = (v+2)(1-\mu)$.

However, the beta distribution is only appropriate for parameters with a parameter space that equals $[0, 1]$. The four-parameter beta distribution extends the beta distribution to

arbitrary intervals and allows a more flexible specification of prior distributions (see also Zeng, 1997). The four-parameter beta distribution $Y \sim \text{Beta4}(a, b, l, u)$ can be obtained by shifting a beta distributed random variable $X \sim \text{Beta}(a, b)$ by lower (l) and upper (u) bounds: $Y = l + (u - l)X$. The density of Y is then given as

$$f(x) = (u - l)^{-1} B_{a,b}^{-1} \left(\frac{x-l}{u-l} \right)^{a-1} \left(\frac{u-x}{u-l} \right)^{b-1}, \quad x \in [l, u] \quad (16)$$

Again, the four-parameter beta distribution can be reparameterized as $Y \sim \text{Beta4}(\mu, v, l, u)$ with prior guess $\mu = a(a + b)^{-1}$, and prior sample size $v = a + b - 2$. The parameters of the original specification $Y \sim \text{Beta4}(a, b, l, u)$ can be obtained as $a = (v + 2)(\mu - l)(u - l)^{-1}$, and $b = (v + 2)(u - \mu)(u - l)^{-1}$.

When specifying the lower and upper bounds of the four-parameter beta distribution, it can improve the numerical stability of parameter estimates to further coerce parameters away from the boundaries by a small value ε (e.g., $\varepsilon = .001$). For a standardized loading λ that can be assumed to be bounded between (0, 1), we suggest a $\text{Beta4}(\mu_\lambda, v_\lambda, \varepsilon, 1 - \varepsilon)$ and interpret μ_λ as a prior guess for the standardized loading and v_λ as the sample size of prior observations on which the prior guess is based (Lüdtke et al., 2018). If only little information is available about the standardized loading or the reliability of a scale, a small value for v_λ would be chosen so that the prior distribution will be only weakly centered around the prior guess μ_λ . Figure 2 (left panel) shows for a prior guess of $\mu_\lambda = .50$ (i.e., standardized loading of .50) how increasing v_λ (i.e., prior sample sizes of 1, 3, and 10) changes the shape of the four-parameter beta distribution. Note that with $\mu_\lambda = .50$ and $v_\lambda = 0$, the four-parameter beta distribution corresponds to a uniform distribution on the interval $[\varepsilon, 1 - \varepsilon]$, which reflects complete ignorance about the size of the loading.

For the latent correlation that is restricted to the interval $[-1, 1]$, we suggest a $\text{Beta4}(\mu_\rho, v_\rho, -1 + \varepsilon, 1 - \varepsilon)$ distribution where μ_ρ is the prior guess for the correlation and v_ρ is again the prior sample size on which the prior guess is based. Figure 2 (right panel) illustrates

the influence of the prior sample size v_ρ (1, 3, and 10) on the shape of the Beta4(μ_ρ , v_ρ , $-1 + \varepsilon$, $1 - \varepsilon$) with a prior guess of $\mu_\rho = .30$. Setting $\mu_\rho = 0$, and $v_\rho = 0$ gives the uniform distribution on $[-1 + \varepsilon, 1 - \varepsilon]$.

3.3. Illustrative Comparison of PML (Constrained ML) and MCMC

To illustrate how the two Bayesian estimation methods, PML and MCMC, can produce different estimates of the correlation, we further simplify the two-factor model and assume that all loadings are equal. Thus, we only need to estimate two parameters³: the correlation ρ (ranging between -1 and 1), and the standardized loading λ (ranging between 0 and 1). Thus, for this simplified model, the likelihood function $L(\rho, \lambda | \mathbf{S})$ is only a function of ρ and λ , given the sufficient statistic $\mathbf{S}$. Furthermore, we assume uniform priors for both parameters (i.e., constant functions with respect to ρ and λ) which results in a joint posterior $p(\rho, \lambda | \mathbf{S})$ that is proportional to the likelihood. In this case, PML is equivalent to constrained ML. We now show how PML (or constrained ML) and MCMC differ in their use of $p(\rho, \lambda | \mathbf{S})$ to obtain point estimates of the latent correlation ρ .

In PML estimation, the mode of the joint posterior distribution is calculated as

$$(\hat{\rho}_{\text{PML}}, \hat{\lambda}_{\text{PML}}) = \arg \max_{(\rho, \lambda)} p(\rho, \lambda | \mathbf{S}) = \arg \max_{(\rho, \lambda)} L(\rho, \lambda | \mathbf{S}), \quad (17)$$

and $\hat{\rho}_{\text{PML}}$ is used as a point estimate of ρ . Note that this is the first component of the multivariate mode, which is calculated by directly maximizing the density of the joint posterior distribution. It becomes clear from Equation 17 that the PML estimate is also the constrained ML estimate because the likelihood function is maximized.

In contrast, in the MCMC method, the univariate mode (MAP) and mean (EAP) are often used as point estimates for ρ . In order to calculate the mean, the marginal posterior

³ For this illustration, we further reduced the number of estimated parameters and assumed that the indicators have a variance of one, and also fixed the variances of the indicators in the analysis model to one. The main purpose of the illustration was to demonstrate the differences between the mode from the joint posterior that is used in PML (or constrained ML), and the EAP that is obtained from the marginal distribution in MCMC. A more systematic and realistic evaluation of the different estimators is provided in the main simulation.

distribution p_ρ of ρ is needed, which is obtained by integrating the joint posterior $p(\rho, \lambda | \mathbf{S})$ with respect to λ :

$$p_\rho(\rho | \mathbf{S}) = \int p(\rho, \lambda | \mathbf{S}) d\lambda = C \int L(\rho, \lambda | \mathbf{S}) d\lambda \quad (18)$$

where $C = 1 / \iint L(\rho, \lambda | \mathbf{S}) d\rho d\lambda$ is the normalizing constant. The corresponding marginal location parameters are given as follows

$$\hat{\rho}_{\text{EAP}} = \int \rho p_\rho(\rho | \mathbf{S}) d\rho = \frac{\iint \rho L(\rho, \lambda | \mathbf{S}) d\rho d\lambda}{\iint L(\rho, \lambda | \mathbf{S}) d\rho d\lambda}, \quad \text{and} \quad (19)$$

$$\hat{\rho}_{\text{MAP}} = \arg \max_{\rho} p_\rho(\rho | \mathbf{S}) = \arg \max_{\rho} \int L(\rho, \lambda | \mathbf{S}) d\lambda \quad (20)$$

In our simple bivariate case, the univariate EAP $\hat{\rho}_{\text{EAP}}$ can be calculated by numerically evaluating the posterior on a bivariate discrete grid of values ρ and λ . The univariate MAP $\hat{\rho}_{\text{MAP}}$ can be obtained by maximization of the integrated likelihood, using numerical integration with respect to λ . However, this would not be practical with a larger number of parameters, and simulation-based MCMC techniques are needed to evaluate high-dimensional integrals (Held & Bové, 2014).

We now employ an idealized scenario to illustrate the difference between the different Bayesian estimates of the latent correlation. In this case, the empirical covariance matrix $\mathbf{S}$ (i.e., the sufficient statistic) obtained from data was set equal to the true covariance matrix $\Sigma = \Sigma(\theta)$. Hence, the likelihood estimates (i.e., the constrained ML and the PML estimate) coincide with the data-generating parameters. The true correlation was $\rho = .7$, and the standardized loading was $\lambda = .50$. Figure 3 shows for a small sample size of $N = 30$ a contour plot of the joint posterior distribution for ρ and λ (upper left panel) and the marginal posterior distribution of ρ (lower left panel). As can be seen, the mode of the joint posterior ($\hat{\rho}_{\text{PML}} = .700$) provides a different Bayesian estimate of the correlation than the mode ($\hat{\rho}_{\text{MAP}} = .710$) or mean ($\hat{\rho}_{\text{EAP}} = .570$) of the marginal posterior. This can be explained by the fact that the marginal posterior is negatively skewed, and the mean is pulled toward zero.

However, with a larger sample of $N = 100$, the Bayesian estimates from the joint posterior (upper right panel) and marginal posterior (lower right panel) more closely agree ($\hat{\rho}_{\text{PML}} = .700$, $\hat{\rho}_{\text{MAP}} = .700$, and $\hat{\rho}_{\text{EAP}} = .680$), and the marginal posterior distribution of ρ is more symmetrically centered around the true value of .70.

The convergence of the different Bayesian estimates to the true parameters as a function of the sample size is further illustrated in Figure 4. We followed the same logic holding the covariance matrix $\mathbf{S} = \mathbf{\Sigma}$ fixed, and only varied the sample size (from $N = 20$ to 100). As can be seen, with increasing sample size, the differences between the mean (EAP) and mode (MAP) of the marginal posterior, and the mode from the joint posterior (PML or the constrained ML) vanishes.

3.3.1. *Illustrative Simulation Study*

To further explore how these differences between the Bayesian point estimates (PML [or constrained ML], MAP, and EAP) affect their frequentist properties, we ran a small simulation study in which we manipulated the sample size ($N = 30, 50, 100$, and 1000), and the magnitude of the true correlation ($\rho = .10, .30, .50, .70$, and $.90$). The standardized loading was set to .50. We generated 1000 replications for each condition and calculated the bias, variability (i.e., the standard deviation of the empirical sampling distribution), and the RMSE (which combines bias and variability into a measure of accuracy) for the different point estimates (PML, MAP, and EAP) of the latent correlation ρ .

The results are shown in Table 2 and confirm the findings from the illustration that the mode from the joint posterior (PML or constrained ML) and the mode from the marginal posterior (MAP) perform very similarly. Both produced approximately unbiased estimates of the latent correlation, with the exception of the condition with a very large correlation ($\rho = .90$), and a small sample size ($N = 30$). By contrast, the mean of the marginal posterior (EAP) provided negatively biased estimates, particularly in conditions with small sample size.

However, the EAP was also less variable (i.e., smaller empirical sampling variability) than the estimates produced by PML and MAP, resulting in overall more accurate estimates in terms of the RMSE, which combines bias and variability. The results also show that there is a turning point that, with a larger true correlation, the bias introduced by the EAP outweighs the gains in efficiency (i.e., less variable estimates of the EAP). Thus, the EAP seems to be most beneficial with a small to moderate true correlation (i.e., $\rho \leq .50$), and will not generally result in more accurate estimates of the latent correlation. Notably, the multivariate mode (PML or constrained ML) performed similarly to the univariate mode (MAP). However, in other models, particularly with strongly depended parameter estimates (Lüdtke et al., 2018), the multivariate and univariate mode can provide substantially different point estimates.

Finally, the findings confirm that with large samples, the different Bayesian point estimates converge and produce almost identical results. It is important to note that the EAP provides more efficient estimates than PML (or constrained ML) even though only uniform (and noninformative) prior distributions were chosen. Hence, the performance gain of Bayesian estimation compared to constrained ML only stems from the fact that the EAP was used as a point estimate. The results of this illustration directly transfer to MCMC estimation because MCMC sampling is only used to approximate the evaluation of integrals needed for computing the EAP (Gelman et al., 2014). With uniform priors, EAPs obtained from MCMC can also be expected to outperform constrained ML (or PML).

4. Simulation Study

The simulation study has two main goals. First, we evaluate the performance of the two Bayesian estimation methods (PML and MCMC) and compare them with traditional ML estimation. For small sample sizes, we expect that unconstrained ML estimation shows serious estimation problems (i.e., nonconvergence or inadmissible parameter estimates). In addition, on the basis of our illustration, we assume that using MCMC together with the EAP

as a point estimate should produce more stable estimates of latent correlations than PML estimation, particularly in conditions with small sample size and small to moderate correlations. Second, we evaluate the extent to which the parameter estimates of the Bayesian approach are sensitive to different specifications of the prior distributions for the standardized loadings and the latent correlation. We assume that by choosing weakly informative and correctly specified prior distributions (i.e., four-parameter beta distributions), the estimates of latent correlations can be stabilized. Furthermore, we explore whether the Bayesian approach produces more accurate estimates, even with mildly misspecified prior distributions. Overall, we expect that the impact of choosing different prior distributions should be more pronounced with small sample sizes.

4.1. Simulation Model and Conditions

The data-generating model was a two-factor CFA, as given by Figure 1. Each factor was measured by three mean-centered and normally distributed indicators. The indicators were assumed to be parallel with standardized loadings of .50, and a variance of one. This resulted in a reliability of $Rel_I = .50$ for each scale (i.e., sum score of the three items), and a reliability of $Rel_1 = .25$ for a single indicator. The following two conditions were manipulated: the latent correlation between the two factors ($\rho = .30, .50, \text{ and } .70$) and the sample size ($N = 30, 50, \text{ and } 100$). For each of the $3 \times 3 = 9$ conditions, we generated 1,000 simulated data sets.

4.2. Analysis Models and Outcomes

Each of the simulated data sets was analyzed with a two-factor CFA model in which the loadings were freely estimated, and the variances of the two factors were each fixed to one. The model has $21 - 13 = 8$ degrees of freedom (the mean structure was assumed to be saturated). Two ML estimation approaches were used. In unconstrained ML estimation, we imposed no constraints on the parameter estimates (loadings, residual variances, and the latent correlation). The variances of the two factors were set to 1. In constrained ML estimation, we

used the parameterization in which standard deviations of the indicators are constrained to be positive, the standardized loadings are restricted to the interval $[0, 1]$, and the latent correlation is restricted to the interval $[-1, 1]$ ⁴. In the Bayesian approach, we used PML and MCMC as estimation methods and varied the prior distribution for the standardized loadings and the latent correlation. PML estimation was implemented using a quasi-newton optimization (employing the `nlmminb` optimizer in R) using the wrapper function `pml_e` from the R package LAM (Robitzsch, 2020). The standard errors were calculated based on the second derivatives of the observed log-likelihood function (ignoring the part of the prior distribution; see also Dolan & Molenaar, 1991; Lüdtke et al., 2018; cf. Held & Bové, 2014). The estimated standard errors were used to calculate 95% confidence intervals.

For the MCMC method, we implemented an adaptive Metropolis-Hastings algorithm in which the proposal distribution is adapted during the burn-in phase (Lüdtke et al., 2018; see also Draper, 2008). This algorithm was implemented using function `amh` from the R package LAM (Robitzsch, 2020). Before running the main simulation study, we investigated the behavior of the MCMC sampler in preliminary simulations by inspecting two criteria: (a) the potential scale reduction factor (PSR; Gelman et al., 2014), and (b) the effective sample size (see Zitzmann & Hecht, 2019, for a discussion). Applying these two criteria suggested that an average chain length of 5,000 iterations with (a burn-period of 2,500 iterations) was sufficient to provide a good approximation of the posterior distribution. The Bayesian point estimate was defined as the mean of the posterior distribution (EAP). Furthermore, the Bayesian credible interval (BCI) was defined by the 2.5th and the 97.5th percentiles of the posterior distribution (Gelman et al., 2014)

⁴ For constrained ML estimation, we also included the equivalent parameterization (see Equation 11) in which the loadings and residual variances were constrained to be positive, and the latent correlation was restricted to the interval $[-1, 1]$. As expected, the results for both parameterizations were virtually identically in every replication (results were numerically identical in 95.5% of the replications for $N = 30$, in 98.4% for $N = 50$, and in 99.9% for $N = 100$). Moreover, the parameterization using standardized instead of unstandardized loadings performed slightly better in terms of RMSE.

For both PML and MCMC, we varied the prior distributions for the standardized loadings and the latent correlation. For each standardized loading, we specified a four-parameter beta distribution $\text{Beta4}(\mu_\lambda, v_\lambda, \varepsilon, 1 - \varepsilon)$ with a prior guess of $\mu_\lambda = .5$, and prior sample sizes of $v_\lambda = 1$, and $v_\lambda = 3$. In addition, we included a prior distribution with $\mu_\lambda = .5$ and $v_\lambda = 0$, which corresponds to a uniform distribution on $[\varepsilon, 1 - \varepsilon]$. For the latent correlation, we specified a four-parameter beta distribution $\text{Beta4}(\mu_\rho, v_\rho, -1 + \varepsilon, 1 - \varepsilon)$ with a prior guess that matched the true correlation of the data-generating model (i.e., $\mu_\rho = \rho$), and two levels of prior sample sizes ($v_\rho = 1$ and 3). We also specified a prior distribution with $\mu_\rho = 0$, and $v_\rho = 0$, which corresponds to a uniform distribution on $[-1 + \varepsilon, 1 - \varepsilon]$. This resulted in $3 \text{ (loadings)} \times 3 \text{ (correlation)} = 9$ different specifications of the prior distributions. Note that these prior distributions were correctly specified and only differed in the amount of information that was incorporated in the prior specification (i.e., prior sample size). It should be emphasized that the different prior sample sizes for the prior distributions carry different amounts of relative information in the different conditions of the simulation (Hamra et al., 2013). For example, with a sample size of $N = 30$, a prior sample size of $v_\rho = 3$ for the prior distribution of the latent correlation would carry more relative information than with a sample size of $N = 100$.

For the standard deviations of the indicator variables, we used improper prior distributions that are constant (Muthén & Asparouhov, 2012). The specification of the improper prior distribution for the standard deviation was held constant across the conditions of the simulation and the analysis models.

We used three criteria to evaluate the different estimation approaches: bias, RMSE, and coverage rate. Bias was calculated by the difference between the mean parameter estimate and the true population parameter value from each design cell. We assessed the overall accuracy with the (empirical) RMSE, which combines the squared empirical bias and variance

of the parameter estimates into a measure of overall accuracy. Finally, we determined the coverage rate of the 95% confidence intervals. A coverage rate between 91% and 98% was considered acceptable (Muthén & Muthén, 2002). The coverage rate was only assessed for the conditions with a larger sample size ($N = 100$).

4.3. Results

We first report the results for the two ML estimation approaches. Second, we compare the two Bayesian estimation methods in the case of correctly specified prior distributions. Third, we investigate the impact of misspecified prior distributions on the performance of the Bayesian approach.

4.3.1. *ML estimation*

For unconstrained ML estimation, a solution was considered to show estimation problems when the algorithm did not converge using the defaults in the `nlminb` optimizer or when the algorithm converged to a solution that included an inadmissible estimate (i.e., correlation smaller than -1 or larger than 1). Table 3 shows that the frequency of estimation problems for unconstrained ML strongly depended on the sample size, and also the magnitude of the true correlation. For example, with $N = 30$ and $\rho = .30$, only 54.5% of the replications converged, and 51.3% of the replications provided converged solutions with admissible estimates. In contrast, for constrained ML estimation, all replications converged. However, in small samples and with a large correlation, a substantial percentage of the solutions for constrained ML showed values at the boundary of the parameter space (e.g., estimated correlation equals 1). Furthermore, the results also show that with increasing sample size, the estimates of unconstrained ML and constrained ML converged to each other. For example, in the condition with $N = 100$ and a large correlation ($\rho = .70$), unconstrained and constrained estimation provided (numerically) identical estimates of the latent correlation in 91.6% of the replications.

Table 4 shows the bias and RMSE for unconstrained and constrained ML estimation as a function of the sample size and the true correlation. For unconstrained estimation, we only included converged solutions, and inadmissible values (i.e., correlations smaller than -1 or larger than 1) were truncated to -1 or 1 . For constrained estimation, all solutions were included (also the ones with boundary values). As can be seen, the two approaches perform very similarly, with a slight advantage for unconstrained ML in conditions with small sample size. However, this is not a fair comparison because the results for unconstrained ML are based on a highly positively selected set of replications given that a substantial proportion of the solutions for unconstrained ML did not converge (see Table 3). Furthermore, the results for unconstrained ML strongly depend on the cut-off values that would be used for the truncation. Figure 5 illustrates the effect of choosing different cut-off values used for trimming (between 1 and 10) on the bias and RMSE of the parameter estimates (in a condition with $N = 30$, and $\rho = .70$). In addition to truncation, the figure also includes trimming, which eliminates solutions with estimates that are above a specific cut-off value. For example, with a cut-off value of 1, a solution that provides an estimate of $\hat{\rho} = 1.30$ would be excluded from trimming (ML trimmed in Figure 3), and the estimate would be set to 1 with truncation (ML truncated in Figure 3). As can be seen, constrained ML estimation clearly outperformed constrained ML with trimming, unless the cut-off value is set to 1, and many replications are excluded from the analyses. In the following, we use constrained ML estimation, which is equivalent to PML estimation with uniform distributions on the admissible parameter space, and compare it with the two Bayesian estimation methods.

4.3.2. Bayesian estimation with correctly specified priors

Table 5 shows the bias for PML estimation and MCMC with noninformative and different correctly specified prior distributions as a function of the sample size and the magnitude of the true correlation. In these specifications, the prior guesses for the correlation

(i.e., $\mu_\rho = .30, .50, \text{ or } .70$) as well as the standardized loading (i.e., $\mu_\lambda = .50$) were set to the true value (when the prior sample sizes of the corresponding priors were at least 1). Note that using a prior sample size of zero, μ_ρ was set to 0, and a uniform distribution was specified. As pointed out, with uniform distributions for the standardized loadings (i.e., $\mu_\lambda = .50$ and $v_\lambda = 0$) and the correlation (i.e., $\mu_\rho = 0$ and $v_\rho = 0$), PML estimation is equivalent to constrained ML estimation.

As can be seen, both PML and MCMC produced biased estimates of the correlation, particularly when the sample sizes were small ($N \leq 50$). In addition, increasing the prior sample from $v_\rho = 1$ to $v_\rho = 3$, and thereby selecting a more informative prior distribution for the latent correlation, decreased the bias for both PML and MCMC. This illustrates the fact that more precise estimates are obtained when correctly specified prior information is added to the model (Smid et al., 2020).

For the RMSE, which combines bias and variability of an estimator, we used PML with uniform prior distributions on the standardized loadings and the latent correlation as a reference method. As this specification of PML is equivalent to constrained ML, it allows a direct comparison of the Bayesian approaches with the best performing ML approach. The RMSE gain in Table 5 reports the relative gain of an estimator compared to the reference method (i.e., values larger/smaller than 100 indicate that the RMSE for the respective method is larger/smaller than for the reference method). The results show that the MCMC methods clearly outperformed PML estimation across all sample size conditions and true values of the latent correlation. As expected from the illustration, the differences between the mode of the joint posterior distribution (PML) and the mean (EAP) of the marginal posterior were most pronounced in conditions with a very small sample size ($N = 30$) and a small true correlation. For example, in the condition with $N = 30$, $\rho = .30$, and uniform prior distributions, the RMSE of the estimates produced by MCMC were only 71% as large as the estimates by PML. This is

an important finding because it clearly shows that, even with (noninformative) uniform distributions on the loadings and the correlation, using the MCMC method stabilizes the parameter estimates compared to PML (or constrained ML). This can be explained by the fact that the MCMC method uses the EAP as a point estimate for the correlation.

To further understand the RMSE differences, we calculated the empirical standard deviation (*SD*) of the estimators across the 1000 replications within each cell. The results show that the estimates of PML were consistently more variable (across the different prior specifications) than those of the MCMC method. For both estimation methods, PML and MCMC, selecting a more informative prior distribution for the correlation (e.g., $v_\rho = 3$ instead of $v_\rho = 1$) had a large positive effect on the accuracy of the parameter estimates. By contrast, choosing a more informative prior distribution for the standardized loadings did not consistently influence the accuracy of the estimates of the latent correlation. Thus, only adding information to the prior distribution for the parameter of interest helped to stabilize estimates of the latent correlation in small sample sizes.

Finally, we assessed the coverage rates for PML and MCMC in conditions with a large sample size ($N = 100$). As can be seen in Table 6, PML provided acceptable coverage rates with percentages close to the nominal 95%. By contrast, the coverage rates produced by MCMC were sometimes too low. However, these low coverage rates can be explained by the fact that MCMC was also more biased in these conditions.

4.3.3. Bayesian estimation with misspecified priors

We also investigated the impact of misspecified prior distributions. To this end, we conducted a simulation with the same data-generating model but held the true correlation constant ($\rho = .5$), and varied the sample size ($N = 30, 50$, and 100). We specified a wide range of four-parameter beta distributions for the standardized loadings and the correlation. For the standardized loadings, we specified a prior distribution with a prior guess of $\mu_\lambda = .50$, and

prior samples sizes of $v_\lambda = 0$, $v_\lambda = 1$, and $v_\lambda = 3$. Also, we included misspecified priors with a prior guess of $\mu_\lambda = .80$, and prior sample sizes of $v_\lambda = 1$, and $v_\lambda = 3$. For the latent correlation, we specified a prior distribution with a prior guess of $\mu_\rho = 0$, and prior samples sizes of $v_\rho = 0$, $v_\rho = 1$, and $v_\rho = 3$, but also included misspecified priors which underestimated (with prior guess $\mu_\rho = .20$) or overestimated (with prior guess $\mu_\rho = .80$) the true correlation. Again, each misspecified prior was included with prior sample sizes of $v_\rho = 1$, and $v_\rho = 3$. In total, we specified 5 (standardized loading) $\times 7$ (correlation) = 35 models with different prior specifications, and estimated them both with PML and MCMC.

Table 7 shows the bias and RMSE for $N = 30$, and $N = 100$. The main findings can be summarized as follows. First, even in the case of misspecified prior distributions, the MCMC method outperformed PML in terms of the accuracy of parameter estimation across all conditions and prior specifications. Second, a misspecified prior distribution for the loading had only a small and sometimes even positive effect on the RMSE. One possible explanation is that we only included misspecified priors that overestimated the true size of the loading (i.e., $\mu_\lambda = .80$). Overestimating the reliability of the indicators by assuming a large positive loading comes close to a manifest approach that ignores the unreliability of scales scores when calculating the correlation. However, with small sample sizes, it has been shown that a manifest approach can produce more accurate estimates of structural relationships than a latent approach that corrects for measurement error (e.g., Ledgerwood & Shrout, 2011; Lüdtke et al., 2008; Savalei, 2019). Third, for the prior distribution of the correlation, the results clearly show that overestimating the true size of the latent correlation (i.e., $\mu_\rho = .80$) had a more negative impact on the accuracy of the estimates in terms of RMSE than underestimating the size of the true correlation (i.e., $\mu_\rho = .20$). More importantly, choosing a small correlation of .20 as a prior guess for the prior distribution, even though misspecified, produced more accurate estimates of the correlation than the Bayesian approach with uniform

priors on the loadings and the correlation, particularly when the sample size was $N = 30$.

Thus, a conservative approach that uses smaller prior guesses for the latent correlation seems to be a promising strategy when the goal is to stabilize estimates of latent correlations with weakly informative prior distributions (i.e., prior sample sizes of 1 or 3).

5. Discussion

In the present article, we showed that a Bayesian approach has the potential to stabilize the parameter estimates of a CFA in small sample size conditions. We compared two Bayesian estimation methods, PML estimation, and MCMC techniques, and clarified that they use different Bayesian point estimates from the joint posterior distribution: the mode (PML) of the joint posterior distribution, and the mean (EAP) or mode (MAP) of the marginal posterior distribution. In a simulation study, we showed that the MCMC method outperformed PML and produced more accurate estimates of latent correlations. These performance gains can be explained by the fact that MCMC used the EAP as the point estimate. However, with larger sample sizes, the differences between PML and MCMC vanished, and both performed similarly. We also suggested the four-parameter beta distribution as a prior distribution for loadings and correlations and argued that it could often be advantageous to choose a parameterization in which the main parameters of interest are bounded. Another main finding of our simulation study was that selecting weakly informative four-parameter beta distributions as priors helped to further stabilize parameter estimates. Importantly, this was also the case when the prior was mildly misspecified.

The main limitation of our simulation study is that we used a very simple CFA with only two latent factors and a small number of items with no cross-loadings (i.e., simple structure). It is straightforward to extend the discussed approaches to models with more latent factors. In constrained ML estimation and PML, appropriate determinant constraints could be implemented to ensure positive definiteness of the latent correlation matrix (Rousseeux & Molenberghs, 1994; Wothke, 1993). For the MCMC method, one could introduce determinant

constraints in the Metropolis-Hastings step to check for the positive definiteness of the correlation matrix (see Browne, 2006).

In addition, we only assessed the quality of statistical inferences (i.e., coverage rates) with moderate sample sizes. Further research is needed to investigate different approaches for estimating standard errors with small sample sizes. One promising approach is to use resampling methods such as Jackknifing (Efron, 1981). Preliminary results have shown that Jackknifing can improve the quality of statistical inference in small samples but more simulation research is needed to evaluate its potential for CFA models.

The results of the present study could be extended into several directions. First, it would be interesting to further explore the potential of PML estimation for CFA models in challenging data constellations (e.g., small samples, complex models; Rosseel, 2020). PML estimation seems to be particularly promising when researchers do not need full access to the posterior distribution but are only interested in obtaining stable point estimates for the parameters of interest. In contrast to simulation-based MCMC techniques which can be slow and difficult to implement, PML estimation shares the advantage of traditional ML that a deterministic optimization of the log-posterior is performed with clear convergence criteria and reasonable computational efficiency (Cousineau & Helie, 2013). Second, data sets in psychological research often have a multilevel structure (e.g., individuals are nested within clusters/groups), and in many applications, it is of interest to analyze relationships among latent constructs at both levels of analysis (e.g., individual level, and group level; Heck & Thomas, 2015). However, it is a well-known finding in the multilevel literature that a substantial number of groups is needed to obtain stable parameter estimates of group-level relationships (Can et al., 2015; Keleva & Brandt, 2014; Li & Beretvas, 2013; Lüdtke et al., 2011). Thus, it would be an important topic for future research to extend the discussed Bayesian approaches to multilevel CFA models (Kim, Dedrick, Cao, & Ferron, 2016). Finally, it would be interesting to compare the two Bayesian methods to other approaches that

have been suggested as solutions for estimation problems in small sample size conditions (Rosseel, 2020). For example, a two-step approach, such as factor score regression, has been suggested as a robust alternative to SEMs in challenging data constellations (Smid & Rosseel, 2020). In addition, one could use alternative error correction approaches that introduce lower bounds to circumvent small estimates of reliability in order to stabilize the estimation of latent correlations (Grilli & Rampichini, 2011). Using these lower bounds l , can be translated into a uniform distribution of standardized loadings on the interval $[l, 1]$, which is a particular case of the four-parameter beta distribution. Since PML (and CML for that matter) with a uniform prior distribution on $[0,1]$ alleviates convergence issues, using lower bounds on the indicator-specific reliability larger than 0 possibly introduces too much information. In addition, parameter estimates could be quite sensitive to the subjective choice of lower bounds.

To conclude, this article showed that Bayesian methods have a great potential for estimating CFA models with small sample sizes. Using simulated data, we showed that the four-parameter beta distribution can be used as prior distribution for standardized loadings and latent correlations in order to stabilize parameter estimates in challenging data constellations. However, in real applications, the specification of prior distributions should be accompanied by a sensitivity analysis that tests how sensitive the resulting parameter estimates are to different specifications of prior information (Depaoli & van de Schoot, 2017).

References

- Anderson, J. C., & Gerbing, D. W. (1984). The effect of sampling error on convergence, improper solutions, and goodness-of-fit indices for maximum likelihood confirmatory factor analysis. *Psychometrika*, 49, 155–173.
- Baldwin, S. A., & Fellingham, G. W. (2013). Bayesian methods for the analysis of small sample multilevel data with a complex variance structure. *Psychological Methods*, 18, 151–164.
- Bollen, K. A. (1989). *Structural equations with latent variables*. New York, NY: Wiley.
- Boomsma, A. (1985). Nonconvergence, improper solutions, and starting values in LISREL maximum likelihood estimation. *Psychometrika*, 50, 229–242.
- Browne, W. J. (2006). MCMC algorithms for constrained variance matrices. *Computational Statistics and Data Analysis*, 50, 1655–1677
- Can, S., van de Schoot, R., & Hox, J. (2015). Collinear latent variables in multilevel confirmatory factor analysis: A comparison of maximum likelihood and Bayesian estimations. *Educational and Psychological Measurement*, 75, 406–427.
doi:10.1177/0013164414547959
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., . . . , Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76. doi:10.18637/jss.v076.i01
- Carroll, R. J., Ruppert, D., Stefanski, A., & Crainiceanu, C. (2006). *Measurement error in nonlinear models: A modern perspective*. London: Chapman & Hall/CRC.
- Chen, F., Bollen, K. A., Paxton, P., Curran, P. J., & Kirby, J. B. (2001). Improper solutions in structural equation models: causes, consequences, and strategies. *Sociological Methods & Research*, 29, 468–508.

- Chen, J., Choi, J., Weiss, B. A., & Stapleton, L. (2014). An empirical evaluation of mediation effect analysis with manifest and latent variables using markov chain monte carlo and alternative estimation methods. *Structural Equation Modeling*, 21, 253–262.
- Choi, J., & Levy, R. (2017). Markov chain monte carlo estimation methods for structural equation modeling: a comparison of subject-level data and moment-level data approaches. *Biometrics and Biostatistics International Journal*, 6, 463–474.
- Chung, Y., Rabe-Hesketh, S., Dorie, V., Gelman, A., & Liu, J. (2013). A nondegenerate penalized likelihood estimator for variance parameters in multilevel models. *Psychometrika*, 78, 685–709.
- Cole, S. R., Chu, H., & Greenland, S. (2014). Maximum likelihood, profile likelihood, and penalized likelihood: A primer. *American Journal of Epidemiology*, 179, 252–260.
- Cowles, M. K., & Carlin, B. P. (1996). Markov chain Monte Carlo convergence diagnostics: A comparative review. *Journal of the American Statistical Association*, 91, 883–904.
- Cousineau, D., & Helie, S. (2013). Improving maximum likelihood estimation using prior probabilities: A tutorial on maximum a posteriori estimation and an examination of the Weibull distribution. *Tutorials in Quantitative Methods in Psychology*, 9, 61–71.
- DeCarlo, L. T., Kim, Y., & Johnson, M. S. (2011). A hierarchical rater model for constructed responses, with a signal detection rater model. *Journal of Educational Measurement*, 48, 333–356.
- de Valpine, P., Turek, D., Paciorek, C. J., Anderson-Bergman, C., Temple Lang, D., & Bodik, R. (2017). Programming with models: Writing statistical algorithms for general model structures with NIMBLE. *Journal of Computational and Graphical Statistics*, 26, 403–413.
- Depaoli, S., & Clifton, J. P. (2015). A Bayesian approach to multilevel structural equation modeling with continuous and dichotomous outcomes. *Structural Equation Modeling*, 22, 327–351.

- Depaoli, S., & van de Schoot, R. (2017). Improving transparency and replication in Bayesian statistics: The WAMBS-checklist. *Psychological Methods*, 22, 240–261.
- Dolan, C. V., & Molenaar, P. C. M. (1991). A comparison of four methods of calculating standard errors of maximum-likelihood estimates in the analysis of covariance structure. *British Journal of Mathematical and Statistical Psychology*, 44, 359–368.
- Draper, D. (2008). Bayesian multilevel analysis and MCMC. In J. deLeuw & E. Meijer (Eds.), *Handbook of multilevel analysis* (pp. 77–139). New York, NY: Springer.
- Efron, B. (1981). Nonparametric estimates of standard error: The jackknife, the bootstrap and other methods. *Biometrika*, 68, 589–599. doi:10.1093/biomet/68.3.589
- Fan, J., Li, R., Zhang, C.-H., & Zou, H. (2020). *Statistical foundations of data science*. Boca Raton: CRC Press.
- Fox, J.-P., Mulder, J., & Sinharay, S. (2017). Bayes factor covariance testing in item response models. *Psychometrika*, 82, 979–1006.
- Firth, D. (1993). Bias reduction of maximum likelihood estimates. *Biometrika*, 80, 27–38.
- Fuller, W. (1987). *Measurement error models*. New York, NY: John Wiley.
- Gagné, P., & Hancock, G. R. (2006). Measurement model quality, sample size, and solution propriety in confirmatory factor models. *Multivariate Behavioral Research*, 41, 65–83.
- Galindo-Garre, F., & Vermunt, J. (2006). Avoiding boundary estimates in latent class analysis by Bayesian posteriormode estimation. *Behaviormetrika*, 33, 43–59.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis*. London, UK: Chapman & Hall.
- Gerbing, D. W., & Anderson, J. C. (1987). Improper solutions in the analysis of covariance structures: Their interpretability and a comparison of alternate respecifications. *Psychometrika*, 52, 99–111.
- Grilli, L., & Rampichini, C. (2011). The role of sample cluster means in multilevel models. *Methodology*, 7, 121–133. doi: 10.1027/1614-2241/a000030

- Gill, J. (2007). *Bayesian methods for the social and behavioral sciences*. Boca Raton, FL: CRC Press.
- Hamra, G. B., MacLehose, R. F., & Cole, S. R. (2013). Sensitivity analyses for sparse-data problems—using weakly informative Bayesian priors. *Epidemiology*, 24, 233–239.
- Harwell, M. R., & Baker, F. B. (1991). The use of prior distributions in marginalized Bayesian item parameter estimation: A didactic. *Applied Psychological Measurement*, 15, 375–389. doi:10.1177/014662169101500409
- Hayashi, K., & Arav, M. (2006). Bayesian factor analysis when only a sample covariance matrix is available. *Educational and Psychological Measurement*, 66, 272–284.
- Hecht, M., Gische, C., Vogel, D., & Zitzmann, S. (2020). Integrating out nuisance parameters for computationally more efficient Bayesian estimation—an illustration and tutorial. *Structural Equation Modeling*, 27, 483–493.
- Heinze, G., & Schemper, M. (2002). A solution to the problem of separation in logistic regression. *Statistics in Medicine*, 21, 2409–2419.
- Heck, R. H., & Thomas, S. L. (2015). *An introduction to multilevel modeling techniques: MLM and SEM approaches using Mplus*. New York: Routledge.
- Held, L., & Bové, S. (2014). *Applied statistical inference*. New York: Springer.
- Hoff, P. D. (2009). *A first course in Bayesian statistical methods*. New York, NY: Springer.
- Hoogland, J. J., & Boomsma, A. (1998). Robustness studies in covariance structure modeling: An overview and a meta analysis. *Sociological Methods & Research*, 26, 329–368.
- Hoyle, R. H. (2012). *Handbook of structural equation modeling*. New York: Guilford Press.
- Hoyle, R. H., & Kenny, D. A. (1999). Sample size, reliability, and tests of statistical mediation. In R. H. Hoyle (Ed.), *Statistical strategies for small sample research* (pp. 195–222). Thousand Oaks, CA: Sage.
- Jackman, S. (2009). *Bayesian analysis for the social sciences*. New York, NY: Wiley.

- Jackson, D. L., Gillaspay, J. A., Jr., & Purc-Stephenson, R. (2009). Reporting practices in confirmatory factor analysis: An overview and some recommendations. *Psychological Methods, 14*, 6–23.
- Jacobucci, R., & Grimm, K. J. (2018). Comparison of frequentist and Bayesian regularization in structural equation modeling. *Structural Equation Modeling, 25*, 639–649.
- Junker, B. W., Patz, R. J., & VanHoudnos, N. M. (2016). Markov chain Monte Carlo for item response models. In W. J. van der Linden (Ed.), *Handbook of item response theory: Vol. 2. Statistical tools* (pp. 271–312). Boca Raton, FL: Chapman & Hall/CRC.
- Kaplan, D. & Depaoli, S. (2012). Bayesian structural equation modeling. In R. Hoyle (Ed.), *Handbook of structural equation modeling* (pp. 650–673). New York: Guilford Press.
- Kelava, A., & Brandt, H. (2014). A general non-linear multilevel structural equation mixture model. *Frontiers in Psychology, 5*, 748. doi:10.3389/fpsyg.2014.00748
- Kim, E. S., Dedrick, R. F., Cao, C., & Ferron, J. M. (2016). Multilevel factor analysis: Reporting guidelines and a review of reporting practices. *Multivariate Behavioral Research, 51*, 881–898.
- Kline, R. B. (2016). *Principles and practice of structural equation modeling*. New York: Guilford Press.
- Kenny, D. A., & Zautra, A. (1995). The trait-state-error model for multiwave data. *Journal of Consulting and Clinical Psychology, 63*, 52–59.
- Ledgerwood, A., & Shrout, P. E. (2011). The trade-off between accuracy and precision in latent variable models of mediation processes. *Journal of Personality and Social Psychology, 101*, 1174–1188.
- Lee, S.-Y. (2007). *Structural equation modeling: A Bayesian approach*. West Sussex, UK: Wiley.
- Lee, S.-Y., & Song, X.-Y. (2004). Evaluation of the Bayesian and maximum likelihood approaches in analyzing structural equation models with small sample sizes.

- Multivariate Behavioral Research*, 39, 653–686.
- Li, X., & Beretvas, S. N. (2013). Sample size limits for estimating upper level mediation models using multilevel SEM. *Structural Equation Modeling*, 20, 241–264.
- Little, T. D. (2013). *Longitudinal structural equation modeling*. New York: Guilford Press.
- Lüdtke, O., Marsh, H. W., Robitzsch, A., Trautwein, U., Asparouhov, T., & Muthén, B. (2008). The multilevel latent covariate model: A new, more reliable approach to group-level effects in contextual studies. *Psychological Methods*, 13, 203–229.
- Lüdtke, O., Marsh, H. W., Robitzsch, A., & Trautwein, U. (2011). A 2×2 taxonomy of multilevel latent contextual models: Accuracy–bias trade-offs in full and partial error correction models. *Psychological Methods*, 16, 444–467.
- Lüdtke, O., Robitzsch, A., Kenny, D. A., & Trautwein, U. (2013). A general and flexible approach to estimating the social relations model using Bayesian methods. *Psychological Methods*, 18, 101–119.
- Lüdtke, O., Robitzsch, A., & Wagner, J. (2018). More stable estimation of the STARTS model: A Bayesian approach using Markov chain Monte Carlo techniques. *Psychological Methods*, 23, 570–593. doi:10.1037/met0000155
- Lynch, S. M. (2007). *Introduction to applied Bayesian statistics and estimation for social scientists*. New York, NY: Springer.
- McNeish, D. (2016). On using Bayesian methods to address small sample problems. *Structural Equation Modeling*, 23, 750–773.
- Merkle, E. C., Fitzsimmons, E., Uanhoro, J., & Goodrich, B. (2020). Efficient Bayesian structural equation modeling in Stan. *arXiv Preprint*, arXiv:2008.07733v1.
- Merkle, E. C., Furr, D., & Rabe-Hesketh, S. (2019). Bayesian comparison of latent variable models: Conditional versus marginal likelihoods. *Psychometrika*, 84, 802–809.
- Miocevic, M., Levy, R., & MacKinnon, D. P. (2020). Different roles of prior distributions in the single mediator model with latent variables. *Multivariate Behavioral Research*.

Advance online publication. doi:10.1080/00273171.2019.1709405

- Mislevy, R. J. (1986). Bayes modal estimation in item response models. *Psychometrika*, 51, 177–195.
- Muthén, B. O., & Asparouhov, T. (2012). Bayesian structural equation modeling: A more flexible representation of substantive theory. *Psychological Methods*, 17, 313–335.
- 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.
- Muthén, L. K., & Muthén, B. O. (2012). *Mplus user's guide* (7th ed.). Los Angeles, CA: Muthén & Muthén.
- Plummer, M. (2003, March). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd international workshop on distributed statistical computing* (Vol. 124, pp. 1–10).
- Rindskopf, D. (2012). Next steps in Bayesian structural equation models: Comments on, variations of, and extensions to Muthén and Asparouhov (2012). *Psychological Methods*, 17, 336–339.
- Robitzsch, A. (2020). *LAM: Some latent variable models*. R package version 0.5-15, <https://CRAN.R-project.org/package=LAM> (accessed May 9, 2020).
- Rosseel, Y. (2012). lavaan: an R package for structural equation modeling. *Journal of Statistical Software*, 48, 1–36. Retrieved from <http://www.jstatsoft.org/v48/i02/>
- Rosseel, Y. (2020). Small sample solutions for structural equation modeling. In R. van de Schoot & M. Miočević (Eds.), *Small sample size solutions* (pp. 226–238). London: Routledge.
- Rousseeuw, P. J. & Molenberghs, G. (1994). The shape of correlation matrices. *The American Statistician*, 48, 276–279. doi:10.1080/00031305.1994.10476079
- Savalei, V. (2019). A comparison of several approaches for controlling measurement error in small samples. *Psychological Methods*, 24, 352–370.

- Savalei, V., & Kolenikov, S. (2008). Constrained vs. unconstrained estimation in structural equation modeling. *Psychological Methods, 13*, 150–170.
- Schoenberg, R. (1997). Constrained maximum likelihood. *Computational Economics, 10*, 251–266.
- Silverman, B. W. (1998). *Density estimation for statistics and data analysis*. Boca Raton, FL: CRC.
- Smid, S. C., McNeish, D., Miocevic, M., & van de Schoot, R. (2020). Bayesian versus frequentist estimation for structural equation models in small sample contexts: A systematic review. *Structural Equation Modeling, 27*, 131–161.
- Smid, S. C., & Rosseel, Y. (2020). SEM with small samples: Two-step modeling and factor score regression versus Bayesian estimation with informative priors. In R. van de Schoot & M. Miocevic (Eds.), *Small sample size solutions* (pp. 239–254). London: Routledge.
- Song, X.-Y., & Lee, S.-Y. (2009). *Basic and advanced Bayesian structural equation modeling: With applications in the medical and behavioral sciences*. Chichester, UK: Wiley.
- Spiegelhalter, D. J., Thomas, A., Best, N. G., Gilks, N. R., & Lunn, D. (2003). *BUGS: Bayesian inference using Gibbs sampling*. MRC Biostatistics Unit, Cambridge, UK.
- Traub, R. E. (1994). *Reliability for the social sciences: Theory and applications*. Thousand Oaks, CA: Sage Publications.
- van Erp, S., Mulder, J., & Oberski, D. L. (2018). Prior sensitivity analysis in default Bayesian structural equation modeling. *Psychological Methods, 23*, 363–388.
- Wolf, E. J., Harrington, K. M., Clark, S. L., & Miller, M. W. (2013). Sample size requirements for structural equation models: An evaluation of power, bias, and solution propriety. *Educational and Psychological Measurement, 73*, 913–934.

- Wothke, W. (1993). Nonpositive definite matrices in structural modeling. In K. A. Bollen & J. S. Long (Eds.), *Testing structural equation models* (pp. 256–293). Newbury Park, CA: Sage.
- Yuan, K.-H., & Bentler, P. M. (2007). Structural equation modeling. In C. R. Rao & S. Sinharay (Eds.), *Handbook of statistics 26: Psychometrics* (pp. 297–358). Amsterdam: North-Holland.
- Yuan, K.-H., & Chan, W. (2008). Structural equation modeling with near singular covariance matrices. *Computational Statistics and Data Analysis*, 52, 4842–4858.
- Zeng, L. (1997). Implementation of marginal Bayesian estimation with four-parameter beta prior distributions. *Applied Psychological Measurement*, 21, 143–156.
doi:10.1177/01466216970212004
- Zitzmann, S., & Hecht, M. (2019). Going beyond convergence in Bayesian estimation: Why precision matters too and how to assess it. *Structural Equation Modeling*, 26, 646–661.
- Zitzmann, S., Lüdtke, O., & Robitzsch, A. (2015). A Bayesian approach to more stable estimates of group-level effects in contextual studies. *Multivariate Behavioral Research*, 50, 688–705.
- Zitzmann, S., Lüdtke, O., Robitzsch, A., & Marsh, H. W. (2016). A Bayesian approach for estimating multilevel latent contextual models. *Structural Equation Modeling*, 23, 661–679.
- Zitzmann, S., Lüdtke, O., Robitzsch, A., & Hecht, M. (2020). On the performance of Bayesian approaches in small samples: A comment on Smid, McNeish, Miocevic, and van de Schoot (2020). *Structural Equation Modeling*. Advance online publication.
doi:10.1080/10705511.2020.1752216

Table 1

Relationship between Standardized Loading, Indicator-specific Reliability, and Reliability of Sum Score for Three Indicators

λ	Rel_I	Rel_T
.35	.13	.30
.43	.18	.40
.50	.25	.50
.58	.33	.60
.66	.44	.70
.76	.57	.80
.87	.75	.90

Note. λ = standardized loading. Rel_I = indicator-specific reliability. Rel_T = reliability of sum score .

Table 2

Illustrating Differences Between the Mode of the Joint Posterior (PML), and the Mode (MAP) and Mean (EAP) of Marginal Posterior: Bias, Standard Deviation, and RMSE as a Function of the Correlation (ρ) and Sample Size (N)

ρ	N	Bias			SD			RMSE		
		PML	MAP	EAP	PML	MAP	EAP	PML	MAP	EAP
0.1	30	0.01	0.01	-0.02	0.41	0.41	0.29	0.41	0.41	0.29
	50	0.01	0.01	0.00	0.31	0.31	0.26	0.31	0.31	0.26
	100	-0.01	-0.01	-0.01	0.20	0.21	0.20	0.20	0.21	0.20
	1000	0.00	0.00	0.00	0.06	0.06	0.06	0.06	0.06	0.06
0.3	30	0.00	0.00	-0.07	0.38	0.37	0.26	0.38	0.37	0.27
	50	-0.01	0.00	-0.04	0.31	0.31	0.25	0.31	0.31	0.26
	100	0.00	0.00	-0.01	0.20	0.20	0.19	0.20	0.20	0.19
	1000	0.00	0.00	0.00	0.06	0.06	0.06	0.06	0.06	0.06
0.5	30	-0.01	-0.01	-0.14	0.37	0.36	0.26	0.37	0.36	0.29
	50	0.00	0.01	-0.06	0.27	0.27	0.22	0.27	0.27	0.23
	100	0.01	0.01	-0.01	0.19	0.19	0.17	0.19	0.19	0.17
	1000	0.00	0.00	0.00	0.06	0.06	0.06	0.06	0.06	0.06
0.7	30	-0.03	-0.03	-0.20	0.30	0.29	0.22	0.30	0.29	0.30
	50	0.00	0.00	-0.11	0.23	0.23	0.18	0.23	0.23	0.21
	100	0.01	0.01	-0.04	0.17	0.17	0.14	0.17	0.17	0.14
	1000	0.00	0.00	0.00	0.05	0.05	0.05	0.05	0.05	0.05
0.9	30	-0.07	-0.07	-0.27	0.22	0.21	0.18	0.23	0.22	0.33
	50	-0.04	-0.04	-0.18	0.17	0.16	0.13	0.17	0.17	0.23
	100	-0.02	-0.02	-0.10	0.12	0.12	0.09	0.12	0.12	0.14
	1000	0.00	0.00	-0.01	0.05	0.05	0.04	0.05	0.05	0.04

Note. SD = standard deviation of empirical sampling distribution; RMSE = root mean square error; ρ = true latent correlation; N = sample size; PML = mode of joint posterior (obtained from maximizing the joint posterior); MAP = mode of marginal posterior (obtained from maximizing the marginal posterior); EAP = mean of marginal posterior (obtained from numerical intergration). Biases smaller than -0.05 or larger than 0.05 are written in bold.

Table 3

Percentage of Solutions with Estimation Problems for Unconstrained Maximum Likelihood (ML) and Constrained Maximum Likelihood (CML) by Magnitude of the True Correlation (ρ), and Sample Size (N)

N	ML Conv			ML Conv+Adm			CML Boundary			ML = CML		
	ρ			ρ			ρ			ρ		
	.3	.5	.7	.3	.5	.7	.3	.5	.7	.3	.5	.7
30	54.5	64.4	67.9	51.3	57.3	53.4	6.3	11.1	21.6	28.0	33.0	34.8
50	73.1	82.7	88.0	70.9	77.5	73.9	3.4	6.3	17.0	53.1	61.7	64.3
100	94.0	97.6	99.6	93.9	96.4	92.8	0.4	1.4	7.2	87.1	92.8	91.6

Note. ML Conv = converged solutions for unconstrained ML; ML Conv+Adm = converged solution with admissible values for unconstrained ML, that is estimated correlation was within the interval $[-1, 1]$; CML Boundary = boundary estimate for correlation with constrained ML; ML=CML = proportion of solutions in which estimated correlation for ML and CML were numerically identical. Note that ML Conv+Adm is not a conditional proportion, and, thus, values cannot exceed ML Conv.

Table 4

Bias and RMSE for Unconstrained Maximum Likelihood (ML) and Constrained Maximum Likelihood (CML) as a Function of the True Correlation (ρ), and the Sample Size (N)

<i>Meth</i>	<i>N</i>	Bias			RMSE		
		ρ			ρ		
		.3	.5	.7	.3	.5	.7
ML	30	0.04	0.02	-0.03	0.40	0.33	0.30
	50	0.04	-0.01	-0.01	0.31	0.30	0.25
	100	0.00	0.01	0.00	0.20	0.20	0.18
CML	30	-0.02	-0.03	-0.07	0.41	0.36	0.34
	50	0.00	-0.05	-0.03	0.33	0.32	0.27
	100	0.00	0.00	0.00	0.21	0.21	0.18

Note. RMSE = root mean square error; ρ = true latent correlation; N = sample size; *Meth* = estimation method (unconstrained ML and constrained ML estimation). Only converged solutions are included for unconstrained ML; values outside $[-1, 1]$ are truncated to -1 or 1 . Biases smaller than $-.05$ or larger than $.05$ are written in bold.

Table 5

Bias and RMSE for Penalized Maximum Likelihood and Markov Chain Monte Carlo with Different Correctly Specified Prior Distributions as a Function of the Sample Size and ρ

<i>N</i>	<i>Meth</i>	v_p	v_λ	Bias			SD			RMSE Gain		
				ρ			ρ			ρ		
				.3	.5	.7	.3	.5	.7	.3	.5	.7
30	PML	0	0	0.00	-0.05	-0.07	0.41	0.39	0.33	100	100	100
		0	1	0.03	-0.01	-0.02	0.43	0.40	0.31	105	101	93
		0	3	0.04	0.00	0.00	0.42	0.39	0.30	103	98	90
		1	0	0.09	0.15	0.16	0.37	0.36	0.26	92	100	91
		1	1	0.12	0.20	0.19	0.37	0.35	0.23	95	101	89
		1	3	0.12	0.21	0.21	0.36	0.35	0.22	92	102	90
		3	0	0.05	0.09	0.17	0.28	0.26	0.23	69	69	85
		3	1	0.09	0.12	0.19	0.29	0.25	0.21	73	69	83
		3	3	0.09	0.12	0.21	0.28	0.24	0.20	71	69	85
	MCMC	0	0	-0.07	-0.15	-0.21	0.28	0.26	0.24	71	76	95
		0	1	-0.02	-0.08	-0.12	0.33	0.30	0.24	80	78	81
		0	3	-0.01	-0.06	-0.10	0.34	0.31	0.25	83	79	79
		1	0	0.06	0.13	0.10	0.22	0.17	0.15	55	54	54
		1	1	0.08	0.15	0.13	0.27	0.21	0.14	67	65	57
		1	3	0.09	0.16	0.14	0.28	0.22	0.14	70	69	58
		3	0	-0.01	0.00	0.09	0.17	0.15	0.11	41	38	41
		3	1	0.02	0.04	0.11	0.19	0.18	0.12	46	46	48
		3	3	0.02	0.05	0.12	0.20	0.18	0.12	48	48	50
50	PML	0	0	-0.02	-0.01	-0.04	0.33	0.30	0.28	100	100	100
		0	1	0.00	0.02	-0.01	0.34	0.29	0.25	102	97	92
		0	3	0.01	0.03	0.00	0.33	0.29	0.25	99	95	89
		1	0	0.04	0.14	0.15	0.31	0.32	0.25	94	116	104
		1	1	0.06	0.17	0.17	0.31	0.31	0.23	94	116	102
		1	3	0.06	0.17	0.19	0.30	0.31	0.22	91	116	103
		3	0	0.02	0.07	0.15	0.25	0.23	0.23	75	81	98
		3	1	0.05	0.09	0.16	0.26	0.22	0.21	78	79	96
		3	3	0.04	0.10	0.18	0.25	0.22	0.20	75	79	98
	MCMC	0	0	-0.07	-0.09	-0.15	0.25	0.23	0.21	77	82	93
		0	1	-0.03	-0.04	-0.09	0.28	0.24	0.21	83	82	82
		0	3	-0.03	-0.02	-0.06	0.28	0.25	0.21	85	83	81
		1	0	0.01	0.10	0.10	0.21	0.19	0.15	64	72	63
		1	1	0.03	0.12	0.12	0.24	0.22	0.14	73	82	66
		1	3	0.03	0.13	0.13	0.25	0.23	0.14	75	86	68
		3	0	-0.02	0.00	0.07	0.17	0.16	0.13	51	52	54
		3	1	0.00	0.03	0.10	0.19	0.17	0.14	57	59	60
		3	3	0.00	0.04	0.10	0.20	0.18	0.14	58	61	62

100	PML	0	0	-0.01	0.00	0.01	0.22	0.21	0.18	100	100	100
		0	1	0.01	0.01	0.01	0.22	0.20	0.18	99	96	97
		0	3	0.01	0.01	0.02	0.22	0.20	0.17	97	95	95
		1	0	0.02	0.10	0.16	0.22	0.27	0.20	98	134	140
		1	1	0.03	0.10	0.17	0.21	0.25	0.19	96	129	139
		1	3	0.04	0.10	0.17	0.21	0.25	0.19	95	126	139
		3	0	0.02	0.04	0.16	0.19	0.18	0.19	87	89	134
		3	1	0.03	0.05	0.16	0.19	0.18	0.18	86	88	134
		3	3	0.03	0.05	0.16	0.19	0.18	0.18	85	87	132
	MCMC	0	0	-0.04	-0.05	-0.06	0.19	0.19	0.17	88	92	96
		0	1	-0.02	-0.03	-0.03	0.20	0.19	0.16	90	91	90
		0	3	-0.01	-0.02	-0.02	0.20	0.19	0.16	90	90	88
		1	0	-0.01	0.05	0.10	0.18	0.20	0.14	82	96	94
		1	1	0.01	0.06	0.11	0.19	0.20	0.13	86	100	96
		1	3	0.01	0.06	0.12	0.19	0.20	0.13	87	101	97
		3	0	-0.02	-0.01	0.07	0.16	0.16	0.14	71	75	86
		3	1	0.00	0.01	0.08	0.16	0.16	0.14	74	75	88
		3	3	0.00	0.02	0.09	0.17	0.16	0.13	75	76	88

Note. RMSE = root mean square error. PML = penalized maximum likelihood; MCMC = Markov

Chain Monte Carlo; v_p = prior sample size for latent correlation; v_λ = prior sample size for standardized loading; biases smaller than $-.05$ or larger than $.05$ are written in bold. For the RMSE gain, the RMSE of PML with uniform prior distributions is used as a reference method; values smaller than 100 indicated that the RMSE for the respective is lower than the RMSE of the reference method.

Table 6

Coverage Rates for Penalized Maximum Likelihood and Markov Chain Monte Carlo with Different Correctly Specified Prior Distributions as a Function of the Magnitude of the True Correlation, and a Sample Size of $N = 100$

v_ρ	v_λ	PML			MCMC		
		ρ			ρ		
		.3	.5	.7	.3	.5	.7
0	0	89.9	92.7	97.0	95.6	95.7	96.8
0	1	92.1	94.8	97.0	95.3	96.1	97.2
0	3	92.8	95.4	97.5	94.8	96.4	96.9
1	0	91.0	92.3	94.4	96.8	96.7	89.9
1	1	92.3	92.5	94.6	96.1	94.9	88.5
1	3	93.0	91.0	95.1	95.5	94.5	87.8
3	0	93.8	95.8	95.1	97.3	98.5	97.0
3	1	95.5	96.3	95.1	97.1	97.8	96.4
3	3	95.8	96.1	95.8	96.6	98.0	95.1

Note. PML = penalized maximum likelihood; MCMC = Markov Chain Monte Carlo; ρ = true latent correlation; v_ρ = prior sample size for latent correlation; v_λ = prior sample size for standardized loading. Coverage rates smaller than 91% or larger than 98% are written in bold.

Table 7

Bias and RMSE for Penalized Maximum Likelihood and Markov Chain Monte Carlo with Different Misspecified Prior Distributions as a Function of the Sample Size

			Bias					RMSE				
			$\mu_\lambda = 0.5$			$\mu_\lambda = 0.8$		$\mu_\lambda = 0.5$			$\mu_\lambda = 0.8$	
Meth	μ_p	v_p	$v_\lambda=0$	$v_\lambda=1$	$v_\lambda=3$	$v_\lambda=1$	$v_\lambda=3$	$v_\lambda=0$	$v_\lambda=1$	$v_\lambda=3$	$v_\lambda=1$	$v_\lambda=3$
$N = 30$												
PML	0	0	-0.04	0.01	0.02	-0.14	-0.02	0.39	0.39	0.38	0.36	0.37
	0.5	1	0.16	0.20	0.20	-0.01	0.12	0.39	0.40	0.40	0.32	0.34
	0.5	3	0.09	0.12	0.12	-0.02	0.09	0.28	0.28	0.28	0.26	0.26
	0.2	1	-0.03	0.02	0.03	-0.11	0.00	0.32	0.32	0.32	0.32	0.32
	0.2	3	-0.09	-0.04	-0.04	-0.14	-0.05	0.27	0.26	0.26	0.29	0.27
	0.8	1	0.30	0.33	0.35	0.11	0.24	0.43	0.44	0.45	0.36	0.38
	0.8	3	0.34	0.36	0.39	0.15	0.27	0.43	0.44	0.45	0.33	0.37
MCMC	0	0	-0.14	-0.07	-0.05	-0.18	-0.10	0.30	0.31	0.31	0.31	0.31
	0.5	1	0.13	0.15	0.16	-0.04	0.06	0.22	0.26	0.28	0.25	0.27
	0.5	3	0.00	0.04	0.05	-0.07	0.01	0.16	0.19	0.20	0.20	0.20
	0.2	1	-0.12	-0.07	-0.06	-0.16	-0.09	0.24	0.25	0.25	0.28	0.27
	0.2	3	-0.18	-0.14	-0.13	-0.19	-0.13	0.25	0.23	0.23	0.27	0.24
	0.8	1	0.36	0.34	0.33	0.15	0.22	0.39	0.38	0.38	0.28	0.32
	0.8	3	0.31	0.32	0.33	0.15	0.23	0.33	0.35	0.35	0.25	0.30
$N = 100$												
PML	0	0	-0.01	0.00	0.01	-0.04	-0.02	0.20	0.20	0.19	0.19	0.18
	0.5	1	0.09	0.09	0.09	0.03	0.03	0.27	0.26	0.26	0.23	0.21
	0.5	3	0.03	0.04	0.05	0.00	0.01	0.18	0.18	0.18	0.17	0.16
	0.2	1	-0.01	0.00	0.01	-0.03	-0.02	0.19	0.18	0.18	0.19	0.17
	0.2	3	-0.04	-0.02	-0.02	-0.05	-0.04	0.17	0.17	0.16	0.18	0.16
	0.8	1	0.19	0.18	0.16	0.09	0.08	0.33	0.32	0.30	0.26	0.23
	0.8	3	0.22	0.22	0.20	0.14	0.12	0.34	0.33	0.32	0.29	0.25
MCMC	0	0	-0.06	-0.03	-0.02	-0.11	-0.07	0.19	0.19	0.19	0.21	0.18
	0.5	1	0.04	0.05	0.06	-0.05	-0.02	0.20	0.21	0.21	0.19	0.18
	0.5	3	-0.02	0.00	0.01	-0.07	-0.03	0.15	0.16	0.16	0.17	0.15
	0.2	1	-0.06	-0.04	-0.03	-0.11	-0.07	0.18	0.17	0.17	0.20	0.17
	0.2	3	-0.09	-0.07	-0.06	-0.13	-0.09	0.17	0.16	0.16	0.19	0.17
	0.8	1	0.19	0.18	0.18	0.04	0.06	0.28	0.27	0.27	0.21	0.20
	0.8	3	0.18	0.17	0.18	0.04	0.06	0.25	0.25	0.25	0.19	0.19

Note. RMSE = root mean square error. PML = penalized maximum likelihood; MCMC = Markov

Chain Monte Carlo; μ_p = prior guess for latent correlation; v_p = prior sample size for latent correlation;

μ_λ = prior guess for loading v_λ = prior sample size for loading; biases smaller than -0.05 or larger than

0.05 are written in bold. For each sample size condition, RMSE values larger than the RMSE for PML

with uniform priors on loadings ($\mu_\lambda = 0$, $v_\lambda = 0$) and the correlation ($\mu_p = 0$, $v_p = 0$) are printed in bold.

Figure 1. Confirmatory factor analysis (CFA) model with two latent factors. The variances of the latent factors are set to 1.0.

Figure 2. Four-parameter beta distributions for a standardized loading λ (left panel), and the correlation ρ (right panel). With a larger prior sample size (v_λ or v_ρ), the distribution puts more probability mass around the prior guess (μ_λ or μ_ρ).

Figure 3. Illustrating the difference between the joint and marginal posterior distribution: the square in the first row indicates the (multivariate) mode of the joint posterior distribution for $N = 30$ (upper left panel), and $N = 100$ (upper right panel); the triangle and the circle in the second row indicate the (univariate) mode and mean of the marginal posterior distribution for the correlation ρ for $N = 30$ (lower left panel), and $N = 100$ (lower right panel). Note that with $N = 30$, the mode of the joint posterior (PML) for ρ strongly deviates from the mean of the marginal posterior (EAP).

Figure 4. Illustrating the convergence of different Bayesian estimates of the correlation (true correlation $\rho = .7$) as a function of the sample size N : mode of the joint posterior distribution (PML), mode of the marginal posterior distribution (MAP), and mean of the marginal posterior distribution (EAP). The covariance matrix of the observed variables was held fixed, and the sample size was varied.

Figure 5. Bias (left panel) and RMSE (right panel) for constrained maximum likelihood (CML), unconstrained maximum likelihood with trimmed estimates (ML trimmed), and unconstrained maximum likelihood with truncated estimates (ML truncated). The x -axis shows the cut-off value (for the estimated correlation) for truncating or trimming (i.e., eliminating) a replication. The results show that choosing a more liberal cut-off value negatively influences the performance of unconstrained ML. Displayed are results for the condition with $N = 30$, and $\rho = .7$.