Prior Sensitivity Analysis in Default Bayesian Structural Equation Modeling

Sara van Erp¹, Joris Mulder¹, and Daniel L. Oberski²

¹Tilburg University

²Utrecht University

©2017, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors permission. The final article will be available, upon publication, via its DOI: 10.1037/met0000162

Abstract

Bayesian structural equation modeling (BSEM) has recently gained popularity because it enables researchers to fit complex models while solving some of the issues often encountered in classical maximum likelihood (ML) estimation, such as nonconvergence and inadmissible solutions. An important component of any Bayesian analysis is the prior distribution of the unknown model parameters. Often, researchers rely on default priors, which are constructed in an automatic fashion without requiring substantive prior information. However, the prior can have a serious influence on the estimation of the model parameters, which affects the mean squared error (MSE), bias, coverage rates, and quantiles of the estimates.

In this paper, we investigate the performance of three different default priors: noninformative improper priors, vague proper priors, and empirical Bayes priors, with the latter being novel in the BSEM literature. Based on a simulation study, we find that these three default BSEM methods may perform very differently, especially with small samples. A careful prior sensitivity analysis is therefore needed when performing a default BSEM analysis. For this purpose, we provide a practical step-by-step guide for practitioners to conducting a prior sensitivity analysis in default BSEM. Our recommendations are illustrated using a well-known case study from the structural equation modeling literature and all code for conducting the prior sensitivity analysis is made available in the online supplemental material.

Keywords: *Bayesian, structural equation models, default priors, sensitivity analysis.*

Prior Sensitivity Analysis in Default Bayesian Structural Equation Modeling

Author notes

Sara van Erp, Joris Mulder, Department of Methodology and Statistics, Tilburg University, The Netherlands.

Daniel L. Oberski, Methodology and Statistics, Utrecht University, The Netherlands.

This research was supported by an NWO Research Talent Grant.

Correspondence concerning this article should be addressed to Sara van Erp, Department of Methodology and Statistics, Tilburg University, P.O. box 90153, 5000 LE Tilburg, The Netherlands. E-mail: s.vanerp_1@tilburguniversity.edu

This work has been presented at the Modern Modeling Methods conference in 2016 and the IOPS summer conference in 2016. A preprint of this paper has been made available on the Open Science Framework.

Introduction

Psychologists and social scientists often ask complex questions regarding group- and individual differences and how these change over time. These complex questions necessitate complex methods such as structural equation modeling (SEM); its Bayesian version (Bayesian structural equation modeling; BSEM), in particular, has recently gained popularity (e.g., Kaplan, 2014) because it potentially resolves some of the difficulties with traditional frequentist SEM. For example, frequentist estimation of multilevel SEMs—often employed when studying multiple classrooms, schools, or countries—has been found to perform badly in terms of bias and power with a small number of groups (Lüdtke, Marsh, Robitzsch, & Trautwein, 2011; Maas & Hox, 2005; Meuleman & Billiet, 2009; Ryu & West, 2009), while BSEM performed well even with small samples (Depaoli & Clifton, 2015; Hox, van de Schoot, & Matthijsse, 2012). BSEM may also reduce issues with nonconvergence (Kohli, Hughes, Wang, Zopluoglu, & Davison, 2015) and inadmissible estimates (Can, van de Schoot, & Hox, 2014; Dagne, Howe, Brown, & Muthén, 2002), it is computationally convenient for models with many latent variables (Harring, Weiss, & Hsu, 2012; Lüdtke, Robitzsch, Kenny, & Trautwein, 2013; Oravecz, Tuerlinckx, & Vandekerckhove, 2011), and BSEM easily yields credible intervals (i.e., the Bayesian version of a confidence interval) on functions of parameters such as reliabilities (Geldhof, Preacher, & Zyphur, 2014) or indirect effects (Yuan & MacKinnon, 2009). Furthermore, BSEM allows researchers to assume that traditionally restricted parameters, such as cross-loadings, direct effects, and error covariances, are approximately rather than exactly zero by incorporating prior information (MacCallum, Edwards, & Cai, 2012; B. O. Muthén & Asparouhov, 2012). However, to take advantage of BSEM, one challenge must be overcome (MacCallum et al., 2012): the specification of the prior distributions. Prior specification is an important but difficult part of any Bayesian analysis. Ideally, the priors should accurately reflect preexisting knowledge about the world, both in terms of the facts and the uncertainty about those facts. Previous research has shown that BSEM has superior performance to frequentist SEM from a subjective Bayesian perspective when priors reflect researchers' beliefs exactly; and from a frequentist perspective,

BSEM outperforms frequentist SEM when priors reflect reality. Priors that do not reflect prior beliefs to infinite accuracy (Bayesian perspective), or that do not correspond to reality (frequentist perspective), however, can lead to severe bias (Baldwin & Fellingham, 2013; Depaoli, 2012, 2013, 2014; Depaoli & Clifton, 2015). Moreover, eliciting priors is a time-consuming task, and even experts are often mistaken and prone to overstating their certainty (e.g., Garthwaite, Kadane, & O'Hagan, 2005; Tversky, 1974). Additionally, in BSEM it is generally more difficult to specify subjective priors due to the many parameters, some of which are not easily interpretable (e.g., latent variable variances). Therefore, instead of relying fully on expert judgements, researchers employing Bayesian analysis often use "default" priors. Default priors can be viewed as a class of priors that (i) do not contain any external substantive information, (ii) are completely dominated by the information in the data, and (iii) can be used in an automatic fashion for a Bayesian data analysis (Berger, 2006). For this reason default priors seem particularly useful for SEM because they allow us to use the flexible Bayesian approach without needing to translate prior knowledge into informative priors.

Previous research has investigated the performance of several default priors for BSEM. Thus far, the BSEM priors studied have been limited to proper priors chosen to equal the true population values in expectation, or chosen purposefully to be biased in expectation by a certain percentage (Depaoli, 2012, 2013, 2014; Depaoli & Clifton, 2015). These studies yielded important insights into the consequences of prior choice. However, commonly suggested alternative default priors in the Bayesian literature, such as noninformative improper priors and empirical Bayes priors (e.g., Carlin & Louis, 2000a; Casella, 1985; Natarajan & Kass, 2000), remain, to our knowledge, uninvestigated. Moreover, while several authors agree that any BSEM analysis should be accompanied by a sensitivity analysis, the available practical guidelines to do so focus on the situation in which substantive information was used to specify the prior (Depaoli & van de Schoot, 2015). In addition, Depaoli and Van de Schoot (2015) provide specific guidelines on checking the sensitivity of the results to different inverse Wishart priors for the covariance matrix. Our contribution is that we specifically focus on prior specification in BSEM, we focus on default

prior specification when prior information is weak or completely unavailable, and we specifically focus on univariate priors (i.e., univariate normal, and inverse gamma) which are easiest to interpret by applied researchers.

In addition to the lack of knowledge regarding priors in BSEM, there appears to be a lack of awareness of the importance of the prior as well. A recent review by Van de Schoot, Winter, Ryan, Zondervan-Zwijnenburg, and Depaoli (2016) identified trends in and uses of Bayesian methods based on 1579 papers published in psychology between 1990 and 2015. Of the 167 empirical papers employing regression techniques (including SEM), only 45% provided information about the prior that was used. 31% of the papers did not discuss which priors were used at all, while 24% did not provide enough information to reconstruct the priors. In terms of the type of prior, 26.7% of the empirical papers used informative priors, of which only 4.5% (2 papers) employed empirical Bayes methods to choose the hyperparameters.

The literature review of Van de Schoot et al. (2016) showed that a substantial part of Bayesian analyses in psychology relies on default priors. Now the problem is that the exact choice of the default prior may affect the conclusions substantially, as has been shown in the general Bayesian literature (e.g., Gelman, 2006; Lambert, Sutton, Burton, Abrams, & Jones, 2005) and will be shown in the context of BSEM in this paper. Different software packages have implemented different default or weakly informative priors as their default software settings (Table 1). With the development of more user-friendly Bayesian software, more non-expert users are trying out Bayesian analysis in general and BSEM in particular and rely on the default software settings, without being fully aware of the influence and importance of the prior distributions. In the 167 empirical papers identified by van de Schoot et al. (2016), WinBUGS (Lunn, Thomas, Best, & Spiegelhalter, 2000) was the most popular software program until 2012, but since 2013 this position has been taken over by the commercial SEM software Mplus (which has Bayesian methods implemented since 2010; L.K. Muthén & Muthén, 1998-2012; van de Schoot et al., 2016).

[TABLE1]

This paper aims to further develop the practice and utility of default BSEM and to raise awareness that the exact choice of default prior in BSEM matters. Specifically, this paper has the following three goals. First, we propose two novel empirical Bayes (EB) prior settings which adapt to the observed data and are easy to implement. Second, we investigate the performance of several default priors, including the novel EB priors, and compare them with the priors studied thus far, thereby investigating prior sensitivity in default BSEM. Third, since the choice of the default prior can have a large effect on the estimates in small samples, we provide a step-by-step guide on how to perform a default prior sensitivity analysis. Note that we focus on frequentist properties of the different default priors, such as bias, mean squared errors, and coverage rates. We take this perspective because it is common to focus on frequentist properties when assessing the performance of default priors (Bayarri & Berger, 2004). A different and popular perspective on Bayesian statistics is that of updating one's prior beliefs with the data. In this perspective, instead of default priors, informative priors are used which contain external subjective information about the magnitude of the parameters before observing the data. Specification of such subjective priors will not be explored in the current paper.

The rest of this article is organized as follows. We first introduce the BSEM model using a running example from the SEM literature. In the subsequent section, we discuss possible priors that have been suggested both in the BSEM and in the wider Bayesian analysis literature. Subsequently, a simulation study investigates the effect these prior choices have on BSEM estimates. We then provide practical guidelines based on the results of the simulation for practitioners who wish to perform their own sensitivity analysis. Finally, we apply these guidelines to empirical data from the running example, providing a demonstration of sensitivity analysis in BSEM.

A Structural Equation Model

Throughout this paper we will consider a linear structural equation model with latent variables from the literature. We have selected this model because it is one of the most popular example

models in structural equation modeling. Furthermore the model includes a mediation effect, which is of interest in substantive research. As a result, investigation of this model will not only result in general insights regarding default priors in BSEM but will also provide specific information about default priors for mediation analysis. The model (Figure 1) describes the influence of the level of industrialization in 1960 (ξ) on the level of political democracy in 1960 (η^{60}) and 1965 (η^{65}) in 75 countries. Industrialization is measured by three indicators and the level of democracy by four indicators at each time point. The indicators for level of democracy consist of expert ratings, and, since some of the ratings come from the same expert at both time points or the same source in the same year, several measurement errors correlate, which we model through pseudo-latent variables D , following Dunson, Palomo, and Bollen (2005).

[FIGURE 1]

The structural model (for $i = 1, \dots, n$) is given by:

$$\begin{aligned} \eta_i &= \alpha + B\eta_i + \Gamma\xi_i + \zeta_i \quad \text{with } \xi_i \sim N(\mu_\xi, \omega_\xi^2), \\ &\quad \text{and } \zeta_i \sim N(0, \Omega_\zeta) \end{aligned}$$

The measurement model is given by:

$$\begin{aligned} y_i &= \nu_y + \Lambda_y\eta_i + D_i + \epsilon_i^y \quad \text{with } D_i \sim N(0, \Omega_D), \\ &\quad \text{and } \epsilon_i^y \sim N(0, \Sigma_y) \\ x_i &= \nu_x + \Lambda_x\xi_i + \delta_i^x \quad \text{with } \delta_i^x \sim N(0, \Sigma_x) \end{aligned}$$

Here, the structural mean and intercepts μ_ξ and α reflect the mean structure in the structural part of the model, while the measurement intercepts ν_y and ν_x reflect the mean structure in the measurement part of the model. The loadings Λ_y and Λ_x represent the relations between the latent variables and their indicators, and the structural regression coefficients B and Γ represent the relations between the latent variables. The residual variances Σ_y and Σ_x reflect the variation in the measurement errors, and the random variances ω_ξ^2 , Ω_ζ , and Ω_D reflect the variation in the latent variables. In practice, researchers are often most interested in the relations between the

latent variables, in this case the direct effect γ_{65} and the indirect effect $\gamma_{60} \cdot b_{21}$ of industrialization in 1960 on political democracy in 1965. Note that this specification is more restrictive than necessary since, following Dunson, Palomo, and Bollen (2005), we fix the nonzero entries in the weight matrix for $\mathbf{D}$ to be 1. As a result negative covariances are not allowed. A solution to this problem is implemented in the R-package *blavaan* (Merkle & Rosseel, 2015; R Core Team, 2015). The model is identified by restricting the first intercept and loading of each latent variable, i.e., $\lambda_y^1 = \lambda_y^5 = \lambda_x^1 = 1$ and $\nu_y^1 = \nu_y^5 = \nu_x^1 = 0$. Appendix A provides the full model in matrix form and a more detailed description of the data and the model can be found in Bollen (1980, 1989). In the application we will use the original data containing observations from 75 countries (available in the *lavaan* package in R; Rosseel, 2012), and a subset of the data containing only the first 35 observations. Maximum likelihood (ML) estimation for this subset gives two warnings: 1) the standard errors of the parameter estimates may not be reliable due to a non-positive definite first-order derivative product matrix, and 2) the latent variable covariance matrix is not positive definite. The first warning can be an indication of weak empirical identification due to the small sample size, whereas the second warning indicates an inadmissible parameter estimate; in this case the estimated variance of the pseudo-latent variable representing the relation between ϵ_4^y and ϵ_8^y , i.e., $\hat{\omega}_{D48}^2$, is negative. These warnings clearly illustrate that when using classical ML estimation, researchers may encounter certain problems which may be overcome by adopting a Bayesian approach, where prior distributions are specified in the subspace of admissible solutions.

Default priors for Bayesian SEM

Default priors have the following three key properties. First, they do not contain external subjective information. Second, they are completely dominated by the data. Third, they can be used in an automatic fashion for a Bayesian data analysis. Because of the second property these priors are also referred to as “noninformative” or “weakly informative” priors. Such priors are often used in Bayesian analysis, including BSEM, when no substantive information is available or when the researcher does not wish to incorporate any substantive information through the prior

distribution. Default priors allow researchers to use the powerful and flexible Bayesian approach without needing to specify an informative prior based on one's prior knowledge, which can be a difficult and time-consuming task. Different software packages use different default priors in their automatic settings based on various heuristic arguments (see Table 1). For example, the commercial SEM software Mplus (L. K. Muthén & Muthén, 1998-2012) specifies a uniform improper prior for variance parameters by default, while the Bayesian modeling software WinBUGS (Lunn et al., 2000) recommends vague proper inverse Gamma priors for the variances. Default priors can roughly be divided in the following three categories: noninformative¹ improper priors, vague proper priors, and empirical Bayes priors. The first two have been used extensively in the BSEM literature, while the latter has, to our knowledge, not been applied to BSEM yet, but it is popular in the general literature on Bayesian modeling (Carlin & Louis, 2000a; Casella, 1985; Natarajan & Kass, 2000). In this study, we will focus on different priors from each of these three commonly used types of default priors. Some priors are chosen because they are the default setting in popular software programs (specifically Mplus and blavaan), while other priors are chosen because they are widely used, or because they have been shown to perform well in certain situations. The EB priors are novel in BSEM and have been included to investigate whether an EB approach can be advantageous in BSEM.

For all three types, we focus on priors that have a conditionally conjugate form. Conditionally conjugate priors have the advantage that they result in fast computation because the resulting conditional posteriors have known distributional forms (i.e., they have the same distribution as the prior) from which we can easily sample. Specifically, the conditionally conjugate prior for a location parameter (e.g., intercepts, loadings, and regression coefficients) is the normal distribution, and for a variance parameter it is the inverse Gamma distribution. Note that these are univariate priors. Generally, it is possible to instead specify a multivariate normal distribution for the vector of location parameters and an inverse Wishart prior for the covariance matrix. However, in this paper we focus specifically on univariate priors for separate parameters. We will now discuss these different default priors in more detail.

Noninformative improper priors

Noninformative improper priors are most commonly used in “objective Bayesian analysis” (Berger, 2006). In a simple normal distribution with unknown mean μ and unknown variance σ^2 , for example, the standard noninformative improper prior $p(\mu, \sigma^2) \propto \sigma^{-2}$ (known as Jeffreys’ prior) yields exactly the same point and interval estimates for the population mean as does classical ML estimation; hence the name “objective Bayes”. An improper prior is not a formal probability distribution because it does not integrate to unity. A potential problem of noninformative improper priors is that the resulting posteriors may also be improper, which occurs when there is too little information in the data (Hobert & Casella, 1996). In the above example of a normal distribution with unknown mean and variance we need at least two distinct observations in order to obtain a proper posterior for μ and σ^2 when starting with the improper Jeffreys’ prior. Currently, little is known about the performance of these types of priors in BSEM. Throughout this paper we will therefore consider the following noninformative improper priors for variance parameters σ^2 :

- $p(\sigma^2) \propto \sigma^{-2}$. This prior is most commonly used in objective Bayesian analysis for variance components. It is equivalent to a uniform prior on $\log(\sigma^2)$. There have been reports, however, that this prior results in improper posteriors for variances of random effects in multilevel analysis (e.g., Gelman, 2006). In a simple normal model with known mean and unknown variance, at least one observation is needed for this prior to result in a proper posterior for the variance.
- $p(\sigma^2) \propto \sigma^{-1}$. This prior was recommended by Berger (2006) and Berger and Strawderman (1996) for variance components in multilevel models. For this prior, at least two observations are needed in a normal model with known mean and unknown variance to obtain a proper posterior.
- $p(\sigma^2) \propto 1$. This prior is the default choice in Mplus (L. K. Muthén & Muthén, 1998-2012). Gelman (2006) noted that it may result in overestimation of the variance. When using this

prior in a normal model with known mean and unknown variance, at least three observations are needed to obtain a proper posterior for the variance.

Each of these noninformative improper priors can be written as the conjugate inverse Gamma prior. The inverse Gamma distribution is given by:

$$p(\sigma^2) = \frac{\beta^\alpha}{\Gamma(\alpha)} (\sigma^2)^{-(\alpha+1)} \exp\left(-\frac{\beta}{\sigma^2}\right) \text{ with shape } \alpha > 0 \text{ and scale } \beta > 0$$

When the shape parameter $\alpha = 0$ and the scale parameter $\beta = 0$, we obtain $p(\sigma^2) \propto \sigma^{-2}$. When the shape parameter $\alpha = -\frac{1}{2}$ and the scale parameter $\beta = 0$, we obtain $p(\sigma^2) \propto \sigma^{-1}$. When the shape parameter $\alpha = -1$ and the scale parameter $\beta = 0$, we obtain $p(\sigma^2) \propto 1$.

Table 2 presents these priors for all variance components in our model. For the intercepts, means, loadings, and regression coefficients, the standard noninformative improper prior is the uniform prior from $-\infty$ to $+\infty$. The vague proper prior $N(0, 10^{10})$ approximates this uniform prior.

Thus, for the intercepts, means, loadings and regression coefficients, we will only investigate vague proper and empirical Bayes priors, which are discussed next.

[TABLE 2]

Vague proper priors

A common solution to avoid improper posteriors while keeping the idea of noninformativeness in the prior is to specify vague proper priors. These priors are formal probability distributions, where the hyperparameters are chosen such that the information in the prior is minimal. In the case of variance parameters, vague proper priors can be specified as conjugate inverse Gamma priors with hyperparameters close to zero, typically 0.1, 0.01, or 0.001. These priors approximate the improper prior $p(\sigma^2) \propto \sigma^{-2}$ (Berger, 2006). The latter option, $IG(\epsilon, \epsilon)$, with $\epsilon = 0.001$ is used as example throughout the WinBUGS manual. We will consider these three typical prior specifications for the variance parameters in our model. Note that smaller hyperparameters lead to a prior that is more peaked around zero. For means and regression parameters, we will investigate a normal prior with a large variance. This vague proper prior approximates a flat prior.

Specifically, we shall use the normal prior $N(0, 10^{10})$, which is the default in Mplus. In addition, we will consider the blavaan default setting for location parameters, which is the normal prior $N(0, 1000)$ for the measurement intercepts and the normal prior $N(0, 100)$ for the loadings, structural intercepts, and structural regression coefficients. We will label this prior coupled with $\pi(\sigma^2) \propto 1$ for variances “vague normal”. The vague proper priors that will be considered throughout this paper are summarized in Table 2.

A potential problem of vague proper priors is that the exact hyperparameters are arbitrarily chosen, while this choice can greatly affect the final estimates. For example, there is no clear rule stating how to specify the shape and scale parameter of the inverse Gamma prior, namely, 0.1, 0.01, 0.001, or perhaps even smaller. Gelman (2006) showed that in a multilevel model with 8 schools on the second level, the posterior for the between-school variance was completely dominated by the inverse Gamma prior with small hyperparameters. In addition, the inverse Gamma prior depends on the scale of the data. Specifically, an $IG(0.1, 0.1)$ prior might be a noninformative choice when the data are standardized, but can be very informative when the data are unstandardized. It is yet unclear how this prior performs in structural equation models, which are considerably more complex than the 8 schools example studied by Gelman (2006), in terms of the number of parameters and the relations between them.

Empirical Bayes priors

The third type of default priors we consider is empirical Bayes (EB) priors. The central idea behind the EB methodology is that the hyperparameters are chosen based on the data at hand (see e.g., Carlin and Louis, 2000a, Ch. 3). This results in a prior with substantial probability mass in the region where the likelihood is concentrated. EB methodology can be seen as a compromise between classical and Bayesian approaches (Casella, 1992). Since all the data are used to inform the prior distribution, EB methods are useful for combining evidence (e.g., across neighborhoods, Carter and Rolph, 1974; or across law schools, Rubin, 1980). In our application, we expect an EB prior informed by the data of all countries to provide better estimates because it adds more

information to the analysis than improper or vague priors, or ML estimation. There is also a computational advantage of EB priors as noted by Carlin and Louis (2000b) who state that the Markov Chain Monte Carlo (MCMC) sampler based on EB priors can be more stable.

We focus on the parametric EB approach, in which a specific distributional form of the prior is assumed, typically conjugate, with only the hyperparameters unknown. Different methods have been proposed to obtain the hyperparameters in this setting. First the hyperparameters can be estimated using the marginal distribution of the data (i.e., the product of the likelihood and the prior integrated over the model parameters: $p(Data) = \int f(Data|\theta)p(\theta)d\theta$, where θ is the vector with parameters, with prior $p(\theta)$, and f denotes the likelihood of the data given the unknown model parameters) (see e.g., Carlin and Louis, 2000a; Casella, 1985, Laird and Ware, 1982). An example is the EB version of the well-known g -prior of Zellner (1986). The g -prior is centered around a reference (or null) value where g controls the prior variance. In EB methodology, g is estimated from the marginal distribution (e.g., see Liang et al., 2008, and the references therein). The resulting EB prior can be seen as the best predictor of the observed data. The difficulty with this approach, however, is that an analytic expression of the marginal distribution of the data may not be available for large complex models. This is also the case in structural equation models, and therefore we will not use the marginal distribution of the data to construct an EB prior in this paper.

A second EB approach, which is simpler to implement, is to specify weakly informative priors centered around the estimates (e.g., ML) of the model parameters. These priors contain minimal information so that the problem of double use of the data is negligible. This type of EB prior has been investigated in generalized multilevel modeling (Kass & Natarajan, 2006; Natarajan & Kass, 2000) and structural equation modeling (Dunson, Palomo, & Bollen, 2005). This approach, however, may perform badly when the estimates that are used for centering the EB priors are unstable, which is the case in structural equation modeling with small samples.

For this reason an alternative EB prior is proposed which is novel in the BSEM literature. The idea is to first center the prior around a reference value to minimize its dependence on unstable

estimates. Subsequently, the other hyperparameters are estimated from the data from a simplified model to keep the solution tractable. The idea of this EB prior was inspired by the constrained posterior prior approach of Mulder et al. (2009, 2010) for Bayesian model selection. As will be shown the proposed EB prior generally contains less information than the information of one observation. For this reason, the double use of the data and the resulting underestimation of the posterior variance, a known problem of EB methodology (e.g., Darnieder, 2011; Carlin & Louis, 2000a; Efron, 1996), is expected to be negligible.

EB priors for intercepts, means, factor loadings, and regression coefficients. Here we discuss the EB prior for the intercepts, means, factor loadings, and regression coefficients for which the conditionally conjugate prior is normally distributed. To estimate the hyperparameters, i.e., prior mean and the prior variance, we simplify the endeavor by constructing a normal prior for α , denoted by $N(\mu_\alpha, \tau_\alpha^2)$, for the following model

$$y_i = \alpha + \epsilon_i \text{ with } \epsilon_i \sim N(0, \sigma^2), \quad (1)$$

for $i = 1, \dots, n$ observations, with unknown error variance σ^2 . Note that α in the model in (1) denotes a location parameter, i.e., an intercept, mean, factor loading, or regression coefficient in the model considered in this paper. The prior mean μ_α is set equal to a reference (or null) point value to avoid the heavy dependence on the data. In the current setting, we set the prior mean equal to zero because of its special meaning of “no effect”. The prior variance τ_α^2 is then estimated as the variance in a restricted model where the prior mean, $\mu_\alpha = 0$, is plugged in for α , i.e., $y_i \sim N(0, \tau_\alpha^2)$, for $i = 1, \dots, n$. The variance in this model can be estimated as $\hat{\tau}_\alpha^2 = n^{-1} \sum_{i=1}^n y_i^2 \approx E[Y^2]$. Subsequently, an expression for the estimate $\hat{\tau}_\alpha^2$ can be obtained by deriving $E[Y^2]$ from the original model (1) according to

$$E[Y^2] = \text{Var}(Y) + (E[Y])^2 = \sigma^2 + \alpha^2. \quad (2)$$

Thus, the prior variance is chosen to be equal to the sum of the estimated error variance and the square of the estimated effect, i.e., to $\hat{\tau}_\alpha^2 = \hat{\sigma}^2 + \hat{\alpha}^2$.

This prior has two important properties. First, this prior has clear positive support where the likelihood is concentrated, which is a key property of an EB prior. This can be seen from the fact that the prior variance will be large (small) when the difference between the observed effect and the prior mean, $\hat{\alpha}$, is large (small). Note however that by centering the prior around a reference value instead of the observed effect, the prior will be less sensitive to the instability of the ML estimates. Second, this EB prior does not contain more information than the information of a single observation. To see this note that the standard error of α is equal to $\sigma/\sqrt{n}$. Furthermore, the EB prior standard deviation is smallest when $\hat{\alpha} = 0$, in which case $\hat{\tau}_{\alpha} = \hat{\sigma}$, which corresponds to the standard error based on a single observation. If $\hat{\alpha} \neq 0$, then $\hat{\tau}_{\alpha} > \hat{\sigma}$, which implies less information than the information in a single observation. For this reason, we expect the problem of using the data twice (i.e., for prior specification and estimation) to be negligible. This behavior is illustrated in Figure 2. In the case of no effect, i.e., $\hat{\alpha} = 0$ (Data 1), the prior variance is equal to the error variance. When $\hat{\alpha} \neq 0$ (Data 2), the prior variance becomes larger, i.e., the error variance plus the squared estimated effect, to ensure positive support around $\hat{\alpha}$. Note that the EB prior behaves similarly to the constrained posterior prior (Mulder, Hoijtink, & Klugkist, 2010) with the difference that the EB prior is simpler to compute.

[FIGURE 2]

This methodology will be used to construct EB priors for the intercepts, means, factor loadings, and regression coefficients. For example for the measurement intercept of y_2 , denoted by ν_2^y , the EB prior variance is equal to the squared ML estimate, i.e., $(\hat{\nu}_2^y)^2$, plus the ML estimate of the variance of the error δ_{y2}^y , i.e., $\hat{\sigma}_{y2}^2$, and thus, the EB prior is distributed as, $\nu_2^y \sim N(0, (\hat{\nu}_2^y)^2 + \hat{\sigma}_{y2}^2)$. An overview of all EB priors for location parameters is given in Table 2.

EB priors for variance components. The above methodology to construct EB priors for location parameters cannot be used for variance components because a clear reference (or null) value is generally unavailable for these types of parameters. Therefore we will consider two alternative approaches for the priors for the variances which will be combined with the EB prior for intercepts, means, factor loadings, and regression coefficients. In the first combination, which

we will label EB1, the priors for the variance components are centered around the ML estimate, as was also considered by Natarajan & Kass (2000). Again we consider the conditionally conjugate prior which has an inverse gamma distribution with a shape parameter and a scale parameter. The shape parameter controls the amount of prior information. By setting the shape parameter to $\frac{1}{2}$, the prior carries the information that is equivalent to one data point (Gelman et al., 2004, p. 50). For this reason the double use of the data is also not a serious concern in this case. The scale parameter is chosen such that the prior median equals the ML estimate, i.e., $\hat{\sigma}^2 \cdot Q^{-1}(\frac{1}{2}, \frac{1}{2})$ with $\hat{\sigma}^2$ denoting the ML estimate of the variance parameter and Q^{-1} denoting the regularized inverse Gamma function (see Table 2). A potential issue of this prior is its heavy dependence on the ML estimates of the variances. Therefore in the second combination, labeled EB2, noninformative uniform priors are specified for the variances. Thus, the EB2 prior can be seen as a hybrid EB prior, where the priors for location parameters depend on the data and the priors for variance parameters are completely independent of the data. ²

A simulation study of default BSEM analyses

Even though all the discussed priors reflect some form of default noninformative BSEM analysis, each choice may result in different conclusions. A simulation study was set up to investigate the performance of the different default priors in the industrialization and political democracy model, a classical SEM application. A common method to check the performance of objective priors is to look at their frequentist properties (Bayarri & Berger, 2004). In particular, we were interested in (1) convergence of the Bayesian estimation procedure, (2) (relative) bias, (3) mean squared error (MSE), (4) coverage rates, (5) quantiles, and (6) type 1 error rates of the direct and indirect effects. We will end the section with a general conclusion regarding the performance of the different default priors.

For the data generation, we considered four different sample sizes: 35, 75, 150, and 500. For such a complex SEM, a sample size of 35 or 75 might seem extremely small. However, BSEM is often recommended especially in situations where the sample size is small (Heerwegh, 2014; Hox et al.,

2012; Lee & Song, 2004). Furthermore, the original sample size was only 75. We expect the influence of the prior to decline as sample size increases. The population values were equal to the ML estimates of the original data (see the online supplemental material for an overview). We manipulated the population values for the direct effect γ_{65} and the indirect effect $\gamma_{60} \cdot b_{21}$, since these are the parameters of substantive interest in the model. We also manipulated two loadings of $\mathbf{y}$ (λ_4^y and λ_8^y) and the variances of the pseudo-latent variables Ω_D , which represent the measurement error correlations. These variances were manipulated because previous research indicates that the vague proper priors for variances are especially influential when the variance parameter is estimated to be close to zero (Gelman, 2006). The manipulations of the population values are shown in Table 3 and were selected in such a way that we obtained a wide range of values. We used a fractional design in which we simulated data under two combinations of population values: 1) combinations of population values for the direct and indirect effect ($3 \times 3 = 9$ conditions), and 2) combinations of population values for the loadings and variances of error correlations ($3 \times 2 = 6$ conditions). In total, this resulted in 15 different populations. From each population, we generated 500 datasets per sample size.³ Table 3 presents an overview of all 60 data-generating conditions.

[TABLE 3]

Each condition was analyzed using the nine different default prior combinations with the same type of prior being specified for all parameters in the model at once, i.e., three noninformative improper priors, three vague proper priors, two EB priors, and the vague normal setting. Note that in the case of the noninformative improper and vague proper priors only the priors on the variance parameters change, while the priors on the mean and regression parameters are specified as the normal prior $N(0, 10^{10})$. In addition, ML estimation was included for each condition, leading to a total of $10 \times 15 \times 4 = 600$ conditions, as shown in Table 3.

The EB priors are based on the ML estimates, which sometimes included Heywood cases (i.e., an estimated negative variance). In the case of negative ML estimates for the variance parameters, we set the prior median for the EB prior on variance parameters equal to 0.001. Preliminary

analyses showed that the precise choice had little effect on the posterior. For the mean, intercept, loading, and regression parameters the residual variances of the model equations were sometimes estimated to be negative, in which case we fixed them to zero for computation of the prior variance. Again, preliminary analyses indicated that the exact choice did not have any clear influence on the posterior, as long as the estimate was fixed to a small number, e.g., 0.001.

For each analysis we ran two MCMC chains and discarded the first half of each chain as burn-in. Convergence was originally assessed using the potential scale reduction (PSR), taking $\text{PSR} < 1.05$ as a criterion, with a maximum of 75,000 iterations. Based on reviewers' comments, we reran the conditions for $N = 35$ with a fixed number of iterations, first 50,000 for each chain and then 100,000 for those replications that did not obtain a $\text{PSR} < 1.05$ (using the population values as starting values). We then assessed convergence by selecting those replications with a $\text{PSR} < 1.1$ ⁴ and manually checked a part of the traceplots. Given that the results did not differ substantially from the original results, we only took this approach for $N = 35$.⁵

Estimation error was assessed using (relative) bias and mean squared error (MSE). Bias was computed as $\frac{1}{S} \sum_{s=1}^S (\hat{\theta}_s - \theta)$ with S being the number of converged replications per cell, θ being the population value of that parameter, and $\hat{\theta}_s$ being the posterior median⁶ for that parameter in a specific replication s . Relative bias was computed as $\frac{1}{S} \sum_{s=1}^S \left(\frac{\hat{\theta}_s - \theta}{\theta} \right)$ and is only defined in those population conditions where $\theta \neq 0$. MSE was computed based on the true population value as: $\frac{1}{S} \sum_{s=1}^S (\hat{\theta}_s - \theta)^2$. We obtained the 95% coverage interval by computing how often the population value was contained in the 95% credible or confidence interval. In addition, to check how well the posteriors reproduced the sampling distributions, we investigated the 2.5% and 97.5% quantiles for every parameter. The 2.5% (97.5%) quantile was computed as the proportion of times that the lower (upper) bound of the 95% confidence/credible interval was higher (lower) than the population value. Ideally these should equal 2.5% and 97.5%, respectively. Finally, we looked at the Type 1 error rates for the direct effect γ_{65} and the indirect effect $\gamma_{60} \cdot b_{21}$. The ML results are included for comparison. All analyses were done in Mplus (version 7.2) and R, using the package MplusAutomation (Hallquist & Wiley, 2014).

Convergence

Table 4 shows the percentage of converged replications for each prior and sample size, averaged across the population values. For $N = 35$, the EB2 prior resulted in the highest convergence (98.9%), followed by the vague normal prior (94.1%). For all priors convergence generally increased with sample size and there was almost no convergence for the improper prior $\pi(\sigma^2) \propto \sigma^{-2}$ when $N \leq 150$. This is not surprising since this prior is known to result in improper posteriors for variances of random effects in multilevel analysis (e.g., Gelman, 2006). Because of the severe nonconvergence under this improper prior we shall not consider it further in this paper. Note that this may imply that the vague proper priors $IG(\epsilon, \epsilon)$, with $\epsilon = 0.1, 0.01$, or 0.001 can perform badly, as they approximate the improper prior $\pi(\sigma^2) \propto \sigma^{-2}$. Specifically, traceplots of converged replications for these vague proper priors showed occasional peaks, resulting in relatively high posterior medians for those replications. We will only present the results in those population conditions with at least 50% convergence and we only consider the converged replications. Convergence percentages for each separate population condition are available in the supplemental material. The ML analysis always converged but often resulted in estimated negative variances. Specifically, in 53.9% of the replications at least one Heywood case occurred.

[TABLE 4]

Bias

Table 5 presents the relative bias, with the bias in brackets, for selected parameters in the model, for $N = 35$ and $N = 75$ averaged across population values. Results for all parameters in the model are available in the online supplemental material. Following L. K. Muthén and Muthén (2002), relative biases exceeding 10% are regarded as substantial and shown in bold. Given that the influence of the prior is greatest for small samples, we will focus on $N = 35$ and $N = 75$ in presenting the results and only mention the results for $N = 150$ and $N = 500$ briefly. The results for $N = 150$ and $N = 500$ can be found in the supplemental material.

For $N = 35$ the vague proper priors resulted in relative biases greater than 0.10 for most location

parameters, followed by the EB priors. The noninformative priors (Mplus default, $\pi(\sigma^2) \propto \sigma^{-1}$) and the vague normal prior resulted in only a few parameters with substantial bias. ML performed best, with relative bias close to zero for all location parameters. For all priors, some parameters showed more bias than others. Specifically, the measurement intercepts ν_y and the structural intercepts α often resulted in large relative bias.

For $N = 75$, the bias decreased for all priors, resulting in only a few location parameters with relative bias exceeding 10% for the noninformative improper, vague proper, and vague normal priors. The EB priors performed worst, with 6 and 7 location parameters showing relative biases greater than 0.10, while ML again resulted in relative bias close to zero for all parameters. Again, the measurement intercepts ν_y showed most bias. For $N = 150$ only the vague proper and EB priors showed relative biases greater than 0.10 for some location parameters, while for $N = 500$ the relative bias was close to zero for all priors and location parameters.

For the variance parameters in the model, when $N = 35$, the vague proper prior $IG(0.1, 0.1)$, the vague normal prior, and the EB2 prior resulted in most cases with substantial bias, followed by the improper priors, and the vague proper prior $IG(0.01, 0.01)$. The vague proper prior $IG(0.001, 0.001)$ and the EB1 prior performed good, with only 3 variance parameters having relative biases greater than 0.10, and ML performed best with only 2 parameters with relevant bias. Generally, the estimated latent variable variances showed more bias than the estimated error variances, especially $\omega_{\zeta 65}^2$ which had relative bias greater than 0.10 for ML and all priors, except the $IG(0.001, 0.001)$ prior.

For $N = 75$, the vague normal and EB2 prior resulted in most biased variance parameters, followed by the Mplus default setting (i.e., the improper prior $\pi(\sigma^2) \propto 1$ combined with the $N(0, 10^{10})$ prior), then the vague proper priors $IG(0.001, 0.001)$ and $IG(0.1, 0.1)$, and next the EB1 prior. The improper prior $\pi(\sigma^2) \propto \sigma^{-1}$ and the vague proper prior $IG(0.01, 0.01)$ performed well, with only 2 variances exceeding the relative bias of 10%, and ML performed best. Again, the estimates for the latent variable variances were generally more biased than the estimates for the error variances. For $N = 150$, all methods resulted in some biased variance parameters,

except the improper prior $\pi(\sigma^2) \propto \sigma^{-1}$ and ML estimation. For $N = 500$, only IG(0.001, 0.001), IG(0.01, 0.01), the Mplus default, and the EB1 prior resulted in some biased variance parameters.

[TABLE 5]

Overall, ML estimation performed best in terms of relative bias for both the variance and location parameters. This can be explained by the fact that ML estimation does not force separate variance components to be positive. For the location parameters, the vague proper priors and the EB priors performed worst and for the variance parameters, the vague normal and EB2 prior performed worst. Of the Bayesian methods, the Mplus default setting performed best for the location parameters and the EB1 prior performed best for the variance parameters.

Mean squared error

Figure 3 shows for each prior and type of parameter the mean squared errors (MSEs) relative to the MSE of ML estimation per population value and parameter on the logarithmic scale, $\ln(\text{MSE}_{\text{Bayes}}/\text{MSE}_{\text{ML}})$. The results are categorized by structural regression coefficients, intercepts and latent mean, factor loadings, and variance parameters. Note that the vertical axis is truncated at $\ln(\text{MSE}_{\text{Bayes}}/\text{MSE}_{\text{ML}}) = 4$, excluding the extreme situations in which the MSE of the prior is more than $\exp(4) \approx 55$ times higher than the MSE of the maximum-likelihood estimate, which occurred for the vague proper priors. Tables with the numerical values for the MSE are available in the supplemental material. For $N = 35$, the vague proper priors performed worst, especially for the intercepts and factor loadings. This can be explained by the occasional extreme values drawn for these priors, as noted previously, and this unstability in the MCMC sampler is due to the fact that the vague proper priors approximate the problematic improper prior $\pi(\sigma^2) \propto \sigma^{-2}$. All other Bayesian methods resulted in smaller or approximately equal MSEs in comparison to ML estimation. In particular, the EB priors and the vague normal prior performed best for the structural regression coefficients. For $N = 75$, $N = 150$, and $N = 500$, there were hardly any clear differences between the MSEs using the different methods.

We can thus conclude that all methods perform similarly in terms of MSEs, except for the vague

proper priors which performed considerably worse.

[FIGURE 3]

Coverage rates

Table 6 shows the coverage rates of the 95% confidence intervals and 95% Bayesian credible intervals for selected parameters. Coverage rates higher than 97.5% or lower than 92.5% are considered as substantially deviating from the desired 95%, and are marked in bold. Coverage rates for all parameters in the model are available in the online supplemental material. For $N = 35$, the EB1 and EB2 priors performed worst with low coverage for 13 and 12 location parameters, respectively, followed by the vague proper prior $IG(0.01, 0.01)$ which showed low coverage for 9 parameters. The vague normal prior showed coverage rates for 8 parameters that were too high, as did the Mplus default setting for 5 parameters. ML estimation resulted in coverage rates that were too low for 4 parameters. The improper prior $\pi(\sigma^2) \propto \sigma^{-1}$ and the vague proper priors $IG(0.001, 0.001)$ and $IG(0.1, 0.1)$ performed best.

For $N = 75$, the vague proper prior $IG(0.001, 0.001)$ performed worst with low coverage rates for 16 location parameters, followed by the EB1 prior with low coverage rates for 14 parameters. The EB2 prior also showed low coverage for 8 parameters. The vague normal prior showed coverage rates for 2 parameters that were slightly too high. The other priors and ML estimation resulted in coverage rates between 92.5% and 97.5% for all parameters. For $N = 150$ and $N = 500$, all methods resulted in coverage rates for the location parameters close to 95%, except for the EB priors when $N = 150$.

We will now discuss the coverage rates for the variance parameters.⁷ For $N = 35$, ML estimation performed worst with low coverage rates for 19 parameters. The EB1 prior also showed low coverage for 14 parameters, as did the EB2 prior for 9 parameters, while the Mplus default setting showed coverage rates for 11 parameters that were too high, as did the vague normal prior for 9 parameters. The improper prior $\pi(\sigma^2) \propto \sigma^{-1}$ and the vague proper priors showed coverage rates for 6 or 8 parameters outside the range of 92.5% and 97.5%. For $N = 75$, coverage rates

improved for most priors and ML estimation, except for the vague proper prior $IG(0.001, 0.001)$ and the EB1 prior which resulted in extreme values for 16 and 15 variance parameters, respectively. For $N = 150$, coverage rates for all variance parameters were close to 95% for the vague normal and EB2 prior. The EB1 prior performed worst, followed by the $IG(0.001, 0.001)$ and $\pi(\sigma^2) \propto \sigma^{-1}$ priors. For $N = 500$, ML estimation performed best and the EB1 prior performed worst, followed by the $\pi(\sigma^2) \propto \sigma^{-1}$ and $IG(0.001, 0.001)$ priors.

In summary, for location parameters the EB priors performed worst in terms of coverage rates, while for the variance parameters ML estimation, the vague proper prior $IG(0.001, 0.001)$, and the EB1 prior performed worst. Across all parameters, the improper prior $\pi(\sigma^2) \propto \sigma^{-1}$ and the vague proper prior $IG(0.1, 0.1)$ performed best.

[TABLE 6]

Quantiles

Even in the case of perfect coverage rates of 95%, it may be that the underestimation of the interval estimate occurs much more often than overestimation (or the other way around). To assess this, we investigated the lower 2.5% and upper 97.5% quantiles. The quantiles were obtained by computing how often the lower 2.5% and upper 97.5% bounds of the credible/confidence intervals were above or below the true population value. Figure 4 shows the quantiles for $N = 35$ with the dashed lines indicating 2.5% and 97.5%. The results are categorized by structural regression coefficients, intercepts and latent mean, factor loadings, and variance parameters. The numerical values on which Figure 4 is based are available in the supplemental material. For the lower quantile, all priors resulted in quantiles close to the desired 2.5% except for the two EB prior settings in the case of intercepts and the vague proper priors in the case of factor loadings. For the upper quantile, the noninformative improper priors and the vague normal prior generally performed best with upper quantiles close to or slightly higher than 97.5%. For the structural regression coefficients and factor loadings, the EB priors performed worst, followed by ML estimation. For the variance parameters, the EB1 prior performed badly,

as did ML estimation. For the intercepts, all priors and ML estimation resulted in quantiles close to the desired 97.5%.

[FIGURE 4]

Figure 5 shows the quantiles for $N = 75$, which were all closer to the desired quantiles compared to $N = 35$. For the lower quantile, the EB priors resulted in too high quantiles for the intercepts. Note also the outliers for the vague proper prior $IG(0.001, 0.001)$ across parameters. For the upper quantile, the EB priors again performed worst for the structural regression coefficients and loadings, followed by ML estimation. For the variance parameters, the EB1 prior and ML estimation also performed badly, while for the intercepts all priors and ML estimation resulted in quantiles close to the desired 97.5%. For $N = 150$, most priors resulted in quantiles close to the desired 2.5% and 97.5%, except for the vague normal and EB2 priors which were slightly higher or lower for some parameters. For $N = 500$, all methods generally resulted in correct quantiles. To conclude, overall the EB priors performed worst in terms of quantiles, while the noninformative improper priors and the vague normal prior performed best.

[FIGURE 5]

Direct and indirect effect

In practice, researchers often are mainly interested in those parameters related to the research question. In this model the parameters of substantive interest are the direct effect γ_{65} and the indirect effect $\gamma_{60} \cdot b_{21}$. Figure 6 shows the MSEs of the direct and indirect effect for the different priors and ML estimation for $N = 35$ and $N = 75$. The figure shows that for $N = 35$, the vague normal and EB priors resulted in the smallest MSEs for the direct effect; the other methods showed slightly worse results. For the indirect effect and $N = 35$, the smallest MSEs were obtained using the EB2 prior, followed by the EB1 prior, and the vague normal prior. ML estimation performed worst.

[FIGURE 6]

Furthermore, we consider the type 1 error rates for the direct and indirect effect, i.e., the

percentage of replications for which zero is not included in the 95% credible interval, or in the 95% confidence interval for the ML estimates, when the population value is zero. Table 7 shows the type 1 error rates for the different priors and different population conditions. For $N = 35$, the error rates for the direct effect were closest to the nominal 5% for the vague proper priors. For the Mplus default setting and the vague normal prior, the rates were too low. For the EB priors and for ML estimation the type 1 error rates were too high. For the indirect effect, the differences between priors were smaller and in general, most priors resulted in error rates close to 5%, except for the vague normal and EB2 priors which resulted in error rates that were too low. For $N = 75$, the EB priors again resulted in error rates that were too high for the direct effect, while the rates for ML estimation were closer to 5%. For the indirect effect, all priors and ML estimation generally resulted in error rates slightly higher than 5%. For $N = 150$ and $N = 500$ all type 1 error rates were generally close to 5%, except for the direct effect for the vague proper priors and EB priors when $N = 150$.

[TABLE 7]

Based on these results, we can conclude that there is not one default prior that performed consistently better than the other priors or than ML estimation across all parameters or outcomes, especially in small samples. When looking across all parameters for $N = 35$, ML estimation performed best in terms of bias; in terms of MSE, the EB priors performed best. However, both the EB priors and ML estimation performed badly in terms of coverage and quantiles. The vague proper priors performed worst in case of bias and MSE but best in terms of coverage and type 1 error rates for the direct effect and indirect effect. Overall, the noninformative improper priors $\pi(\sigma^2) \propto 1$ and $\pi(\sigma^2) \propto \sigma^{-1}$ performed best across all parameters and outcomes with good coverage rates and quantiles, and average bias and MSE. One disadvantage of the improper priors in small samples is their high nonconvergence percentage, especially for $\pi(\sigma^2) \propto \sigma^{-1}$. Of the noninformative improper priors, the Mplus default setting (i.e., $\pi(\sigma^2) \propto 1$ for variances coupled with $N(0, 10^{10})$ for location parameters) performed best in terms of bias for the location

parameters, but worse than $\pi(\sigma^2) \propto \sigma^{-1}$ for the variance parameters. In terms of coverage rates, the Mplus default setting outperformed the $\pi(\sigma^2) \propto \sigma^{-1}$ prior, especially for the variance parameters. For the MSE and quantiles, they did not differ substantially. When considering the direct and indirect effect, the parameters of practical interest, both noninformative improper priors performed good in terms of bias and coverage. Finally, although the type 1 error rate for the direct effect was slightly too low for the Mplus default setting when $N = 35$, this result was not available for the $\pi(\sigma^2) \propto \sigma^{-1}$ prior, due to high nonconvergence. Based on these differences in performance between the noninformative improper priors, we thus recommend the Mplus default priors as general choice for BSEM, with the important observation that this setting does not perform perfectly.

Even though the different default priors that were investigated in this section are routinely used in practice, their performance varies greatly across conditions. Therefore a prior sensitivity analysis is highly recommendable, especially in small samples when clear prior information is absent. The next section will provide guidelines on how to perform such an analysis in default BSEM.

A practical guide to prior sensitivity analysis

Given the results of the simulation study and in line with recommendations by B. O. Muthén and Asparouhov (2012, p. 320), prior sensitivity analysis is an important step in BSEM. The goal of a prior sensitivity analysis is to assess whether the results of the BSEM analysis are influenced by the specific default prior that is used. When the conclusions are similar using the different default priors, we can be confident that the results are reliable and robust to default prior specification. On the other hand, if the different default approaches result in substantially different conclusions, some care must be taken regarding the reliability of the results.

A prior sensitivity analysis can be conducted by rerunning the analysis with different choices for the prior. Due to the large number of parameters, possible prior choices, and possible settings for each choice, conducting a sensitivity analysis can become quite involved. Nevertheless, prior sensitivity analysis is particularly relevant in the context of SEM. Due to the complex

relationships inherent in structural equation models, a prior on a specific parameter can indirectly influence other parts of the model as well. Consequently, the effects of the prior for different parameters may cancel out, but can also accumulate. Moreover, some parameters are more influenced by their prior distribution (e.g., variances of latent variables) compared to other parameters (e.g., residual variances). If a parameter is highly influenced by its prior distribution, this influence can carry through the model and affect other parameters as well. Depaoli and Van de Schoot (2015) provided guidelines on conducting prior sensitivity analyses for Bayesian analyses with informative priors. The goal of this section is to provide a step-by-step guide on how to conduct a prior sensitivity analysis in BSEM using default priors. This analysis is recommended when prior information is weak, or when a researcher prefers to exclude external information in the statistical analysis. We will illustrate the guidelines on the democracy and industrialization model from Section 2.

Step 1: Decide which parameters to investigate

The first step in conducting a sensitivity analysis for structural equation models is to decide which parameters to focus on. Although it is important to change the prior on each parameter, there are generally only a few parameters of substantive interest (e.g., the direct and indirect effect in the model considered throughout this paper). Therefore, we recommend to focus primarily on the parameters of substantive interest in determining the sensitivity to the prior. Which parameters are of interest will, of course, vary across different applications. In addition to the parameters of substantive interest we recommend to focus on the latent variable variances in the model as well since these are generally most sensitive to the choice of the prior. These variance parameters can therefore unduly influence other parameters as well. In this first step, it is also helpful to consider which magnitudes of changes in the parameter values would constitute meaningful differences in the parameters. These magnitudes will be used in Step 4 to determine when a parameter is sensitive to the choice of the prior.

Step 2: Decide which priors to include

The second step consists of deciding which priors to include. Software such as Mplus limits the choice of possible priors by allowing a limited set of prior choices, such as normal priors for location parameters (e.g., intercepts, regression coefficients) and inverse Gamma priors for variance parameters. Given the large number of parameters in the model, it is infeasible to alter the prior for each parameter one at a time. In addition, it is more realistic to change the prior for all parameters in the model simultaneously because in general researchers will specify the same type of default prior for all parameters in the model.

Based on the results of the simulation study, we recommend to include the following default priors in the prior sensitivity analysis: the noninformative improper priors $\pi(\sigma^2) \propto \sigma^{-1}$ and $\pi(\sigma^2) \propto 1$, and the vague proper priors $IG(\epsilon, \epsilon)$ with $\epsilon = 0.001, 0.01$, and 0.1 for variance parameters, combined with the vague proper prior $N(0, 10^{10})$ for location parameters; the vague normal prior; and the EB priors. Note that it is important to consider multiple values for ϵ when considering the vague proper prior $IG(\epsilon, \epsilon)$, since the choice of ϵ can have a large influence on the results. When the results are not robust to the exact choice of ϵ , we do not recommend using these priors for drawing substantive conclusions. On the other hand, when the results are robust to the choice of ϵ , the results are reliable for drawing substantive conclusions. We do not recommend to include the improper prior $\pi(\sigma^2) \propto \sigma^{-2}$, due to its severe nonconvergence.

Furthermore, when prior knowledge is available, a researcher can use an informative prior. This prior can be specified by choosing the hyperparameters in such a way that the resulting prior has high probability on those parameter values deemed plausible by previous research or by an expert in the field. The challenge in specifying informative priors is to specify the hyperparameters such that the prior probability that the parameter falls in a plausible parameter region equals a certain percentage, e.g., 95%. When informative priors are used, we follow the recommendation of Depaoli and Van de Schoot (2015) to compare the results of the informative priors to results obtained using default priors.

Step 3: Technical implementation (Mplus)

The R package `MplusAutomation` (Hallquist & Wiley, 2014) can be used to automatically create and run the Mplus input files for the analyses with different priors. Subsequently, the results of all analyses can be read into R simultaneously. In the supplemental material we provide the code for our sensitivity analysis, which can be used as a template for a prior sensitivity analysis using `MplusAutomation`.

One issue when conducting the analyses in an automatic way is how to assess convergence. When using Markov Chain Monte Carlo (MCMC) sampling, it is important to ensure that the chains converge to the posterior distribution. Mplus provides an automatic criterion based on the potential scale reduction (PSR). Sampling will continue until the cut-off for the PSR defined in the `BCONVERGENCE` option is reached or before that if the maximum number of iterations is reached (specified through the `BITERATIONS` option). The maximum number of iterations should depend on the model under consideration, with more complex models requiring a larger number of iterations and preliminary analyses can be conducted to get an indication of the required number of iterations. In addition, when using the PSR it is recommended to rerun the analysis with twice as many iterations to avoid preliminary fulfillment of the PSR criterion. More information on the PSR can be found in Gelman and Rubin (1992) or the Mplus User guide (L. K. Muthén & Muthén, 1998-2012). A second option is to not rely on the automatic criterion to determine the number of iterations, but to specify a fixed number of iterations (using the `FBITERATIONS` options) and subsequently assessing convergence diagnostics, such as the PSR. Again, preliminary analyses can be conducted to get an indication of the required number of iterations. Note that there exist other methods to automatically assess convergence, for example, `blavaan` has a setting which relies on the PSR in combination with Raftery and Lewis' (1992) convergence diagnostic to determine the number of draws. Regardless of which automated criterion is used to assess convergence, it is highly recommended to check the traceplots of the posterior draws for all parameters.

Step 4: Interpretation of the results

The marginal posterior distributions for each parameter in the model can be summarized in different ways. As Bayesian point estimates the mean, median, or mode can be used. By default, Mplus provides the posterior median, which is also the summary we used. In addition, we considered the 95% credible interval. As noted in Step 1, in order to conclude whether the results are sensitive to the prior, the researcher must first decide what constitutes a meaningful difference in parameters of interest, based on the application at hand. In other words, boundaries must be specified for the changes in results across the priors; if a change in a parameter exceeds this boundary, the parameter can be classified as sensitive. Changes can be evaluated by comparing the results obtained with a specific prior to the results obtained with the original prior distribution. To define a meaningful boundary, it may be helpful to set the bounds on the standardized estimates, which are generally easier to interpret. In addition, because the standardized estimates automatically include the scale of the variables, only the sensitivity of the latent mean, intercept, loading, and regression parameters needs to be considered. However, other options are possible, for example the threshold can be based on qualitative differences. One such option is to classify a parameter as sensitive if a different prior results in a different sign of the estimate, or if the EPC (Saris, Satorra, & van der Veld, 2009), or EPC-interest (Oberski, 2014) exceeds a certain cut-off. The EPC or "expected parameter change" estimates the change in a parameter when relaxing a constraint, while the EPC-interest estimates the expected change in the parameter of interest. Sensitivity with respect to other outcomes may also be evaluated, such as whether the credible interval includes zero or not; or whether model fit measures, such as the posterior predictive p-value (Gelman, Meng, & Stern, 1996), exceed a threshold such as 0.05.

The results of the prior sensitivity analysis can fall into one of three categories: (1) the results are not sensitive to the choice of the prior; (2) the results obtained using default priors do not vary, but these results differ from the result obtained using informative priors; and (3) the results vary across all priors, both default and informative. In scenario (1) we can conclude that the results are robust to the choice of the prior. In scenario (2) the information embedded in the informative prior

influences the results, as would be expected. As noted by Berger (2006), subjective prior elicitation is difficult and can generally only provide certain characteristics of the prior (e.g., the location), whereas other features (e.g., the parametric form) are typically chosen in a convenient way. Thus, in scenario (2), the researcher should be certain that the chosen informative prior is an accurate reflection of one's prior belief, in terms of the prior guess (i.e., the prior mean) and the prior uncertainty (i.e., the prior variance and prior's distributional form). If the certainty about the informative prior cannot be warranted, the analysis based on default priors is recommended for substantive conclusions. If the informative prior is an accurate representation of the prior beliefs, the results from this prior can be used for final conclusions, while the results of the default analyses can be used as a reference.

In scenario (3) there is an additional difficulty that the results are also not robust to the different default priors. This may occur when the sample is relatively small and the prior is (possibly unintentionally) relatively informative. In this situation, one option is to collect more data. The advantage of Bayesian analysis is that we can simply collect additional data and combine it with the original data, whereas classical methods, such as confidence intervals and p-values, require a fixed sampling plan before data analysis (e.g., Robert, 2007, p. 23). Thus, the researcher can continue to collect additional observations until the results are no longer sensitive to the priors. However, it is not always feasible or possible to collect more data. In that case, we recommend that the researcher reports all results or the range of results obtained using the different default priors. The range can be computed by first combining all posterior draws from the different priors, and then computing the median and the lower and upper bounds of the 95% credible interval of the combined set of posterior draws. In addition, differences between default priors can be examined graphically, for example using boxplots such as those shown in Figure 7. The interpretation should then focus on how the substantive conclusions (e.g., the effect size of the indirect effect) vary across the default priors. This mirrors the recommendation of Leamer (1983) and relates to robust Bayesian analysis (Berger, 2006), where the results from multiple prior distributions are combined to obtain a range of results.

Note that the three scenarios can occur for different parameters. If results between priors vary only for the nuisance parameters and not for the parameters of interest (i.e., the direct and indirect effect), reliable conclusions can be drawn about these parameters of interest. Only when the researcher wishes to draw conclusions about the complete fitted model, sensitivity of the nuisance parameters to the priors should be taken into account.

Empirical application: democracy and industrialization data

We applied these steps to the original data from the democracy and industrialization application, which has a sample size of 75. In addition, we took the first 35 observations of the original data to illustrate a prior sensitivity analysis in a situation where the results are quite sensitive to the choice of the prior. For comparison, we also include the ML estimates in the results.

Step 1. In this specific application, the parameters of substantive interest are the direct effect γ_{65} and the indirect effect $\gamma_{60} \cdot b_{21}$. We decided that a standardized change of 0.1 would constitute a meaningful difference in the parameters. Note that the variables have been standardized with respect to the variances of both $\mathbf{y}$ and $\mathbf{x}$.

Step 2. We include in our sensitivity analysis the noninformative improper prior $\pi(\sigma^2) \propto \sigma^{-1}$ (combined with the vague proper prior $N(0, 10^{10})$ for location parameters), the Mplus default setting (i.e., $\pi(\sigma^2) \propto 1$ for variance parameters and $N(0, 10^{10})$ for location parameters), the vague proper priors (with $\epsilon = 0.1, 0.01$, and 0.001 , combined with the vague proper prior $N(0, 10^{10})$ for location parameters), the vague normal prior, and the EB priors. The Mplus default prior setting is used as baseline to which the other priors are compared. In general, however, the original prior distribution will serve as baseline. Although we do not recommend the use of the improper prior $\pi(\sigma^2) \propto \sigma^{-2}$, we did include this prior in the sensitivity analysis for illustrative purposes.

In addition, informative priors are available in Dunson, Palomo, and Bollen (2005) for this model based on expert knowledge, which we included in our prior sensitivity analysis (see the online supplemental material). We assume that these priors reflect the available prior knowledge and the certainty about this knowledge correctly.

Step 3. We ran the analyses using MplusAutomation. All files for our analyses are available in the supplemental material. We did not rely on the PSR criterion to determine the amount of iterations, but instead specified a fixed number of 75,000 iterations (through the FBITERATIONS option). For each analysis, we checked whether the $PSR < 1.1$ for all parameters and we eyeballed each traceplot to ensure convergence. An example of a traceplot illustrating convergence and the corresponding estimated posterior densities is available in the online supplemental material.

Step 4. To assess sensitivity, we compared the standardized median for each prior with the standardized median obtained when using the Mplus default prior settings (the baseline). The Mplus default settings correspond to a normal prior, $N(0, 10^{10})$, for means and regression parameters and an improper prior, $\pi(\sigma^2) \propto 1$, for variances, implemented as an inverse Gamma prior, $IG(-1, 0)$. Note that, in practice, the original prior distributions will serve as baseline. As noted in Step 1, a standardized change of 0.1 would constitute a meaningful difference. Consequently, if the standardized median of a prior deviated more than 0.1 from the standardized median obtained under the baseline, we concluded that the results are sensitive to the prior. We will first discuss the results for the original data ($N = 75$) followed by the results for $N = 35$.

Results original data ($N = 75$)

Table 8 shows the standardized and unstandardized ML estimates and posterior medians for the direct effect for each analysis and sample size, as well as the 95% confidence and credible intervals. Standardized estimates that deviate more than 0.1 from the estimates obtained with the Mplus default setting are presented in bold. It is clear that for $N = 75$, none of the estimates exceed the cut-off and thus we can conclude that the direct effect is not sensitive to the prior. There are differences between the priors in terms of credible intervals. Specifically, for the informative prior, the lower bound is negative while it is positive for all other priors and ML estimation. Consequently, a test for the direct effect using the informative prior would result in the conclusion that the effect is not substantially different from zero, while the other priors and ML estimation would lead to this conclusion. Note that the informative prior results in the

smallest credible interval because of the additional information added through the prior. In addition, the EB priors have slightly smaller credible intervals compared to the other default priors, because the EB priors include more prior information than the default priors.

[TABLE 8]

Table 9 shows the point estimates and 95% confidence and credible intervals for the indirect effect for each analysis and sample size. None of the estimates exceed the cut-off of 0.1 and thus the indirect effect is not sensitive to the choice of the prior. The confidence and credible intervals for the indirect effect show less variation compared to the direct effect so that conclusions based on interval testing for the indirect effect would not differ across the Bayesian methods or ML estimation. Only the width of the credible intervals differ, with the informative prior showing the smallest credible intervals, followed by the EB priors.

[TABLE 9]

Results subset original data ($N = 35$)

For the subset of 35 observations from the original data, the ML analysis led to empirical weak identification and inadmissible estimates due to the small sample size and thus these results are excluded. The standardized and unstandardized posterior medians, and 95% credible intervals for the direct effect are presented in Table 8. The standardized median for the informative prior is shown in bold, indicating that this estimate differs more than 0.1 from the estimate obtained under the Mplus default setting. In addition, the informative prior is the only prior resulting in a negative lower bound of the 95% credible interval. Thus, for $N = 35$, the direct effect is sensitive to the choice of the prior. The results for the indirect effect are presented in Table 9. Again, the indirect effect is less sensitive to the choice of the prior. None of the standardized estimates exceed the cut-off and the only clear difference between the priors is that the width of the credible interval is smallest for the informative prior, followed by the EB priors.

Conclusions from the prior sensitivity analysis

Based on the scenarios described in Step 4 of the prior sensitivity guide, we can thus conclude that for $N = 75$, the estimates of the parameters of interest (i.e., the direct and indirect effect) are robust to the choice of the prior (scenario (1)). However, the credible interval for the direct effect included zero for the informative prior whereas it did not include zero for the default priors. Thus, when testing the direct effect, we find ourselves in scenario (2). The same holds for $N = 35$, where both the estimate and credible interval of the direct effect were sensitive to the informative prior. Therefore, careful consideration of the informative prior for the direct effect is necessary. The informative prior for the direct effect was the normal prior $N(0.5, 2)$, which results in 95% prior probability on the interval $(-3.43, 4.42)$ (see the online supplemental material). Compared to the default priors, which are more spread out, the informative prior shrinks the estimate for the direct effect towards the prior mean, resulting in a smaller estimate. If the informative prior has been specified with care and accurately reflects the prior beliefs (we assume this was the case), the results obtained with the informative prior can be used for substantive conclusions, which implies no significant direct effect. The default analysis, which suggest a significant direct effect, can be reported as a reference analysis to show that the information in the data implies a significant direct effect.

For both sample sizes, the nuisance parameters were sensitive to the default priors as well. Thus, if the goal of the analysis is to draw conclusions about the full model, scenario (3) is applicable. If no informative priors were specified, and if it is not possible to collect more data, the researcher should consider and report the (range of) results from all default priors. By combining the posterior draws from all default priors and computing the median and bounds of the 95% credible interval, we can obtain a range for all parameters, which is reported in Table 10. Some of the credible intervals based on all posterior draws are very wide. This common behavior of a robust Bayesian analysis (e.g., Berger, 2006) can be explained by the fact that there is very little information in the data to fit the relatively complex SEM model.

Additionally, we can examine the differences between the default priors graphically, for example

by plotting the standardized posterior medians for each parameter, as is done in Figure 7 for the structural intercept α_{60} . From Figure 7 we can see that for $N = 35$, the estimated medians vary from -1.4 to -2 and the researcher should further examine these differences between the priors. For example, in this case, the smallest estimates are obtained using the EB priors, whereas the improper and vague proper priors generally result in estimates close to -2, and the vague normal prior lies in between. Of the default priors, the EB priors are most informative, as they include information regarding the ML estimates. The improper and vague proper priors are least informative, since they have the largest posterior variance, and the vague normal prior lies in between. Thus, for more informative default priors, the estimate for α_{60} becomes smaller. If informative priors were specified and scenario (3) is applicable, the researcher should carefully consider each of the informative priors and if in doubt whether the prior accurately reflects the prior belief, the researcher should consider the results of all default priors.

[TABLE 10]

[FIGURE 7]

Discussion

Bayesian methods are a useful alternative to ML estimation for structural equation models. In the case of small samples, ML estimation can result in empirical weak identification and inadmissible estimates whereas BSEM analyses can prevent these problems. In order to use the BSEM framework, however, prior distributions must be specified for the model parameters. In this paper, we focused on default priors that can be applied in an automatic fashion for a BSEM analysis when prior knowledge is absent or if a researcher does not wish to include external information. Based on the results, we recommend the Mplus default setting (i.e., the noninformative improper prior $\pi(\sigma^2) \propto 1$ for variance parameters, combined with the vague proper prior $N(0, 10^{10})$ for location parameters) as general default prior for BSEM. In general, we recommend against the use of the improper prior $\pi(\sigma^2) \propto \sigma^{-2}$, since it suffers from major convergence problems, and against the vague proper priors which approximate this improper prior and consequently lead to

instable MCMC estimation. The vague proper priors can be considered in the prior sensitivity analysis, only when multiple values for the hyperparameters are included and the results do not vary across these choices. The performance of the different default priors varied greatly across conditions. For this reason it is highly recommended to consider several default priors when performing a default BSEM analysis, to assess robustness of the results to the choice of the prior. For $N = 35$, ML estimation performed better than the Bayesian methods in terms of bias. This can be explained by the fact that ML estimation does not force the separate variances to be positive, while the priors we have considered are only defined in the region where the separate variances are positive. In the case of small samples, the likelihood has support for negative variances while the default priors give probability zero to negative variances to obtain interpretable estimates at the cost of introducing some bias. Given that variance parameters are often nuisance parameters, one might argue that minimizing bias is preferred over interpretable estimates. It would be interesting to adopt a Bayesian approach where priors for the variances have support in the negative subspace of certain variance parameters and compare the bias to ML estimation (e.g., Mulder & Fox, 2013). Generally, however, ML estimation cannot be recommended for small samples ($N = 35$) due to the low coverage rates for the variance parameters and high type 1 error rates for testing direct effects. For larger samples (i.e., $N = 75, 150, 500$) ML performed good in terms of all outcome measures and generally outperformed the Bayesian methods. Therefore, if ML estimation is feasible, it is recommended for large samples ($N \geq 75$). Note, however, that for more complex models than the SEM studied here new simulations have to be performed to check whether the sample size of the data at hand is large enough to fit this more complex model using ML. Additionally, ML estimation has the disadvantage that confidence intervals depend on the sampling plan. This implies that optional stopping or deciding to collect more data because of inaccurate results (i.e., wide confidence intervals) is not straightforward to incorporate in a classical analysis using confidence intervals (e.g., Robert, 2007, p. 23). Bayesian credible intervals, on the other hand, abide by the likelihood principle (Berger & Wolpert, 1984), and therefore the results are invariant of the sampling plan.

We have proposed two EB priors which are novel in BSEM. Although the EB priors performed best in terms of MSE, they did not perform well in terms of bias, coverage rates, quantiles, and Type 1 error rates. Several studies have found that EB priors can result in an underestimation of the posterior variance (Carlin & Louis, 2000a; Darnieder, 2011; Efron, 1996), which can partly explain the low coverage rates. Furthermore, the EB priors considered in this paper were developed to be generally applicable. For example, the proposed EB prior for location parameters is centered around zero with the prior variance chosen such that the prior has positive support where the likelihood is concentrated. However, for some parameters such as intercepts, the prior mean of zero might not be realistic and can lead to biased estimates. In addition, it may be that in certain extreme situations (e.g., when the error variances are approximately zero), data dependent priors, such as our EB priors, should not be used. In general, we believe that the EB methodology offers interesting possibilities for BSEM, however, more research is needed for further development of good EB priors.

We provided guidelines on how to conduct a (default) prior sensitivity analysis in Mplus and illustrated these guidelines on a structural equation model from the literature. An important step is choosing which parameters are of substantive interest. If these parameters are insensitive to the default priors, the estimates can be readily interpreted even if the estimates of some nuisance parameters show prior sensitivity. If the estimated parameters of interest are sensitive to the default priors, i.e., they differ more than the chosen threshold value, the (range of) results of all default priors should be reported. To obtain robust bounds for the interval estimates, we recommend combining the posterior draws of the different default Bayesian analyses, and subsequently, reporting the upper and lower bounds of the 95% credible intervals based on the combined set of draws (e.g., Berger, 2006).

We investigated only conditionally conjugate priors since these are available in Mplus. However, many non-conjugate priors have been proposed in the Bayesian literature as more robust (i.e., less influential) alternatives. For example, Gelman (2006) and Polson and Scott (2012) proposed the half-Cauchy prior for random effects variances, which can be implemented in a Gibbs sampler

relatively easy through parameter expansion. A second option for random effects variances is a Gamma prior in combination with posterior mode estimates which has been proposed in the context of meta-analysis by Chung, Rabe-Hesketh, and Choi (2013). Note that the choice of the prior for residual variances is considerably less important than the prior for the variances of latent variables (e.g., Polson & Scott, 2012). For intercept, mean, and regression parameters a robust alternative is the t -distribution, which has been proposed as prior for logistic models by Gelman, Jakulin, Pittau, and Su (2008) and as error distribution to obtain robust models (e.g., robust growth curve models; Zhang, Lai, Lu, & Tong; 2013). The t -distribution includes the Cauchy distribution as special case when the number of degrees of freedom is set to 1. These priors should be investigated in the context of BSEM to assess their performance and determine whether they can be used as default priors.

Footnotes

¹Here, noninformative refers to the ultimate goal of the prior rather than its actual behavior. In fact, many different noninformative priors exist (Kass & Wasserman, 1996), which often result in slightly different estimates.

²We also considered a third combination: an EB prior for the variances (i.e., an inverse Gamma prior centered around the ML estimate) combined with the $N(0, 10^{10})$ prior for the location parameters. Thus, this combination can also be seen as a hybrid EB prior, where the priors for the variance parameters depend on the data and the priors for the location parameters are completely independent of the data. However, this combination performed similarly to the uniform prior for variances combined with the $N(0, 10^{10})$ prior for location parameters and we have therefore not considered it further.

³We created cumulative average plots to assess whether 500 replications were enough to attain Monte Carlo convergence, which was the case.

⁴Although less conservative, we used this cut-off because after 100,000 iterations some parameters had a PSR slightly above 1.05 while the traceplots indicated convergence based on eyeballing. An example is given in the supplemental material.

⁵Some priors and conditions were added later during the review process. Specifically, for $N = 150$ and $N = 500$, two vague proper and two noninformative improper priors were added later in those population conditions where the direct and indirect effect were manipulated, as well as the vague normal and EB2 prior in all population conditions. For these conditions, we simply ran all replications with 100,000 iterations and assessed convergence by selecting those replications with a PSR < 1.1 , given that this strategy had proven correct for $N = 35$.

⁶We compared the posterior medians with the posterior means. Overall, these two posterior summaries did not differ substantially. However, in some replications the posterior mean for a specific parameter was extremely high. This can happen when the MCMC sampler samples some extreme values which have a large influence on the posterior mean, but not on the posterior median. Therefore, we used the posterior median as point estimate.

⁷We deleted some values for the coverage rates for the vague normal prior, the vague proper priors, the Mplus default, and the EB2 prior which were equal to zero. These values occurred only for Ω_D when the population value was zero. This happened because the lower bound of the credible interval was always greater than zero for these priors and thus a population value of zero is by definition never contained in the credible interval. For the EB1 prior and the improper prior $\pi(\sigma^2) \propto \sigma^{-1}$ the lower bound did equal zero in several replications thereby resulting in a coverage not equal to zero in this situation. Note that, since the lower bound did not equal zero in all replications, the resulting coverage for these priors was low when the population values were equal to zero. For ML, the lower bound of the confidence interval can be negative.

Table 1

Overview of default priors in software packages for BSEM

Software	Type of parameter	Default prior form	Default prior hyperparameters
Mplus (L. K. Muthén & Muthén, 1998-2012)	Intercepts/loadings/slopes	Normal	$N(0, 10^{10})$
	Thresholds	Normal	$N(0, 5)$
	Variance parameters	inverse Gamma	$IG(-1, 0)$
	Variance covariance matrices continuous variables	inverse Wishart	$IW(0, -p - 1)$
	Variance covariance matrices categorical variables	inverse Wishart	$IW(I, p + 1)$
	Class proportions mixture models	Dirichlet	$D(10, 10, \dots, 10)$
Blavaan (Merkle & Rosseel, 2015)	Measurement intercepts	Normal	$N(0, 1000)$
	Structural intercepts/loadings/regression coefficients	Normal	$N(0, 100)$
	Precision residuals	Gamma	$G(1, 0.5)$
	Blocks of precision parameters	Wishart	$W(I, p + 1)$
	Correlations	Beta*	$B(1, 1)$
Stan (Carpenter et al., 2017)	All parameters	Uniform	
Amos (Arbuckle, 2013)	All parameters	Uniform	
WinBUGS (Lunn et al., 2000)	All parameters	No defaults; manually specify proper priors	
JAGS (Plummer, 2003)	All parameters	No defaults; manually specify proper priors	

Note. p represents the number of variables. I represents the identity matrix. * The Beta prior for correlations in blavaan has support $(-1, 1)$ instead of $(0, 1)$.

Table 2

Overview of the default prior specifications per parameter considered throughout this paper.

Parameter type	Parameter	Prior			
		Default	Noninformative improper	Vague proper	Empirical Bayes (EB)
Latent variable variances	ω_ξ^2	$\pi(\omega_\xi^2) \propto 1$	$\pi(\omega_\xi^2) \propto 1$	IG(0.001, 0.001)	$\text{IG}\left(\frac{1}{2}, \hat{\omega}_\xi^2 \cdot Q^{-1}\left(\frac{1}{2}, \frac{1}{2}\right)\right)$
			$\pi(\omega_\xi^2) \propto \omega_\xi^{-1}$	IG(0.01, 0.01)	
			$\pi(\omega_\xi^2) \propto \omega_\xi^{-2}$	IG(0.1, 0.1)	
	ω_ζ^2	$\pi(\omega_\zeta^2) \propto 1$	$\pi(\omega_\zeta^2) \propto 1$	IG(0.001, 0.001)	$\text{IG}\left(\frac{1}{2}, \hat{\omega}_\zeta^2 \cdot Q^{-1}\left(\frac{1}{2}, \frac{1}{2}\right)\right)$
			$\pi(\omega_\zeta^2) \propto \omega_\zeta^{-1}$	IG(0.01, 0.01)	
			$\pi(\omega_\zeta^2) \propto \omega_\zeta^{-2}$	IG(0.1, 0.1)	
	ω_D^2	$\pi(\omega_D^2) \propto 1$	$\pi(\omega_D^2) \propto 1$	IG(0.001, 0.001)	$\text{IG}\left(\frac{1}{2}, \hat{\omega}_D^2 \cdot Q^{-1}\left(\frac{1}{2}, \frac{1}{2}\right)\right)$
			$\pi(\omega_D^2) \propto \omega_D^{-1}$	IG(0.01, 0.01)	
			$\pi(\omega_D^2) \propto \omega_D^{-2}$	IG(0.1, 0.1)	
	σ_y^2	$\pi(\sigma_y^2) \propto 1$	$\pi(\sigma_y^2) \propto \sigma_y^{-2}$	IG(0.001, 0.001)	$\text{IG}\left(\frac{1}{2}, \hat{\sigma}_y^2 \cdot Q^{-1}\left(\frac{1}{2}, \frac{1}{2}\right)\right)$
			$\pi(\sigma_y^2) \propto \sigma_y^{-1}$	IG(0.01, 0.01)	
			$\pi(\sigma_y^2) \propto \sigma_y^{-2}$	IG(0.1, 0.1)	
Residual variances	σ_x^2	$\pi(\sigma_x^2) \propto 1$	$\pi(\sigma_x^2) \propto \sigma_x^{-2}$	IG(0.001, 0.001)	$\text{IG}\left(\frac{1}{2}, \hat{\sigma}_x^2 \cdot Q^{-1}\left(\frac{1}{2}, \frac{1}{2}\right)\right)$
			$\pi(\sigma_x^2) \propto \sigma_x^{-1}$	IG(0.01, 0.01)	
			$\pi(\sigma_x^2) \propto \sigma_x^{-2}$	IG(0.1, 0.1)	
	α	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{\alpha}^2 + \hat{\omega}_\zeta^2)$
			-	$\text{N}(0, 100)$	
	b	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{b}^2 + \hat{\omega}_\zeta^2)$
			-	$\text{N}(0, 100)$	
	γ	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{\gamma}^2 + \hat{\omega}_\zeta^2)$
			-	$\text{N}(0, 100)$	
	μ_ξ	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{\mu}_\xi^2 + \hat{\omega}_\xi^2)$
			-	$\text{N}(0, 100)$	
	ν_y	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{\nu}_y^2 + \hat{\sigma}_y^2)$
			-	$\text{N}(0, 1000)$	
Measurement intercepts	ν_x	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{\nu}_x^2 + \hat{\sigma}_x^2)$
			-	$\text{N}(0, 1000)$	
	λ_y	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{\lambda}_y^2 + \hat{\sigma}_y^2)$
			-	$\text{N}(0, 100)$	
	λ_x	$\text{N}(0, 10^{10})$	-	$\text{N}(0, 10^{10})$	$\text{N}(0, \hat{\lambda}_x^2 + \hat{\sigma}_x^2)$
			-	$\text{N}(0, 100)$	

Q^{-1} : denotes the regularized inverse Gamma function.

Table 3

Overview of the data generating and analysis conditions included in the simulation study.

Variable	# Levels	Values
<i>Data generating conditions</i>		
Sample size	4	$N \in \{35, 75, 150, 500\}$
Direct effect	3	$\gamma_{65} = 0$ $\gamma_{65} = 1$ $\gamma_{65} = 2$
Indirect effect	3	$\gamma_{60} \cdot b_{21} = 0 \times 0.837$ $\gamma_{60} \cdot b_{21} = 1 \times 0.837$ $\gamma_{60} \cdot b_{21} = 2 \times 0.837$
Loadings	3	$\lambda_4^y, \lambda_8^y = 0$ $\lambda_4^y, \lambda_8^y = 1$ $\lambda_4^y, \lambda_8^y = 2$
Error covariances	2	$\mathbf{\Omega}_D = \mathbf{0}$ $\mathbf{\Omega}_D = \mathbf{1}$
<i>Analysis conditions</i>		
Priors	9	<p>Noninformative improper:</p> <p>$\pi(\sigma^2) \propto 1$ and $N(0, 10^{10})$ (Mplus default)</p> <p>$\pi(\sigma^2) \propto \sigma^{-1}$ and $N(0, 10^{10})$</p> <p>$\pi(\sigma^2) \propto \sigma^{-2}$ and $N(0, 10^{10})$</p> <p>Vague proper:</p> <p>$IG(0.001, 0.001)$ and $N(0, 10^{10})$</p> <p>$IG(0.01, 0.01)$ and $N(0, 10^{10})$</p> <p>$IG(0.1, 0.1)$ and $N(0, 10^{10})$</p> <p>Vague normal:</p> <p>$N(0, 1000)$ and $N(0, 100)$ and $\pi(\sigma^2) \propto 1$</p> <p>Empirical Bayes (EB):</p> <p>EB1: $IG\left(\frac{1}{2}, \hat{\sigma}^2 \cdot Q^{-1}\left(\frac{1}{2}, \frac{1}{2}\right)\right)$ and $N(0, \hat{\mu}^2 + \hat{\sigma}^2)$</p> <p>EB2: $\pi(\sigma^2) \propto 1$ and $N(0, \hat{\mu}^2 + \hat{\sigma}^2)$</p>
Maximum Likelihood	1	

Note. The maximum likelihood (ML) estimates from the original data were: $\hat{\gamma}_{65} = 0.572$;

$\hat{\gamma}_{60} \cdot \hat{b}_{21} = 1.483 \cdot 0.837$; $\hat{\lambda}_4^y = 1.265$; $\hat{\lambda}_8^y = 1.266$; $\hat{\omega}_{D15}^2 = 0.624$; $\hat{\omega}_{D24}^2 = 1.313$; $\hat{\omega}_{D26}^2 = 2.153$;

$\hat{\omega}_{D37}^2 = 0.795$; $\hat{\omega}_{D48}^2 = 0.348$; and $\hat{\omega}_{D68}^2 = 1.356$.

Table 4

Percentage converged replications for the default priors in the simulation study, averaged across population values.

N	Mplus default	$\pi(\sigma^2) \propto \sigma^{-1}$	$\pi(\sigma^2) \propto \sigma^{-2}$	IG(0.001, 0.001)	IG(0.01, 0.01)	IG(0.1, 0.1)	Vague normal	EB1	EB2
35	73.7	53.3	0	51.4	68.9	88.7	94.1	83.3	98.9
75	99.6	99.5	1.03	97.2	99.9	99.9	99.0	99.6	100
150	100	100	1.77	100	100	100	95.0	99.9	99.6
500	100	99.8	45.6	99.8	100	100	98.7	99.8	98.8

Note. N = sample size. EB1 = Empirical Bayes prior location and variance parameters. EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB priors.

Table 5

Relative bias and bias (in brackets) for each default prior and ML estimation, averaged across population values, for selected parameters

	γ_{60}	γ_{65}	b_{21}	$\gamma_{60} \cdot b_{21}$	α_{60}	α_{65}	λ_2^y	ν_2^y	$\omega_{\zeta 60}^2$	$\omega_{\zeta 65}^2$	ω_{D15}^2	σ_{y6}^2	σ_{x2}^2
Sample size = 35													
Mplus default	0.044 (0.053)	0.026 (0.024)	-0.016 (-0.013)	0.016 (0.017)	0.13 (-0.264)	0.05 (-0.118)	0.038 (0.047)	0.069 (-0.179)	0.067 (0.266)	1.207 (0.208)	0.332 (0.239)	0.099 (0.142)	0.427 (0.051)
$\pi(\sigma^2) \propto \sigma^{-1}$	0.051 (0.068)	-0.009 (-0.002)	-0.003 (-0.003)	0.035 (0.045)	0.185 (-0.336)	0.006 (-0.034)	0.026 (0.053)	0.078 (-0.181)	0.014 (0.159)	0.373 (0.082)	0.011 (0.133)	-0.227 (-0.095)	0.134 (0.017)
IG(0.001, 0.001)	0.037 (0.056)	0 (0.005)	0.271 (0.227)	0.035 (0.043)	0.137 (-0.278)	0.611 (-1.426)	0.468 (0.588)	1.411 (-3.683)	-0.019 (-0.074)	0.005 (0.001)	-0.053 (0.017)	-0.117 (-0.17)	-0.111 (-0.013)
IG(0.01, 0.01)	-0.05 (-0.048)	-0.047 (-0.021)	0.191 (0.274)	-0.042 (-0.027)	-0.185 (0.251)	0.364 (-0.388)	0.752 (1.133)	2.015 (-2.444)	-0.134 (-0.631)	0.161 (0.033)	0.378 (0.238)	-0.208 (-0.097)	0.068 (0.006)
IG(0.1, 0.1)	-0.04 (-0.046)	0.01 (0.004)	0.034 (0.046)	-0.042 (-0.038)	-0.15 (0.244)	0.119 (-0.163)	0.155 (0.255)	0.417 (-0.715)	-0.157 (-0.728)	0.667 (0.128)	0.269 (0.25)	-0.238 (-0.111)	0.339 (0.04)
Vague normal	-0.055 (-0.062)	0.002 (0.001)	0.007 (0.006)	-0.069 (-0.067)	-0.162 (0.328)	0.012 (-0.028)	0.068 (0.086)	0.101 (-0.264)	-0.042 (-0.167)	1.186 (0.204)	0.424 (0.294)	0.099 (0.143)	0.436 (0.052)
EB1	-0.077 (-0.095)	-0.149 (-0.103)	0.011 (0.009)	-0.079 (-0.08)	-0.23 (0.468)	-0.2 (0.467)	-0.019 (-0.023)	-0.162 (0.422)	-0.04 (-0.159)	0.212 (0.036)	-0.037 (-0.002)	-0.089 (-0.128)	0.036 (0.004)
EB2	-0.079 (-0.100)	-0.155 (-0.109)	-0.011 (-0.009)	-0.099 (-0.102)	-0.236 (0.479)	-0.251 (0.585)	-0.06 (-0.075)	-0.229 (0.599)	0.084 (0.334)	1.491 (0.256)	0.268 (0.200)	0.096 (0.139)	0.499 (0.060)
ML	0.021 (0.030)	0.002 (0.006)	0.017 (0.015)	0.032 (0.035)	0.074 (-0.142)	0.036 (-0.088)	0.018 (0.024)	0.038 (-0.101)	-0.061 (-0.235)	-0.246 (-0.043)	-0.078 (-0.053)	-0.107 (-0.159)	-0.033 (-0.004)
Sample size = 75													
Mplus default	-0.016 (-0.018)	0.036 (0.029)	-0.013 (-0.01)	-0.036 (-0.035)	-0.041 (0.084)	0.039 (-0.092)	0.026 (0.033)	0.035 (-0.09)	-0.003 (-0.01)	0.688 (0.118)	0.189 (0.142)	0.016 (0.023)	0.211 (0.025)
$\pi(\sigma^2) \propto \sigma^{-1}$	-0.018 (-0.034)	0.038 (0.017)	-0.017 (-0.003)	-0.041 (-0.04)	-0.058 (0.166)	0.01 (-0.045)	0.019 (0.041)	0.043 (-0.123)	-0.016 (-0.107)	0.232 (0.037)	0.118 (0.075)	-0.094 (-0.002)	0.034 (0.001)
IG(0.001, 0.001)	-0.034 (-0.04)	-0.002 (0.002)	0.007 (0.006)	-0.038 (-0.037)	-0.094 (0.192)	0.004 (-0.009)	0.06 (0.075)	0.104 (-0.273)	-0.055 (-0.216)	-0.117 (-0.02)	0.027 (0.03)	-0.062 (-0.089)	-0.148 (-0.018)
IG(0.01, 0.01)	-0.022 (-0.039)	0 (0.008)	0.004 (0.007)	-0.022 (-0.03)	-0.078 (0.196)	0.004 (-0.053)	0.032 (0.056)	0.089 (-0.194)	-0.052 (-0.243)	-0.022 (0.002)	0.078 (0.036)	-0.097 (-0.085)	-0.066 (-0.005)
IG(0.1, 0.1)	-0.005 (-0.027)	0.023 (0.019)	-0.006 (0)	-0.016 (-0.031)	-0.018 (0.143)	0.01 (-0.075)	0.026 (0.058)	0.075 (-0.194)	-0.051 (-0.253)	0.334 (0.079)	0.063 (0.088)	-0.153 (-0.091)	0.144 (0.018)
Vague normal	-0.036 (-0.044)	0.025 (0.019)	-0.009 (-0.008)	-0.051 (-0.052)	-0.113 (0.23)	0.028 (-0.065)	0.022 (0.028)	0.036 (-0.093)	-0.004 (-0.017)	0.632 (0.109)	0.18 (0.135)	0.032 (0.046)	0.169 (0.02)
EB1	-0.068 (-0.087)	-0.102 (-0.076)	0.012 (0.01)	-0.061 (-0.067)	-0.205 (0.416)	-0.141 (0.33)	-0.012 (-0.016)	-0.104 (0.272)	-0.023 (-0.09)	-0.146 (-0.025)	-0.046 (-0.015)	-0.069 (-0.099)	-0.092 (-0.011)
EB2	-0.068 (-0.088)	-0.074 (-0.051)	-0.01 (-0.008)	-0.083 (-0.088)	-0.211 (0.429)	-0.131 (0.307)	-0.027 (-0.034)	-0.128 (0.333)	0.044 (0.174)	0.761 (0.131)	0.133 (0.108)	0.009 (0.013)	0.259 (0.031)
ML	-0.009 (-0.011)	0.020 (0.018)	0.008 (0.007)	-0.003 (-0.003)	-0.029 (0.000)	0.045 (-0.021)	0.010 (0.005)	0.024 (-0.017)	-0.029 (-0.087)	-0.135 (-0.043)	-0.041 (0.000)	-0.061 (-0.010)	-0.005 (-0.038)

Note. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior location and variance parameters. EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB priors. Values for the relative bias greater than 0.10 are shown in bold.

Table 6

Coverage rates of 95% Bayesian credible intervals and 95% confidence intervals for each default prior and ML estimation, averaged across population values, for selected parameters

	γ_{60}	γ_{65}	b_{21}	$\gamma_{60} \cdot b_{21}$	α_{60}	α_{65}	λ_2^y	ν_2^y	$\omega_{\zeta 60}^2$	$\omega_{\zeta 65}^2$	ω_{D26}^2	σ_{y4}^2	σ_{x2}^2
Sample size = 35													
Mplus default	96.5	97.2	98.2	97	96.9	97.3	98.2	97.7	97.1	96.6	98	98	97.7
$\pi(\sigma^2) \propto \sigma^{-1}$	96.2	96.3	96.7	96.4	96.2	96	97.2	97	96.4	98.2	83.3	95.3	98.8
IG(0.001, 0.001)	94.8	95	95.4	96	94.9	95	95.4	95.8	94.9	98.2	91.2	92.5	98.1
IG(0.01, 0.01)	91.8	94.5	94	93.6	91.9	94.3	90.8	91.3	86.2	98.8	94	94	99
IG(0.1, 0.1)	93.4	96.3	96	94.9	93.6	95.8	94.5	94.7	88.2	98.1	95.7	95.3	98.9
Vague normal	96.7	97.8	98.1	97	97	97.5	98.1	97.7	95.1	96.9	97.9	98.1	97.6
EB1	94.9	77.6	91.6	93.4	95	72.4	88.5	87.1	91.7	64.6	79.3	88.8	84.7
EB2	96.3	81.4	96.0	95.0	96.9	74.2	91.1	89.2	95.0	95.3	77.8	98.0	96.4
ML	93	90.2	93.7	93.8	92.7	90.9	92.9	93.1	87.4	90.7	90.4	92	94
Sample size = 75													
Mplus default	94.3	95.8	95.6	94.1	94.5	94.8	95.3	95.4	94.7	94.5	95.8	94.5	95.6
$\pi(\sigma^2) \propto \sigma^{-1}$	94.4	95.2	95.8	94.4	94.7	94.9	95	95.4	94.3	97.6	92.8	93.8	96.3
IG(0.001, 0.001)	91.6	92.2	92.7	91.6	91.8	91.8	91.5	91.8	90.5	94.7	72.5	87.9	89.9
IG(0.01, 0.01)	94.4	95	95.4	94.2	94.5	94.5	94.4	94.7	93.3	97.7	93	91.5	96.5
IG(0.1, 0.1)	94.4	95.8	95.3	94.3	94.4	95.2	94.4	94.8	93.5	97.1	95.1	92	97.9
Vague normal	95.3	96.6	97	95.1	95.6	96.3	96.2	95.9	95.2	96.3	96.4	94.9	96.3
EB1	92.8	82.5	89.1	90.8	93.3	80.6	91.3	90.3	92.6	73.2	81.5	89.1	86.3
EB2	94.2	87.4	95.1	92.5	94.8	83.3	92.9	92.2	94.6	94.1	76.8	95.1	95.2
ML	93.9	93.6	95.3	93.8	94.1	93.3	94.5	94.7	92.0	92.7	93.0	92.8	95.2

Note. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior location and variance parameters.

EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB priors. Coverage rates lower than 92.5% and higher than 97.5% are shown in bold.

Table 7

Type 1 error rates for the direct and indirect effect, for the different default priors in different population conditions

Parameter	Mplus default	$\pi(\sigma^2) \propto \sigma^{-1}$	IG(0.001, 0.001)	IG(0.01, 0.01)	IG(0.1, 0.1)	Vague normal	EB1	EB2	ML
Sample size = 35									
Direct effect	2.6	NA	5.9	4.1	3.5	2.0	15.8	12.1	10.9
Indirect effect	4.2	4.0	NA	5.1	5.2	2.9	4.4	2.8	5.1
Sample size = 75									
Direct effect	3.5	4.0	4.6	4.3	3.4	3.0	13.7	9.7	6.0
Indirect effect	6.1	6.1	6.5	6.4	6.1	5.4	7.1	6.1	6.3

Note. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior location and variance parameters. EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB priors. Some results are not available (NA) for conditions that did not have at least 50% convergence.

Table 8

Standardized and unstandardized point estimates and 95% confidence and credible intervals for the direct effect γ_{65} in the prior sensitivity analysis

Prior	Standardized estimate	Unstandardized estimate	Lower bound 95% CI	Upper bound 95% CI	Width 95% CI
Sample size = 35					
Mplus default	0.299	1.137	0.132	2.193	2.061
$\pi(\sigma^2) \propto \sigma^{-1}$	0.284	1.052	0.090	2.087	1.997
IG(0.001, 0.001)	0.270	0.990	0.074	2.041	1.967
IG(0.01, 0.01)	0.278	1.019	0.059	2.029	1.970
IG(0.1, 0.1)	0.283	1.052	0.088	2.053	1.965
Vague normal	0.293	1.085	0.114	2.086	1.972
EB1	0.270	0.975	0.160	1.741	1.581
EB2	0.274	0.997	0.137	1.812	1.675
Informative	0.090	0.225	-0.427	0.791	1.218
Sample size = 75					
Mplus default	0.183	0.574	0.098	1.092	0.994
$\pi(\sigma^2) \propto \sigma^{-1}$	0.177	0.554	0.082	1.046	0.964
IG(0.001, 0.001)	0.174	0.552	0.089	1.023	0.934
IG(0.01, 0.01)	0.175	0.549	0.078	1.031	0.953
IG(0.1, 0.1)	0.177	0.555	0.084	1.052	0.968
Vague normal	0.182	0.569	0.096	1.080	0.984
EB1	0.153	0.475	0.055	0.873	0.818
EB2	0.158	0.488	0.064	0.910	0.846
Informative	0.109	0.288	-0.126	0.678	0.804
ML	0.182	0.572	0.114	1.030	0.916

Note. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior location and variance parameters. EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal, EB, and informative priors. Standardized estimates deviating more than 0.1 from the estimate obtained under the Mplus default prior settings are shown in bold.

Table 9

Standardized and unstandardized point estimates and 95% confidence and credible intervals for the indirect effect $\gamma_{60} \cdot b_{21}$ in the prior sensitivity analysis

Prior	Standardized estimate	Unstandardized estimate	Lower bound 95% CI	Upper bound 95% CI	Width 95% CI
Sample size = 35					
Mplus default	0.430	1.572	0.556	2.938	2.382
$\pi(\sigma^2) \propto \sigma^{-1}$	0.446	1.599	0.580	2.962	2.382
IG(0.001, 0.001)	0.459	1.627	0.603	2.992	2.389
IG(0.01, 0.01)	0.456	1.622	0.609	2.996	2.387
IG(0.1, 0.1)	0.452	1.622	0.623	2.976	2.353
Vague normal	0.422	1.505	0.537	2.798	2.261
EB1	0.405	1.417	0.537	2.546	2.009
EB2	0.387	1.353	0.473	2.487	2.014
Informative	0.461	1.125	0.593	1.897	1.304
Sample size = 75					
Mplus default	0.385	1.191	0.522	1.982	1.460
$\pi(\sigma^2) \propto \sigma^{-1}$	0.394	1.221	0.548	2.017	1.469
IG(0.001, 0.001)	0.399	1.241	0.569	2.050	1.481
IG(0.01, 0.01)	0.397	1.229	0.561	2.035	1.474
IG(0.1, 0.1)	0.393	1.208	0.553	2.011	1.458
Vague normal	0.383	1.177	0.522	1.950	1.428
EB1	0.376	1.150	0.592	1.801	1.209
EB2	0.364	1.108	0.554	1.751	1.197
Informative	0.377	0.988	0.582	1.482	0.900
ML	0.396	1.242	0.542	1.941	1.399

Note. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior location and variance parameters. EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal, EB, and informative priors. Standardized estimates deviating more than 0.1 from the estimate obtained under the Mplus default prior settings are shown in bold.

Table 10

Posterior medians and lower and upper bounds of the 95% credible interval based on all posterior draws from sensitivity analyses with default priors.

Parameter	Lower bound 95% CI	Median	Upper bound 95% CI	Lower bound 95% CI	Median	Upper bound 95% CI
	Sample size = 35			Sample size = 75		
γ_{60}	0.211	0.538	0.771	0.211	0.428	0.617
γ_{65}	0.032	0.281	0.530	0.032	0.172	0.323
b_{21}	0.541	0.780	0.953	0.541	0.885	0.971
$\gamma_{60} \cdot b_{21}$	0.114	0.420	0.735	0.114	0.379	0.599
λ_1^y	0.609	0.845	0.955	0.609	0.845	0.919
λ_2^y	0.492	0.737	0.886	0.492	0.692	0.811
λ_3^y	0.456	0.721	0.871	0.456	0.717	0.826
λ_4^y	0.637	0.842	0.951	0.637	0.831	0.909
λ_5^y	0.683	0.850	0.936	0.683	0.802	0.882
λ_6^y	0.528	0.758	0.887	0.528	0.726	0.832
λ_7^y	0.668	0.839	0.931	0.668	0.818	0.894
λ_8^y	0.483	0.718	0.861	0.483	0.812	0.892
λ_1^x	0.887	0.950	0.986	0.887	0.917	0.957
λ_2^x	0.935	0.985	1.000	0.935	0.973	0.998
λ_3^x	0.742	0.868	0.935	0.742	0.867	0.919
μ_ξ	5.781	7.775	10.120	5.781	7.415	9.038
α_{60}	-3.750	-1.694	0.898	-3.750	-0.723	0.972
α_{65}	-3.458	-1.791	-0.179	-3.458	-0.999	0.022
ν_2^y	-1.959	-0.741	0.015	-1.959	-0.600	-0.086
ν_3^y	-0.637	0.258	1.091	-0.637	0.213	0.768
ν_4^y	-1.575	-0.308	0.395	-1.575	-0.692	-0.226
ν_6^y	-1.600	-0.850	-0.223	-1.600	-0.857	-0.370
ν_7^y	-0.649	0.038	0.663	-0.649	-0.108	0.384
ν_8^y	-1.076	-0.328	0.359	-1.076	-0.704	-0.238
ν_2^x	-5.936	-4.254	-3.001	-5.936	-4.060	-3.095
ν_3^x	-5.515	-3.969	-2.627	-5.515	-3.901	-2.955

Figure 1. Structural equation model describing the influence of industrialization in 1960 (ξ) on political democracy in 1960 (η^{60}) and 1965 (η^{65}).

Figure 2. Illustration EB priors for location parameters. The EB prior for α is $N(0, \tau_\alpha^2)$ with τ_α^2 equal to $\hat{\sigma}^2$ for Data 1, and equal to $\hat{\sigma}^2 + \hat{\alpha}^2$ for Data 2.

Figure 3. Mean squared error (MSE) for Bayesian estimation using different default priors divided by the MSE for maximum likelihood (ML) estimation on the natural logarithmic scale. EB1 = Empirical Bayes prior location and variance parameters. EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB priors. Vertical dashed lines indicate where the MSE for the Bayesian estimates equals the MSE for the ML estimates.

Figure 4. 2.5% and 97.5% quantiles for each default prior and maximum likelihood (ML)

estimation for $N = 35$. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior

location and variance parameters. EB2 = EB prior location parameters combined with

$\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters

combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$

prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal

$N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters.

IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters.

Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB

priors. Vertical dashed lines indicate the desired 2.5% and 97.5%.

Figure 5. 2.5% and 97.5% quantiles for each default prior and maximum likelihood (ML)

estimation for $N = 75$. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior

location and variance parameters. EB2 = EB prior location parameters combined with

$\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters

combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$

prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal

$N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters.

IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters.

Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB

priors. Vertical dashed lines indicate the desired 2.5% and 97.5%.

Figure 6. Mean squared error (MSE) for each default prior and maximum likelihood (ML) estimation for the direct effect γ_{65} and the indirect effect $\gamma_{60} \cdot b_{21}$. ML = maximum likelihood estimation. EB1 = Empirical Bayes prior location and variance parameters. EB2 = EB prior location parameters combined with $\pi(\sigma^2) \propto \sigma^{-1}$ prior variance parameters. Vague normal = $\pi(\sigma^2) \propto 1$ prior for variance parameters combined with the normal $N(0, 1000)$ prior for measurement intercepts and the normal $N(0, 100)$ prior for the other location parameters. Mplus default = $\pi(\sigma^2) \propto 1$ combined with the normal $N(0, 10^{10})$ prior. $\pi(\sigma^2) \propto \sigma^{-1}$ = noninformative improper priors variance parameters. IG(0.001, 0.001), IG(0.01, 0.01), IG(0.1, 0.1) = vague proper priors variance parameters. Location parameters have the normal $N(0, 10^{10})$ prior, except for the vague normal and EB priors.

Figure 7. Standardized posterior medians for α_{60} in the prior sensitivity analysis.

Appendix

Industrialization and political democracy model in matrix form

The structural model (for $i = 1, \dots, n$) is given in matrix form as:

$$\begin{pmatrix} \eta_i^{60} \\ \eta_i^{65} \end{pmatrix} = \begin{pmatrix} \alpha^{60} \\ \alpha^{65} \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ b_{21} & 0 \end{pmatrix} \begin{pmatrix} \eta_i^{60} \\ \eta_i^{65} \end{pmatrix} + \begin{pmatrix} \gamma^{60} \\ \gamma^{65} \end{pmatrix} \xi_i + \begin{pmatrix} \zeta_i^{60} \\ \zeta_i^{65} \end{pmatrix}$$

With ξ_i representing industrialization level in country i in 1960, and η_i^{60} and η_i^{65} representing political democracy in country i in 1960 and 1965, respectively. The parameters of interest in this model are the direct and indirect effect of industrialization in 1960 on political democracy in 1965, γ_{65} and $\gamma_{60} \cdot b_{21}$, respectively.

The measurement model for political democracy is given by:

$$\begin{pmatrix} y_{1i} \\ y_{2i} \\ y_{3i} \\ y_{4i} \\ y_{5i} \\ y_{6i} \\ y_{7i} \\ y_{8i} \end{pmatrix} = \begin{pmatrix} 0 \\ \nu_2^y \\ \nu_3^y \\ \nu_4^y \\ 0 \\ \nu_6^y \\ \nu_7^y \\ \nu_8^y \end{pmatrix} + \begin{pmatrix} 1 & 0 \\ \lambda_2^y & 0 \\ \lambda_3^y & 0 \\ \lambda_4^y & 0 \\ 0 & 1 \\ 0 & \lambda_6^y \\ 0 & \lambda_7^y \\ 0 & \lambda_8^y \end{pmatrix} \begin{pmatrix} \eta_i^{60} \\ \eta_i^{65} \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{pmatrix} \begin{pmatrix} D_i^{15} \\ D_i^{24} \\ D_i^{26} \\ D_i^{37} \\ D_i^{48} \\ D_i^{68} \end{pmatrix} + \begin{pmatrix} \epsilon_{1i}^y \\ \epsilon_{2i}^y \\ \epsilon_{3i}^y \\ \epsilon_{4i}^y \\ \epsilon_{5i}^y \\ \epsilon_{6i}^y \\ \epsilon_{7i}^y \\ \epsilon_{8i}^y \end{pmatrix}$$

With D representing a vector of pseudo-latent variables used to model the correlations between measurement errors in such a way that the covariance matrix Σ_y remains a diagonal matrix.

The measurement model for industrialization level is given by:

$$\begin{pmatrix} x_{1i} \\ x_{2i} \\ x_{3i} \end{pmatrix} = \begin{pmatrix} 0 \\ \nu_2^x \\ \nu_3^x \end{pmatrix} + \begin{pmatrix} 1 \\ \lambda_2^x \\ \lambda_3^x \end{pmatrix} \xi_i + \begin{pmatrix} \delta_{1i}^x \\ \delta_{2i}^x \\ \delta_{3i}^x \end{pmatrix}$$

References

- Arbuckle, J. L. (2013). IBM SPSS Amos 22 user's guide.
- Asparouhov, T., & Muthén, B. (n.d.). *Bayesian analysis using Mplus: Technical implementation*. Retrieved 18-04-2015, from <https://www.statmodel.com/download/Bayes3.pdf>
- 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(2), 151-164.
- Bayarri, M. J., & Berger, J. O. (2004). The interplay of Bayesian and frequentist analysis. *Statistical Science*, 19(1), 58-80.
- Berger, J. O. (2006). The case for objective bayesian analysis. *Bayesian Analysis*, 3, 385-402.
- Berger, J. O., & Strawderman, W. E. (1996). Choice of hierarchical priors: Admissibility in estimation of normal means. *The Annals of Statistics*, 24(3), 931-951.
- Berger, J. O., & Wolpert, R. L. (1984). *The likelihood principle*. Hayward, CA: Institute of Mathematical Statistics.
- Bollen, K. (1980). Issues in the comparative measurement of political democracy. *American Sociological Review*, 45(3), 370-390.
- Bollen, K. (1989). *Structural equations with latent variables*. New York: John Wiley and Sons.
- Can, S., van de Schoot, R., & Hox, J. (2014). Collinear latent variables in multilevel confirmatory factor analysis: A comparison of maximum likelihood and Bayesian estimations. *Educational and Psychological Measurement*, 75(3), 406-427.
- Carlin, B. P., & Louis, T. A. (2000a). *Bayes and empirical Bayes methods for data analysis* (2nd ed.). New York: Chapman and Hall/CRC.
- Carpenter, B., Gelman, A., Hoffman, M., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76(1).
- Casella, G. (1985). An introduction to empirical Bayes data analysis. *The American Statistician*, 39(2), 83-87.

- Chung, Y., Rabe-Hesketh, S., & Choi, I.-H. (2013). Avoiding zero between-study variance estimates in random-effects meta-analysis. *Statistics in Medicine*, 32(23), 4071-4089.
- Dagne, G. A., Howe, G. W., Brown, C. H., & Muthén, B. O. (2002). Hierarchical modeling of sequential behavioral data: An empirical Bayesian approach. *Psychological Methods*, 7(2), 262.
- Darnieder, W. F. (2011). *Bayesian methods for data-dependent priors* (Unpublished doctoral dissertation). The Ohio State University.
- Depaoli, S. (2012). Measurement and structural model class separation in mixture CFA: ML/EM versus MCMC. *Structural Equation Modeling: A Multidisciplinary Journal*, 19(2), 178–203.
- Depaoli, S. (2013). Mixture class recovery in GMM under varying degrees of class separation: Frequentist versus Bayesian estimation. *Psychological methods*, 18(2), 186-219.
- Depaoli, S. (2014). The impact of inaccurate informative priors for growth parameters in Bayesian growth mixture modeling. *Structural Equation Modeling: A Multidisciplinary Journal*, 21(2), 239-252.
- Depaoli, S., & Clifton, J. P. (2015). A bayesian approach to multilevel structural equation modeling with continuous and dichotomous outcomes. *Structural Equation Modeling: A Multidisciplinary Journal*, 1-25.
- Depaoli, S., & van de Schoot, R. (2015). Improving transparency and replication in Bayesian statistics: The WAMBS-checklist. *Psychological Methods*.
- Dunson, D. B., Palomo, J., & Bollen, K. (2005). *Bayesian structural equation modeling* (Tech. Rep.). SAMSI. Retrieved from <http://www.samsi.info/sites/default/files/tr2005-05.pdf>
- Efron, B. (1996). Empirical Bayes methods for combining likelihoods. *Journal of the American Statistical Association*, 91(434), 538-550.
- Garthwaite, P. H., Kadane, J. B., & O'Hagan, A. (2005). Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association*, 100(470),

680–701.

- Geldhof, G. J., Preacher, K. J., & Zyphur, M. J. (2014). Reliability estimation in a multilevel confirmatory factor analysis framework. *Psychological Methods*, 19(1), 72-91.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian Analysis*, 1(3), 515-534.
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. (2004). *Bayesian data analysis* (2nd ed.). London, England: Chapman and Hall/CRC.
- Gelman, A., Jakulin, A., Pittau, M. G., & Su, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2(4), 1360-1383.
- Gelman, A., Meng, X.-L., & Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, 6, 733-807.
- Gelman, A., & Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4), 457-472.
- Hallquist, M., & Wiley, J. (2014). MplusAutomation: Automating Mplus model estimation and interpretation [Computer software manual]. Retrieved from <http://CRAN.R-project.org/package=MplusAutomation> (R package version 0.6-3)
- Harring, J. R., Weiss, B. A., & Hsu, J.-C. (2012). A comparison of methods for estimating quadratic effects in nonlinear structural equation models. *Psychological methods*, 17(2), 193-214.
- Heerwegh, D. (2014). Small sample Bayesian factor analysis. In *Phuse*. Retrieved from <http://www.lexjansen.com/phuse/2014/sp/SP03.pdf>
- Hobert, J. P., & Casella, G. (1996). The effect of improper priors on Gibbs sampling in hierarchical linear mixed models. *Journal of the American Statistical Association*, 91(436), 1461-1473.
- Hox, J., van de Schoot, R., & Matthijsse, S. (2012). How few countries will do? Comparative

- survey analysis from a Bayesian perspective. *Survey Research Methods*, 6(2), 87-93.
- Kaplan, D. (2014). *Bayesian statistics for the social sciences*. Guilford Publications.
- Kass, R. E., & Natarajan, R. (2006). A default conjugate prior for variance components in generalized linear mixed models (comment on article by Browne and Draper). *Bayesian Analysis*, 1(3), 535-542.
- Kass, R. E., & Wasserman, L. (1996). The selection of prior distributions by formal rules. *Journal of the American Statistical Association*, 91(435), 1343-1370.
- Kohli, N., Hughes, J., Wang, C., Zopluoglu, C., & Davison, M. L. (2015). Fitting a linear-linear piecewise growth mixture model with unknown knots: A comparison of two common approaches to inference. *Psychological methods*, 20(2), 259-275.
- Lambert, P. C., Sutton, A. J., Burton, P. R., Abrams, K. R., & Jones, D. R. (2005). How vague is vague? a simulation study of the impact of the use of vague prior distributions in MCMC using WinBUGS. *Statistics in Medicine*, 24(15), 2401-2428.
- Leamer, E. E. (1983). Let's take the con out of econometrics. *The American Economic Review*, 73(1), 31-43.
- 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(4), 653-686.
- 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(4), 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(1), 101-119.
- Lunn, D. J., Thomas, A., Best, N., & Spiegelhalter, D. (2000). WinBUGS - a Bayesian modelling framework: Concepts, structure, and extensibility. *Statistics and Computing*, 10, 325-337.
- Maas, C. J., & Hox, J. J. (2005). Sufficient sample sizes for multilevel modeling. *Methodology*,

I(3), 86–92.

MacCallum, R. C., Edwards, M. C., & Cai, L. (2012). Hopes and cautions in implementing Bayesian structural equation modeling. *Psychological Methods*, 17(3), 340-345.

Merkle, E. C., & Rosseel, Y. (2015). blavaan: Bayesian structural equation models via parameter expansion. *arXiv preprint arXiv:1511.05604*.

Meuleman, B., & Billiet, J. (2009). A Monte Carlo sample size study: How many countries are needed for accurate multilevel SEM? *Survey Research Methods*, 3(1), 45-58.

Mulder, J., & Fox, J.-P. (2013). Bayesian tests on components of the compound symmetry covariance matrix. *Statistics and Computing*, 23, 109-122.

Mulder, J., Hoijtink, H., & Klugkist, I. (2010). Equality and inequality constrained multivariate linear models: Objective model selection using constrained posterior priors. *Journal of Statistical Planning and Inference*, 140, 887-906.

Mulder, J., Klugkist, I., van de Schoot, R., Meeus, W. H. J., Selfhout, M., & Hoijtink, H. (2009). Bayesian model selection of informative hypotheses for repeated measurements. *Journal of Mathematical Psychology*, 53, 530-546.

Muthén, B. O., & Asparouhov, T. (2012). Bayesian structural equation modeling: A more flexible representation of substantive theory. *Psychological Methods*, 17(3), 313-335.

Muthén, L. K., & Muthén, B. O. (1998-2012). *Mplus user's guide. seventh edition*. Los Angeles, CA: Muthén and Muthén.

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: A Multidisciplinary Journal*, 9(4), 599-620.

Natarajan, R., & Kass, R. E. (2000). Reference Bayesian methods for generalized linear mixed models. *Journal of the American Statistical Association*, 95(449), 227-237.

Oberski, D. L. (2014). Evaluating sensitivity of parameters of interest to measurement invariance in latent variable models. *Political Analysis*, 22(1), 45-60.

Oravecz, Z., Tuerlinckx, F., & Vandekerckhove, J. (2011). A hierarchical latent stochastic

- differential equation model for affective dynamics. *Psychological methods*, 16(4), 468-490.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd international workshop on distributed statistical computing* (Vol. 124).
- Polson, N. G., & Scott, J. G. (2012). On the half-Cauchy prior for a global scale parameter. *Bayesian Analysis*, 7(4), 887-902.
- R Core Team. (2015). R: A language and environment for statistical computing. [Computer software manual]. R Foundation for Statistical Computing, Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Raftery, A. E., Lewis, S., et al. (1992). How many iterations in the gibbs sampler. *Bayesian statistics*, 4(2), 763–773.
- Robert, C. P. (2007). *The Bayesian choice: From decision-theoretic foundations to computational implementations*. New York, NY: Springer Science+Business Media.
- Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2), 1-36. Retrieved from <http://www.jstatsoft.org/v48/i02/>
- Ryu, E., & West, S. G. (2009). Level-specific evaluation of model fit in multilevel structural equation modeling. *Structural Equation Modeling*, 16(4), 583–601.
- Saris, W. E., Satorra, A., & van der Veld, W. M. (2009). Testing structural equation models or detection of misspecifications? *Structural Equation Modeling*, 16(4), 561-582.
- Tversky, A. (1974). Assessing uncertainty. *Journal of the Royal Statistical Society. Series B (Methodological)*, 148–159.
- van de Schoot, R., Winter, S., Ryan, O., Zondervan-Zwijnenburg, M., & Depaoli, S. (2016). A systematic review of Bayesian papers in psychology: The last 25 years. *Psychological Methods*, 4(21).
- Yuan, Y., & MacKinnon, D. P. (2009). Bayesian mediation analysis. *Psychological methods*, 14(4), 301-322.
- Zhang, Z., Lai, K., Lu, Z., & Tong, X. (2013). Bayesian inference and application of robust

growth curve models using Student's t distribution. *Structural Equation Modeling: A Multidisciplinary Journal*, 20(1), 47-78.