Classical and Bayesian Uncertainty Intervals for the Reliability of Multidimensional Scales

Julius M. Pfadt¹, Don van den Bergh², Morten Moshagen¹

¹ *Department of Psychological Research Methods, Ulm University, Germany*

² *Department of Psychology, University of Amsterdam, The Netherlands*

Word count: 11,577

Disclosure statement: The authors report there are no competing interests to declare.

Data availability: The data are available at an open science repository for this manuscript.

Data deposition: All associated files for this manuscript are available at <https://osf.io/xyq4n/>.

Acknowledgements: The authors thank Eric-Jan Wagenmakers for helpful comments and feedback on this manuscript.

Correspondence: Don van den Bergh, University of Amsterdam, Department of Psychological Methods. Postbus 15906, 1001 NK Amsterdam, The Netherlands. E-Mail should be sent to: donvdbergh@hotmail.com.

Classical and Bayesian Uncertainty Estimates for the Reliability of Multidimensional Scales

Abstract

The reliability of a multidimensional test instrument is commonly estimated using coefficients ω_t (total) and ω_h (hierarchical). However, point estimates for the coefficients are rarely accompanied by uncertainty estimates. In this study, we compare bootstrap and normal-theory confidence intervals. In addition, we develop Bayesian versions of coefficients ω_t and ω_h by sampling from a second-order factor model. Results from a comprehensive simulation study show that the studied confidence intervals performed well when the sample size was sufficiently large ($N \geq 500$). The Bayesian estimates performed well across most studied conditions. When the sample size was small and the reliability low, only the bias-corrected and accelerated bootstrap confidence interval approached a satisfactory coverage among all intervals. This study guides on ω_t and ω_h confidence intervals and introduces ω_t and ω_h credible intervals that are easy to use and come with the benefits of Bayesian parameter estimation.

Keywords: Coefficient omega total/hierarchical, confirmatory factor analysis, bootstrap, second-order factor model

Many constructs in the social and behavioral sciences assume a multidimensional structure with a general factor along with more specific subdimensions called group factors. For example, the short version of the UPPS-P (Cyders et al., 2014) is a measurement of impulsivity and comprises 20 items. The subscales (negative and positive urgency, (lack of) premeditation, (lack of) perseverance, sensation seeking) are measured by four items each, so that the more general impulsivity domain is partitioned into five more specific subdomains. Depending on the research question and design, the reliabilities of approximate unidimensional subsets of the composite test score may be of interest. For instance, one may divide the impulsivity scale score into unidimensional subscale scores of four items each and estimate the reliability for each scale using, e.g., coefficient α (Cronbach, 1951). However, this approach is unsuited when either the total or the general factor reliability of the composite test score (all items) is of interest. Likewise, estimating coefficient α for the complete impulsivity test score underestimates the true reliability (Sijtsma, 2009). Proper coefficients for the reliability of a multidimensional scale are the coefficients ω_t , indicating the total reliability of a scale, and ω_h , indicating the reliability of the general factor, that is, the extent to which all items of a scale measure a common underlying latent variable (McDonald, 1999;

Zinbarg et al., 2005).¹

When reporting reliability coefficients, researchers usually present a point estimate without an uncertainty estimate such as a confidence interval. This practice disregards the sampling error when estimating a reliability parameter from a finite sample. To quantify uncertainty, confidence intervals (in a frequentist setting) or credible intervals based on posterior distributions (in a Bayesian setting) can be used. Here, we investigate several approaches to obtain confidence intervals for coefficients ω_t and ω_h and describe a method to estimate a second-order factor model in the Bayesian framework based on work by S.-Y. Lee (2007) that can be employed to obtain Bayesian versions of ω_t and ω_h and their posterior distributions.

To our knowledge, confidence intervals for coefficients ω_t and ω_h , both analytic and bootstrapped ones, have not been investigated before.² Furthermore, we are unaware of any attempts concerning Bayesian versions of coefficients ω_t and ω_h . This makes our work novel in several ways. We are the first to investigate various confidence intervals for coefficients ω_t and ω_h . Also, we are the first to develop and evaluate Bayesian versions of ω_t and ω_h by adopting methodology from Bayesian structural equation modeling (SEM). Finally, the implementation of the developed methods into an R-package is another novel contribution that increases the impact of this work.

The outline of this paper is as follows. First, we present the factor model approach to reliability and describe how the coefficients ω_t and ω_h can be obtained from the parameters of the second-order factor model. Second, we detail how the second-order factor model can be estimated in the Bayesian framework, which yields the parameters needed to compute the posterior distributions of ω_t and ω_h . Third, we discuss uncertainty estimation in reliability analysis. Fourth, we compare the frequentist and Bayesian reliability estimates in a simulation study before showcasing a (Bayesian) reliability analysis with the impulsivity example. Finally, we discuss the results of the simulation study, practical implications, and future directions of research.

The factor model approach to reliability

Historically, reliability has been defined in the realm of classical test theory (CTT). In CTT, the test score is divided into a true score and an error score: $X = T + E$. The reliability is the relative amount of the test score variance that is true score variance which equals the correlation between parallel test administrations (Lord & Novick, 1968). Parallel tests can be described as administering the same test at least twice to the same sample of participants with them having no recollection of the previous administration.

¹The greatest lower bound (glb; Woodhouse & Jackson, 1977) is also a reliability coefficient that can be estimated for multidimensional scales. However, the glb suffers from capitalization on chance, which results in a severe positive bias of the coefficient (Pfadt, van den Bergh, et al., 2022; Sijtsma, 2009; Ten Berge & Sočan, 2004), so we do not consider the glb further.

²Kelley and Pornprasertmanit (2016) investigated confidence intervals for a hierarchical coefficient ω , however, the coefficient in their study is conceptually different from the ω_h -coefficient researched in this work.

Reliability $\rho_{XX'}$ can be written as:

$$\rho_{XX'} = \frac{\sigma_T^2}{\sigma_X^2} = 1 - \frac{\sigma_E^2}{\sigma_X^2}, \quad (1)$$

with σ_X^2 as the test score variance, σ_T^2 and σ_E^2 as the true score variance and the error score variance, respectively (Lord & Novick, 1968). Note that CTT makes no assumptions about the dimensionality of the data (e.g., Sijtsma & Pfadt, 2021).

Coefficients ω_t and ω_h are based on a factor analytic (FA) approach to decomposing a test score and inferring reliability. In a typical factor model, for example, the single-factor model (also unidimensional/congeneric model; Spearman, 1904), one divides the test score into a part explained by one common latent variable or factor, and a remaining error part that cannot be explained by the common factor: $X = \lambda F + E$, with λ as the loadings of the item scores that comprise X on factor F . We notice that in contrast to CTT, the definition of reliability in the FA approach changes as the definition of the true score variance changes. Reliability is then defined as the relative amount of test score variance that is due to one or more factors (e.g., Bollen, 1989; Jöreskog, 1971; McDonald, 1970). In the single-factor model, there is only one common factor that explains the responses in the test, thus the true score variance is the variance explained by the common factor. A single-factor model is at odds with assuming a multidimensional construct that incorporates a general factor along with group factors and thereby only provides a rather crude approximation of the general factor. Considering the impulsivity-example with its five dimensions of impulsivity, the variance due to only a single factor falls short of capturing the true score variance explained by all facets of impulsivity; hence, a more sophisticated factor model is appropriate.

Second-order factor model

We assume that the test score can be explained by several primary latent variables (group factors) that are each common to some but not all items and the primary latent variables are themselves explained by one second-order general factor. This model is also sometimes referred to as a “higher-order factor model”. For instance, one can structure the impulsivity-scale into a general second-order factor, that is, the overall impulsivity that influences the outcomes of the manifest variables through the group factors (or dimensions), specifically the subscales negative and positive urgency, premeditation, perseverance, and sensation seeking.

The second-order factor model allows decomposing the test score into convenient parts (see Figure 1 for an exemplary drawing of a second-order factor model of the impulsivity-example). When we assume data $\mathbf{Y}$ with dimensions $n \times k$ (observations $\times$ items) to be multivariate normal and with mean zero for all $i = 1, \dots, n$, the second-order model with two levels and a single common factor at the second level is defined as:

$$\mathbf{y}_i = \mathbf{\Lambda} \boldsymbol{\eta}_i + \boldsymbol{\epsilon}_i, \quad (2)$$

Figure 1

Impulsivity Example: Second-Order Factor Model With Five Group Factors

$$\eta_i = \beta \xi_i + \delta_i. \quad (3)$$

Vector $\mathbf{y}_i$ ($k \times 1$) contains the scores of person i on k items. The loading matrix $\mathbf{\Lambda}$ ($k \times q$) contains the primary loadings of the manifest variables on the group factors; q denotes the number of group factors common to some but not all items. In the primary loadings matrix $\mathbf{\Lambda}$, we fix the loadings of manifest variables on “outgroup” factors (factors they are not associated with) to 0. The combined residuals ϵ ($k \times 1$) have a diagonal covariance matrix $\mathbf{\Psi}_\epsilon$ ($k \times k$) containing k residual variances ψ_ϵ ($k \times 1$). The vector $\boldsymbol{\eta}_i$ ($q \times 1$) contains the group factor scores of each person. On the second level, the factor scores $\boldsymbol{\eta}_i$ ($q \times 1$) are explained by a loading vector $\boldsymbol{\beta}$ ($q \times 1$), one exogenous common factor ξ_i (1×1), and the error vector $\boldsymbol{\delta}_i$ ($q \times 1$). Combined, $\boldsymbol{\delta}$ has a diagonal covariance matrix $\mathbf{\Psi}_\delta$ ($q \times q$) containing the residual variances ψ_δ ($q \times 1$) of the group factors.

To compute coefficients ω_t and ω_h , the loadings of the manifest variables on the general factor are required, which are only indirectly obtained in the second-order model. It is thus useful to transform the second-order model into a bi-factor structure instead (see Figure S1 in the supplementary materials for an exemplary drawing of a bi-order factor model for the impulsivity scale.) The bi-factor model is a more general factor model that relates group factors to items similar to the second-order model but relates the general factor directly to the manifest variables and not to the group factors. The second-order factor model is nested in the bi-factor model, and the models

are equivalent when the proportionality condition is met, that is when the ratios of variance explained by the general factor to the variance explained by the group factor are identical for all indicators of a group factor (Bader & Moshagen, in press; Mansolf & Reise, 2017; Mulaik & Quartetti, 1997).

Through a Schmid-Leiman transformation, the loadings of a second-order factor model can be transformed to yield the loadings of a corresponding (and thus equivalent) bi-factor model (Schmid & Leiman, 1957). The transformation works as follows: The loadings of the manifest variables on the general factor, γ ($k \times 1$), are obtained by equation $\gamma = \Lambda\beta$; the loadings of the items on the group factors in a bi-factor model, Π ($k \times q$), equal $\Pi = \Lambda\sqrt{\Psi_\delta}$; the residual variances of the manifest variables in the bi-factor model equal the residual variances in the second-order model, ψ_ϵ . Now, one may compute:

$$\omega_t = \frac{\mathbf{1}'\gamma\gamma'\mathbf{1} + \mathbf{1}'\Pi\Pi'\mathbf{1}}{\mathbf{1}'\gamma\gamma'\mathbf{1} + \mathbf{1}'\Pi\Pi'\mathbf{1} + \mathbf{1}'\Psi_\epsilon\Psi_\epsilon'\mathbf{1}} \quad (4)$$

and

$$\omega_h = \frac{\mathbf{1}'\gamma\gamma'\mathbf{1}}{\mathbf{1}'\gamma\gamma'\mathbf{1} + \mathbf{1}'\Pi\Pi'\mathbf{1} + \mathbf{1}'\Psi_\epsilon\Psi_\epsilon'\mathbf{1}} \quad (5)$$

In the work of Zinbarg et al. (2005), the denominator in (4) and (5) equals the variance of the test score, σ_X^2 . We assume that the sum of all the variance parts in the denominator of Equations 4 and 5 will equal the total variance of the test score when the model fits in terms of the likelihood ratio test, which represents a strictly model-based definition. When the model fit is poor, one should refrain from using the coefficients at all (McDonald, 1999).

The choice of whether coefficient ω_t or coefficient ω_h is the proper reliability coefficient depends on whether one considers the variance due to the group factors as part of the true score variance or the error score variance. Following the reasoning of Zinbarg et al. (2005) and assuming that variance due to group factors is systematic variance and thus true score variance, one arrives at coefficient ω_t as an estimate for the total reliability. If one is only interested in the general factor variance as the true score variance one considers ω_h as a measure of the general factor reliability. Coefficient ω_h then also informs about the extent to which the test scores can be explained by a common factor in the presence of group factors, that is, the general factor saturation (e.g., Rodriguez et al., 2016).

Zinbarg et al. (2005) summarize some general properties of ω_t and ω_h : Coefficient ω_t equals the true total reliability of the composite test score when no item specific factors are present in the data; when item-specific factors are present, the coefficient is an underestimate of the reliability. Coefficient ω_t equals ω_h when the data fit a unidimensional model; when the data fit a multidimensional model ω_h will be smaller than ω_t . Reanalyzing previous studies on the performance of various reliability

coefficients, Cho (2022) found coefficient ω_t (named ρ_{SOF}) based on a second-order factor model to perform well with multidimensional data.

The Bayesian approach to coefficients ω_t and ω_h

The goal of reliability estimation in the Bayesian framework is to obtain the posterior distributions of coefficients ω_t and ω_h . As a probability distribution, the posterior represents the relative probabilities for all possible parameter values after the data have been observed. One obtains the posterior distribution by updating the prior distribution of the parameter with the knowledge gained from the observed data. The prior distribution represents the probabilities for different ranges of parameter values before the data have been observed. Thus, the posterior distribution of, for example, coefficient ω_t contains the probabilities for all possible ranges of values of ω_t and encapsulates the precision with which ω_t is estimated.

The prior distribution plays an important role as it quantifies the prior confidence in different values of a parameter and thereby affects the parameter retrieval in dynamic ways. The posterior distribution is obtained by updating the prior with the data. If the evidence in the data is not very convincing, for example, because there is a lot of noise or because the sample size is relatively small, then the prior distribution has more influence on the posterior distribution. For instance, in a Bayesian reliability analysis, we could specify prior distributions that imply a prior for ω_t that resembles a normal distribution with a peak around .5. We might draw a data sample of $n = 50$ yielding a frequentist ω_t estimate of .85. The Bayesian point estimate for ω_t would likely be somewhere between .5 and .8 as the evidence in the data is too weak to cause a major shift of the prior confidence and thus away from .5. The associated credible interval would be considerably wide reflecting the large uncertainty in the estimation with such a small sample. We do not regard such a scenario as a disadvantage of Bayesian parameter estimation, rather it encourages researchers to 1) focus on the width of the posterior distribution and the corresponding uncertainty interval instead of how close the point estimate is to a cutoff value, 2) contemplate the confidence they have in different reliability values before seeing the data, and 3) to collect more data to improve the precision of the reliability estimation.

The prior distributions of coefficients ω_t and ω_h are not directly specified as the coefficients are functions of loadings and residual variances. Instead, one specifies the prior distributions of the second-order factor model parameters to control the priors of ω_t and ω_h . The factor model prior distributions are then updated with the data by Markov chain Monte Carlo sampling (MCMC; e.g., Gilks et al., 1995), more specifically Gibbs sampling (Geman & Geman, 1984). One draws samples from the conditional posterior distributions of the factor model parameters that are needed to compute coefficients ω_t and ω_h . In this procedure, we follow S.-Y. Lee (2007, Chapter 4), who provides the prior and conditional posterior distributions for a Gibbs Sampler of a general structural equation model. At the level of the measurement model, we adopted

the standard distributions for Bayesian SEM (S.-Y. Lee, 2007; Muthén & Asparouhov, 2012); at the level of the structural model, we simplified the equations, since there is only a single secondary factor predicting the primary factors instead of multiple possibly hierarchically structured latent variables. Specifically, the loadings of the primary group factors on the general factor reduce to a vector instead of a matrix representation for general SEM (see S.-Y. Lee, 2007); the factor scores of the secondary factor appear in vector form instead of a matrix; and the covariance matrix of the secondary factor scores simplifies to a single variance value.

We draw starting values for the Gibbs sampling from the prior distributions of the second-order model parameters. Following standard Bayesian SEM procedures (S.-Y. Lee, 2007; Muthén & Asparouhov, 2012), we adopt prior distributions for ψ_ϵ , Λ , ψ_δ , β , Φ , η , and ξ . Let n be the sample size, k be the number of manifest variables, and q the number of group factors. The prior distribution of the residual variances of the manifest variables is the inverse gamma distribution, $p(\psi_\epsilon) = IG(a_0, b_0)$ with shape and scale parameters. The prior distribution of the loadings of the manifest variables on the group factors, the non-zero elements in loading matrix Λ , is a normal distribution with mean λ_0 and variance $A_0 \psi_\epsilon$, $p(\Lambda | \psi_\epsilon) = N(\lambda_0, A_0 \psi_\epsilon)$. The prior distribution of the latent residual variances is an inverse gamma, $p(\psi_\delta) = IG(c_0, d_0)$, and the prior of the loadings of the group factors on the general factor is a normal distribution, $p(\beta | \psi_\delta) = N(\beta_0, B_0 \psi_\delta)$. We combine the scores of the primary and second-order factors, η , and ξ , in Ω since they have a joint prior distribution, $p(\Omega) = N(0, \Phi)$. The diagonal covariance matrix of the combined factor scores, Φ , is composed of the general factor variance ϕ and the residual variances ψ_δ of the group factors. The prior distribution of the general factor variance is an inverse gamma, $p(\phi) = IG(p_0/2, R_0/2)$.³

The prior distributions of the factor model parameters are conjugate, which means that the conditional posterior distributions are in the same family of distributions as the prior distributions. We insert the drawn starting values from the prior distributions of the factor model parameters into the conditional posterior distributions of the factor model parameters (cf. S.-Y. Lee, 2007):

$$p(\psi_\epsilon | \Omega) = IG(n/2 + a_0, \mathbf{b}), \quad (6)$$

$$p(\Lambda | \Omega, \psi_\epsilon) = N(\mathbf{a}, \psi_\epsilon \mathbf{A}), \quad (7)$$

with $\mathbf{A} = (A_0^{-1} + \Omega_1' \Omega_1)^{-1}$; Ω_1 is the subset of factor scores due to the group factors; $\mathbf{a} = \mathbf{A}(\Omega_1 \mathbf{Y})$; and $\mathbf{b} = \text{diag}(\frac{b_0}{2} (\mathbf{Y}'\mathbf{Y} - \mathbf{a}'\mathbf{A}^{-1}\mathbf{a}))$.

$$p(\psi_\delta | \Omega) = IG(n/2 + c_0, \mathbf{d}), \quad (8)$$

$$p(\beta | \Omega, \psi_\delta) = N(\mathbf{c}\sqrt{\phi}, \psi_\delta \mathbf{C}), \quad (9)$$

³The inverse gamma prior for the factor variance, $p(\phi) = IG(p_0/2, R_0/2)$, is proportional to the inverse Wishart prior described in S.-Y. Lee (2007, Chapter 4) when the number of exogenous factors equals one.

with $C = (B_0^{-1} + \mathbf{\Omega}'_2 \mathbf{\Omega}_2)^{-1}$, $\mathbf{\Omega}_2$ is the subset of factor scores due to the general factor; $c = C(\mathbf{\Omega}_2 \mathbf{\Omega}_1)$; $d = \text{diag}(\frac{d_0}{2} (\mathbf{\Omega}'_1 \mathbf{\Omega}_1 - c' C^{-1} c))$; ϕ is the variance of the general factor.

$$p(\mathbf{\Omega} | Y, \psi_\epsilon, \psi_\delta, \mathbf{\Lambda}, \beta) = N \left[(\mathbf{\Gamma}^{-1} + \mathbf{\Sigma})^{-1} \mathbf{\Lambda}_0 \mathbf{\Psi}_\epsilon^{-1} Y, (\mathbf{\Gamma}^{-1} + \mathbf{\Sigma})^{-1} \right] \quad (10)$$

$$p(\phi | Y, \mathbf{\Omega}) = IG \left(\frac{n + p_0}{2}, \frac{\mathbf{\Omega}'_2 \mathbf{\Omega}_2 R_0^{-1}}{2} \right). \quad (11)$$

Matrix $\mathbf{\Gamma} = (\mathbf{I} - \mathbf{B})^{-1} \mathbf{\Psi}_\delta^+ (\mathbf{I} - \mathbf{B})'^{-1}$, $\mathbf{I}$ is the identity matrix. Matrix $\mathbf{B}$ is a $(q + 1)$ -dimensional square matrix with only zeros except for the first column, which contains $(0, \beta)$ as entries. Matrix $\mathbf{\Psi}_\delta^+$ is a $(q + 1)$ -dimensional diagonal matrix with the general factor variance, ϕ , set to 1, and the group factor residual variances, ψ_δ , as entries; and $\mathbf{\Sigma} = \mathbf{\Lambda}'_0 \mathbf{\Psi}_\epsilon^{-1} \mathbf{\Lambda}_0$, where $\mathbf{\Lambda}_0$ equals $\mathbf{\Lambda}$ with an added column of zeros at the beginning. We normalize the factors by dividing the factor scores in $\mathbf{\Omega}$ by their standard deviation in every iteration of the Gibbs sampler to ensure identification.

In the first iteration of the Gibbs sampler, we use the drawn starting values from the prior distributions of the factor model parameters to successively draw samples from the conditional posterior distributions of the various factor model parameters (see Equations 6 to 11). Once we obtained samples from each of the conditional posterior distributions in the first iteration, we plug the just drawn values into the second iteration of the sampler. We repeat that process until sufficient values of each factor model parameter have been drawn, for example, 10,000 times. We restart the Gibbs sampler with different starting values, to assess convergence to the stationary distribution (the posterior). The limiting distribution of samples obtained from an adequately constructed Gibbs sampler is the target distribution (i.e., the posterior distribution). The samples sufficiently approximate the target distribution if the Markov chains converged. Convergence is assessed by checking traceplots and computing the $\hat{R}$ -statistic, which compares the within-chain variance to the between-chain variance (Gelman et al., 2014, Chapter 11.5).

In Gibbs sampling, there can be considerable autocorrelation within a posterior sample of a factor model parameter. Autocorrelation can be reduced by *thinning* the posterior samples. For instance, a thinning interval of 4 means only every 4th value of the posterior sample in one chain is used. We can check the level of autocorrelation by computing the effective sample size (ESS). The ESS indicates the number of independent posterior samples that contain as much information as the complete set of autocorrelated posterior samples (Geyer, 2011; Stan Development Team, 2022).

Once the MCMC samples converged to the posterior distributions of the loadings $\mathbf{\Lambda}$, β , and residual variances ψ_ϵ , ψ_δ , we use the Schmid-Leiman transformation to obtain posterior samples of the parameters γ and $\mathbf{\Pi}$. From there, we obtain posterior samples of coefficients ω_t and ω_h using Equations 4 and 5.

Uncertainty estimation

When one estimates a parameter in statistical analysis, it is good practice to account for the uncertainty of that parameter in the form of a standard error and/or an uncertainty interval. An interval adds an area of uncertainty to a point estimate to account for sampling error and informs about the underlying population parameter. Although one estimates the reliability of a test to account for measurement error, the reliability estimate, i.e., coefficient ω_t , is affected by error itself and thus necessitates an uncertainty estimate.

Confidence intervals for ω_t and ω_h

The most popular candidate for an uncertainty interval is the frequentist confidence interval, or confidence interval for short. Generating a confidence interval follows a general procedure: obtain a measure of error from the sampling distribution of a parameter and add that error to the point estimate of the parameter to create an interval. Contrary to the common misconception that one can be 95% confident that the 95% confidence interval contains the true parameter value, i.e., ω_t (e.g., Cumming, 2014), a 95% confidence interval does, in fact, cover the underlying population parameter in 95% of the cases when one would repeat the process of sampling and computing the 95% confidence interval for ω_t numerous times (Morey et al., 2016; Neyman, 1937).

A complication concerning the uncertainty estimation of ω_t and ω_h is that obtaining their limiting distribution is unfeasible, so approximations are required. In the present study, we examine six different confidence intervals: five bootstrap confidence intervals, and one Wald-type confidence interval. The procedure to obtain a bootstrap confidence interval for coefficients ω_t and ω_h is straightforward. One resamples the data of size n , for example, 1,000 times, each yielding n observations, and computes the coefficients ω_t and ω_h for each of the 1,000 resampled data sets using an FA method, that is, in this study a confirmatory factor analysis (CFA). Thus, one obtains an empirical sampling distribution for each reliability coefficient (Efron, 1979). The empirical sampling distribution yields a measure of uncertainty, for example, a standard error, to construct an interval around the point estimate of ω_t or ω_h .

We consider five intervals that are constructed this way. First, for the standard error interval (SE), one constructs an interval around the point estimate of ω_t (or ω_h) by adding the z-value from a standard normal distribution multiplied by the standard error of ω_t (or ω_h). The standard error of ω_t (or ω_h) equals the standard deviation of the empirical sampling distribution. Second, one obtains the standard error bias corrected interval (SE_{bias}) in the same way as the SE -interval, except that one incorporates the difference between the mean of the empirical sampling distribution and the point estimate into the formation of the interval to correct for possible bias in the empirical sampling distribution (for details, see Davison & Hinkley, 1997, Chapter 5: *normal approximation*).⁴ Third, to obtain the standard error logistic transformed interval (SE_{ln}), one generates

⁴In the R-package *boot* (Canty & Ripley, 2021) the SE_{bias} interval equals *norm*-interval.

an interval by applying the natural logistic transformation to the point estimate of ω_t (or ω_h) as well as to the standard error as obtained for the SE interval. The interval is then transformed back to the normal scale to form the SE_{ln} confidence interval (for details, see Kelley & Pornprasertmanit, 2016). Contrary to the SE and SE_{bias} intervals, the SE_{ln} interval is not symmetric. Fourth, for the percentile bootstrap interval ($Perc$), one computes the lower and upper quantiles of the empirical sampling distribution of ω_t (or ω_h). Fifth, one obtains the bias-corrected and accelerated bootstrap interval (BCA) by accounting for the bias and skewness of the empirical sampling distribution by applying the jackknife-resampling technique (see DiCiccio & Efron, 1996). We chose these specific intervals because of their popularity in common R-packages, *boot* (Canty & Ripley, 2021) and *lavaan* (Rosseel, 2012), and their promising performance in Kelley and Pornprasertmanit (2016) and Padilla and Divers (2016).

For the Wald-type interval (*Wald*), one fits the second-order factor model in a CFA to multidimensional data. In that process, one can define coefficients ω_t and ω_h as to be estimated parameters in the maximum likelihood estimation and thereby obtain standard errors for the coefficients. These one can use to construct Wald-type confidence intervals.

In simulation studies for coefficient ω_u , the ω -version for unidimensional data, the SE , SE_{ln} , and BCA confidence intervals displayed good coverage performance; the $Perc$ interval performed slightly worse but still satisfactory (Kelley & Pornprasertmanit, 2016; Padilla & Divers, 2016). Kelley and Pornprasertmanit (2016) also investigated the performance of the SE , SE_{ln} , $Perc$, and BCA intervals for coefficient ω_h and found them to perform well across various conditions. Note that Kelley and Pornprasertmanit computed ω_h by fitting a unidimensional factor model to obtain single-factor loadings and dividing the squared sum of those loadings by the test score variance, σ_X^2 . However, with multidimensional data, coefficient ω_h from Kelley and Pornprasertmanit (2016) will not be equivalent to coefficient ω_h as defined in this study or in Zinbarg et al. (2005). To date, we are unaware of any simulation studies that examined confidence intervals for coefficient ω_t or Zinbarg et al.'s version of coefficient ω_h .

Credible intervals for ω_t and ω_h

The credible interval for either ω_t or ω_h is obtained from their respective posterior samples. We compute the X% credible intervals as the highest posterior density region (Box & Tiao, 1973) that contains X% of the posterior mass of ω_t or ω_h . The resulting intervals contain ω_t or ω_h with X% probability. To date, we are unaware of any research on credible intervals for ω_t and ω_h .

To summarize, coefficients ω_t and ω_h are representatives of the FA approach to reliability and define the true score variance of a composite test score in different ways. Coefficient ω_t includes the true score variance that is due to one general and multiple group factors; coefficient ω_h only includes the true score variance that is due to a general factor. Although one may estimate both coefficients from familiar structural

equation models, estimates for their uncertainty have been scarcely researched, let alone reported. Thus, we examine the statistical performance of six confidence intervals as well as credible intervals for ω_t and ω_h in a simulation study. The knowledge gained from the simulation study shall provide researchers with the necessary tools to report uncertainty measures in reliability analysis. We provide the Bayesian routines, as well as the Wald-type confidence interval in the R-package *Bayesrel* (Pfadt, Van den Bergh, et al., 2022).

Simulation study

To compare several confidence intervals and credible intervals for coefficients ω_t and ω_h , we conducted a simulation study with multiple conditions. In the study, we relate the coefficients to their population values. Although we did not control population reliability per se, we note that the population value of ω_t equals population reliability when one adopts the FA approach to reliability and excludes item-specific factors.

Priors

The choice of the prior hyperparameters is important in Bayesian structural equation modeling. The hyperparameters define the shape of the prior distributions of the factor model parameters. The default prior hyperparameters of software solutions for Bayesian SEM have in common that they include very large variances: In Mplus, the prior for the factor loadings is $N(0, 10^{10})$, and the prior for the residual variances is $IG(-1, 0)$ (Asparouhov & Muthén, 2021); in blavaan, an R-package for Bayesian SEM, the prior loadings are normally distributed with $N(0, 10)$, and the prior residual variances follow an inverse gamma distribution, $IG(1, 0.5)$ (Merkle & Rosseel, 2018). Both the priors in Mplus and blavaan yield very large variances, which implies highly uninformative priors. For our Gibbs sampler, prior distributions with such large variances are inefficient. For the simulation study, we aimed to define the hyperparameters so that the prior distributions of the factor model parameters are relatively uninformative. Relatively uninformative priors compromise between efficient sampling and exploring the parameter space in its entirety. Particularly, the priors for the factor model parameters include a variance that is large enough so all plausible values of the parameters are likely to be sampled, but small enough so that reaching the target distribution does not take days. In the following, we describe the reasoning for the hyperparameters of the relatively uninformative setting, which are the basis of the main simulation results.

For a relatively uninformative prior on the residual variances, S.-Y. Lee (2007) recommends an inverse gamma distribution with shape and scale parameters as $IG(6, 10)$, resulting in mean residual variances of 1 with a variance of 2. For the residual variances of the manifest and latent variables, we chose prior hyperparameters $a_0 = 2$, $b_0 = 1$, $c_0 = 2$, and $d_0 = 1$ so that the prior residual variances have a slightly smaller

mean but much larger variances than Lee's recommendation.⁵ For the prior loadings, we defined hyperparameters $\lambda_0 = 0$, $A_0 = 1$, and $\beta_0 = 0$, $B_0 = 2.5$, resulting in loadings, both manifest and latent, with a mean of 0, and variance determined by the residual variances multiplied by a scalar, A_0 and B_0 , respectively. We chose that multiplier to be higher for the latent loadings since the uncertainty at higher levels of the model is larger. For the hyperparameters of the general factor variance, we chose $p_0 = q^2 - q$, and $R_0 = k$. Thus, with simulation conditions of $q = 3/k = 9$ and $q = 5/k = 30$, the prior hyperparameters for the general factor variance yield an inverse gamma prior with mean = 2.25 and variance = $5\frac{1}{16}$ for three group factors, and an inverse gamma with mean = $1\frac{2}{3}$ and variance = $\frac{25}{72}$. This way, we ensure that the prior variance of the general factor scales down when the number of group factors grows. When the number of group factors is small we assume that the posterior variance of the g-factor is smaller than with more group factors.

Since the choice of prior can affect the parameter estimation (Smid & Winter, 2020), we assume that transparently defining relatively uninformative priors is a constructive way to introduce a new method. Whether or not the relatively uninformative prior specification holds up in practice is tested in a simulation study. To explore different prior settings, the simulation study also includes two additional sets of prior hyperparameters: A highly uninformative set, and a highly informative set. For the highly uninformative set we defined the hyperparameters to increase the mean and variance of prior residual variances, the inverse gamma priors, that is, $a_0 = c_0 = 3$, $b_0 = d_0 = 6$, and $p_0 = 5$, $R_0 = 9$ (all other hyperparameters stay the same as above).⁶ With the highly informative set of hyperparameters, we reduced the variance of the priors. We use the population values of the manifest and latent loadings from the simulation study (see below) for prior loadings λ_0 and β_0 , and define $a_0 = c_0 = 9$, $b_0 = d_0 = 4$ (all other hyperparameters stay the same as above).⁷

It is important to keep in mind, that the informativeness of the prior distributions in the simulation study only concerns the factor model parameters. In particular, uninformative prior loadings and residual variances, both manifest and latent, do not necessarily imply uninformative priors on the coefficients ω_t and ω_h for two reasons: a) The prior distributions on ω_t and ω_h are functions of the prior loadings and residual variances of a bi-factor model, which are in turn functions of the priors of the second-order model parameters as specified above. b) The second-order factor loadings are linked to the magnitude of the prior residual variances, that is, larger prior residual variances increase the variance of the prior factor loadings. Researchers should beware of this intricate link between the factor model parameters and the ω -coefficients when

⁵The inverse gamma distribution with shape α and scale β has a mean of $\beta/(\alpha - 1)$, and a variance of $\beta^2/[(\alpha - 1)^2(\alpha - 2)]$.

⁶For the uninformative case, the hyperparameters $a_0 = c_0 = 3$, $b_0 = d_0 = 6$ yield inverse gamma priors for the residual variances with mean = 3 and variance = 9, and the hyperparameters $p_0 = 5$, $R_0 = 9$ yield an inverse gamma prior for general factor variance with mean = 3 and variance = 18.

⁷For the informative case, the hyperparameters $a_0 = c_0 = 9$, $b_0 = d_0 = 4$ yield an inverse gamma distribution for the residual variances with mean = 0.5 and variance = $\frac{1}{28}$.

exploring different prior specifications in the future. Any prior adjustments to the parameters of the second-order model can be performed using the R-package *Bayesrel*.

Method

We performed all analyses in R (v4; R Core Team, 2021) and provide the R-code for the simulation in an online repository at <https://osf.io/xyq4n/>. We constructed data generating covariance matrices based on the second-order factor model using the equation $\Lambda_0(I - B)^{-1}\Psi_\delta^+(I - B)^{-1}\Lambda_0' + \Psi_\epsilon$. We assumed a simple structure for the group factors, that is, no cross-loadings between a group factor of one set of items and an item from another set, and an equal number of items loading on each group factor. The population values of the coefficients were computed from the parameters that constructed the data generating covariance matrices.

We manipulated model size (9 items with three group factors and 30 items with five group factors) and reliability (low and high). To obtain a low ($\omega_t = .6$, $\omega_h = .4$) and a high ($\omega_t = .9$, $\omega_h = .7$) reliability condition, we drew standardized general factor loadings and group factor loadings from uniform distributions until we reached the desired ω_t and ω_h values.⁸ We generated random multivariate normal data sets centered on 0 with the data generating covariance matrices in different sample sizes ($n = 100/500/2000$). This setup resulted in a total of 12 conditions. Within the literature investigating the ω -coefficients in simulation studies, researchers use various study designs and factor compositions (cf. Cho, 2022). We assume our conditions imitate a variety of real-world scenarios. Within the substantive literature using CFA, the model sizes we chose reflect typical sets (DiStefano & Hess, 2005; Jackson et al., 2009).

For each condition, we generated 1,000 data sets that yielded properly converging models. In each replication, we set the Gibbs sampler to three chains with 10,000 iterations each (burn-in at 2000) and thinning interval at 4.⁹ The number of bootstrap replications was 2000. For the bootstrapped intervals, we always performed 2000 replications and based the CI on the converged solutions only. For the MCMC samples, we ensured convergence by computing the $\hat{R}$ -statistic (Gelman et al., 2014, Chapter 11.5) for the posterior samples of the coefficients. We assumed an $\hat{R}$ -value close to 1 and no larger than 1.1 indicated convergence (Gelman et al., 2014). When an MCMC run yielded a posterior sample with an $\hat{R}$ -value larger than 1.1, we restarted the sampling until the value was smaller than 1.1. We report the $\hat{R}$ values as well as the ESS in the results section.

For the frequentist coefficients, we used the R-package *lavaan* (Rosseel, 2012) to

⁸For each condition of reliability level and number of items the data generating standardized group factor loadings, Λ , as well as the standardized general factor loadings, β , changed: Low reliability, 9 items, three group factors, $\Lambda \in [.3, .6]$, and $\beta \in [.5, .8]$; high reliability, 9 items, three group factors, $\Lambda \in [.6, .9]$, and $\beta \in [.6, .9]$; low reliability, 30 items, five group factors, $\Lambda \in [.2, .4]$, and $\beta \in [.4, .8]$; high reliability, 30 items, five group factors, $\Lambda \in [.4, .7]$, and $\beta \in [.5, .8]$. The data generating covariance matrices can be found in an open science repository <https://osf.io/xyq4n/>.

⁹With 10,000 iterations, a burn-in of 2000, and thinning of 4, the resulting chain contains $(10000 - 2000)/4 = 2000$ values.

fit the second-order factor model in a CFA using maximum likelihood. This way, we obtained point estimates of ω_t and ω_h , the *Wald* confidence interval, and – using the boot-package (Canty & Ripley, 2021) – we obtained bootstrap samples of coefficients ω_t and ω_h to compute the remaining confidence intervals. The credible intervals are highest posterior density intervals (HPD). The point estimates of the Bayesian coefficients are the means of their posterior distributions. We added the R-code to obtain Bayesian versions of ω_t and ω_h to the R-package Bayesrel. The package also allows estimating frequentist point estimates and *Wald* intervals for ω_t and ω_h .

We considered two study outcomes, the bias of point estimation using the root mean square error (RMSE) and the coverage of the intervals. The RMSE is defined as $\sqrt{\frac{1}{T} \sum_{t=1}^T (\hat{\theta}_t - \theta)^2}$, T denotes the number of simulation runs (1,000), θ is the population value of the coefficient, and $\hat{\theta}_t$ are the point estimates of the coefficients across the T simulation runs. Within each condition, the RMSE can be interpreted as the standard deviation of the point estimates around the population value across simulation runs. Since the RMSE indicates not only the bias of a coefficient but also the sampling variability we expect a well-performing method to have small values close to but not equal to 0. The coverage was determined as the percentage of intervals across simulation runs that included the population value of the coefficient. At a confidence level of 95%, we expected a well-performing method to cover the population value in 95% of the cases.

Results

In Table 1, we see the RMSE results for the frequentist CFA and Bayesian point estimation (mean) for ω_t and ω_h . For most conditions, the frequentist and Bayesian coefficients performed equally well with the Bayesian methods displaying slightly higher RMSE values throughout most conditions. When the sample size was reasonably large, both methods achieved equally well point estimation performance. For both methods, point estimation was poorest when the reliability was low and the sample size small, which was also the condition when the largest differences between frameworks emerged: The Bayesian point estimation performed slightly worse than the frequentist, an indication of the influence of the prior distribution when evidence in the data is weak.

The coverage results of the confidence and credible intervals for ω_t and ω_h are in Tables 2 and 3, respectively. As expected, the coverage of all intervals improved with increasing sample size. When the reliability was high all confidence intervals showed good coverage performance for both coefficients. When the reliability was low none of the intervals, except for the *BCA* interval, could reach satisfactory coverage for both coefficients. The *BCA* interval showed the best coverage across all conditions. However, with 9 items, low reliability, and a sample size of 2000, the *BCA* interval could be estimated in only 185 out of 1,000 simulation runs as the adjustment constant for the skewed bootstrap sampling distribution could not be estimated. The *SE*, *SE_{bias}*, and *SE_{ln}* intervals displayed similar satisfactory coverage performance for both coefficients

Table 1

RMSE Results of Coefficients ω_t and ω_h

ρ	k	n	ω_t		ω_h	
			CFA	Bayes	CFA	Bayes
low	9	100	0.067 ₁₀₁	0.075	0.092 ₁₀	0.108
		500	0.026	0.028	0.044	0.045
		2000	0.014	0.014	0.022	0.022
	30	100	0.070 ₁₇₆	0.096	0.116 ₂	0.127
		500	0.026	0.030	0.041	0.044
		2000	0.013	0.014	0.021	0.021
high	9	100	0.015	0.018	0.051	0.049
		500	0.007	0.007	0.022	0.022
		2000	0.003	0.003	0.011	0.011
	30	100	0.015	0.016	0.047	0.046
		500	0.007	0.007	0.021	0.020
		2000	0.003	0.003	0.010	0.010

Note. The RMSE equals the standard deviation of the point estimates around the population value across simulation runs. The ρ column denotes the reliability level and represents the population values of the coefficients, low: $\omega_t = .6$, $\omega_h = .4$; high: $\omega_t = .9$, $\omega_h = .7$. The subscripts in the CFA columns denote the number of simulation runs in which the point estimates were outside the 0 to 1 range. No subscript means all estimates were between 0 and 1.

across almost all conditions. However, when reliability was low and the sample size small, the SE and SE_{bias} intervals for coefficient ω_t did not yield bounds within the 0 to 1 range in any of the simulation runs, which is visible in the 100% coverage rates (see Table 2). The issue also occurred for medium sample size, 30 items, and low reliability, that is, less than half of the simulation runs resulted in proper interval bounds for ω_t . The issue persisted for ω_h with 9 items when reliability was low and sample size small: Only about a fifth of the interval bounds fell inside the 0 to 1 range. Similar to the SE and SE_{bias} intervals, the SE_{ln} interval displayed excessive coverage in the same conditions that the SE and SE_{bias} intervals failed. In these conditions, the SE_{ln} interval yielded interval bounds very close to 0 and 1 since the SE_{ln} interval employs the same bootstrap estimated standard error as the SE and SE_{bias} intervals but its logistic transformation prevents it from leaving the 0 to 1 range. The *Perc* and *Wald* intervals for coefficient ω_t performed very similarly, only displaying insufficient coverage when sample size and reliability were low. In precisely those conditions the coverage of ω_h of both intervals was unsatisfactory but in opposite directions: The *Perc* interval covered the population value of ω_h too frequently, whereas the *Wald* interval covered it too rarely.

The Bayesian intervals performed well, except when reliability was low and sample size small: With 9 and 30 items the credible intervals could not reach satisfactory coverage performance, similar to most of the confidence intervals. Again, this could be due to the influence of the prior distribution when the data are not influential enough to shift confidence away from the prior. The convergence diagnostics for the Gibbs sampler yielded satisfactory values (see Table S1). The mean $\hat{R}$ values for the simulation

Table 2

Coverage Results of the 95% Confidence and Credible Intervals for Coefficient ω_t

ω_t	k	n	confidence intervals						Bayes
			SE	SE_{bias}	SE_{ln}	$Perc$	BCA	$Wald$	
.6	9	100	1 [†]	1 [†]	1	.795	.907	.722	.927
		500	.953	.958	.958	.949	.942	.925	.947
		2000	.946	.948	.946	.946	.941 [‡]	.951	.957
	30	100	1 [†]	1 [†]	1	.994	.979	.739	.861
		500	.982 [†]	.981 [†]	.982	.973	.967	.901	.933
		2000	.940	.940	.938	.938	.936	.937	.932
.9	9	100	.942	.943	.956	.944	.934	.939	.945
		500	.950	.950	.953	.954	.952	.950	.959
		2000	.945	.946	.948	.947	.945	.948	.952
	30	100	.954	.945	.957	.949	.941	.939	.929
		500	.954	.952	.950	.946	.948	.950	.943
		2000	.936	.937	.937	.935	.933	.935	.936

Note. The ω_t column contains the population value. Except for the *Wald* interval, all confidence intervals are bootstrapped with a CFA. The *Wald* interval is based on a single CFA. The † denotes the conditions in which both upper and lower interval bounds fell outside the 0 to 1 range in more than half of the 1,000 simulation runs. The ‡ denotes the conditions in which the *BCA* interval estimation failed in more than half of the 1,000 simulation runs.

conditions ranged from 1 to 1.003 for ω_t and from 1.001 to 1.003 for ω_h . The mean ESS for ω_t ranged from 2304 (low reliability, $k = 30$, $n = 100$) to 5965 (high reliability, $k = 9$, $n = 2000$) and for ω_h from 1729 (low reliability, $k = 30$, $n = 100$) to 5507 (high reliability, $k = 30$, $n = 2000$).

Since the prior distributions affect the performance of the coefficients, at the least when the sample size is small and the reliability low, we also ran the simulation study with a set of highly informative and a set of very uninformative priors. The detailed results for these conditions are in the supplementary materials in Tables S2 and S3, respectively. With highly informative priors the point estimation of both reliability coefficients improves, as well as the coverage of the credible intervals for ω_h . The coverage of the intervals for ω_t does not change meaningfully, except when the sample size is small, then the coverage is slightly worse than with the original set of priors. With very uninformative prior hyperparameters, the performance of the reliability coefficients is substantially worse than with the original set. The coefficients reach a satisfactory coverage rate when the sample size approaches 2000.

Data Example

To illustrate a (Bayesian) reliability analysis with ω_t and ω_h , we use the impulsivity data set of the UPPS-P we gathered from Lozano et al. (2018) containing 455 observations, 443 are complete. The data set is included in the R-package *Bayesrel* under the name

Table 3

Coverage Results of the 95% Confidence and Credible Intervals for Coefficient ω_h

ω_h	k	n	confidence intervals						Bayes
			SE	SE_{bias}	SE_{ln}	$Perc$	BCA	$Wald$	
.4	9	100	.975 [†]	.968 [†]	.975	.989	.939	.693	.901
		500	.947	.939	.947	.958	.948	.923	.957
		2000	.945	.939	.947	.947	.957 [‡]	.943	.949
	30	100	.993	.972	.996	.997	.945	.607	.832
		500	.970	.962	.971	.969	.967	.900	.946
		2000	.952	.949	.949	.956	.950	.945	.950
.7	9	100	.924	.923	.943	.944	.943	.929	.955
		500	.952	.952	.957	.957	.958	.950	.955
		2000	.940	.940	.942	.942	.947	.941	.944
	30	100	.950	.944	.962	.963	.951	.946	.937
		500	.949	.948	.951	.955	.953	.945	.939
		2000	.949	.946	.948	.953	.951	.951	.947

Note. The ω_h column contains the population value. Except for the *Wald* interval, all confidence intervals are bootstrapped with a CFA. The *Wald* interval is based on a single CFA. The † denotes the conditions in which both upper and lower interval bounds fell outside the 0 to 1 range in more than half of the 1,000 simulation runs. The ‡ denotes the conditions in which the *BCA* interval estimation failed in more than half of the 1,000 simulation runs.

“upps”.¹⁰ Interested readers can find the R-code for the calculations in the supplementary materials.

Frequentist analysis

First, we compute the reliability values of the various subscales of the impulsivity-scale. By means of coefficient α (frequentist) we obtain: negative urgency $\hat{\alpha} = .665$, lack of perseverance $\hat{\alpha} = .694$, lack of premeditation $\hat{\alpha} = .748$, sensation seeking $\hat{\alpha} = .754$, positive urgency $\hat{\alpha} = .796$. Further, we wish to obtain the reliability of the full scale; initially, we estimate $\hat{\alpha} = .809$. Since we are aware of the tendency of coefficient α to underestimate the reliability when data are multidimensional (e.g., Dunn et al., 2014), we proceed and compute coefficients ω_t and ω_h . We obtain $\hat{\omega}_t = .865$ and $\hat{\omega}_h = .644$ by means of a CFA with full-information maximum likelihood, and *Wald*-type confidence intervals ω_t -95% CI [.847, .883] and ω_h -95% CI [.594, .693]. Because of the frequentist nature of the confidence interval, we are unable to interpret the bounds of a specific interval at hand. For the ω_t -interval of [.847, .883], we cannot conclude that the true ω_t , and thus the reliability, lies within the interval bounds with 95% probability. We can conclude that if we were to repeat the sampling procedure and the estimation of the intervals numerous times, the *Wald* intervals would cover the population values of ω_t in 95% of the cases. When trying to make inferences about exceeding a cutoff value, a

¹⁰The data set is also available at https://figshare.com/articles/dataset/Concordance_between_the_original_and_short_version_of_the_Impulsive_Behaviour_Scale_UPPS-P_using_an_IRT_model/6013217.

frequentist might transform the confidence interval into a hypothesis test. In particular, the null hypothesis of the reliability being, for example, .80 is rejected at a significance level of 5% since the 95% confidence interval does not contain the value of .80. To sum up, we learn that the reliability of the scale is higher than indicated by coefficient α , and we also notice that there is uncertainty around the point estimates.

Bayesian analysis

For the Bayesian analysis, we specify the same prior hyperparameters as for the simulation study which are also the default parameters in the `Bayesrel`-package. The data contain 13 missing values, which we address with a Bayesian imputation scheme. We treat the missing values as to be estimated parameters in the MCMC sampling process (e.g., Schafer, 1999).¹¹ We estimate the Bayesian coefficients ω_t and ω_h with 10,000 iterations, 3 chains, a burn-in of 2000, and a thinning interval of 4 (the settings also used in the simulation study). The posterior distributions yield point estimates of $\hat{\omega}_t = .863$ and $\hat{\omega}_h = .640$, and credible intervals $\omega_t \in 95\% \text{ CI } [.844, .88]$ and $\omega_h \in 95\% \text{ CI } [.589, .688]$. We are 95% certain that the true ω_t lies in between .844 and .88; similarly, the probability that coefficient ω_h lies in between .589 and .688 is 95%. Figure 2 shows the prior and posterior distributions of ω_t and ω_h for the impulsivity data set. The posterior plot offers a simple display of the probability for different values of coefficients ω_t and ω_h . The amount of posterior mass in the figure depicted by the grey area is proportional to the 95% probability of different reliability values. The plot also shows that the prior distribution of ω_t is relatively uniform, whereas the prior distribution of ω_h is slightly peaked around $\omega_h = 0.1$. Both prior distributions depend on the same prior parameters of the second-order factor model (see Section “The Bayesian approach to ω_t and ω_h ”).

Furthermore, we can compute the probability that ω_t and ω_h for the impulsivity data set exceed certain cutoff values, for example, .80 and .60, respectively. Inspecting both the prior and posterior probabilities for the cutoffs, we notice the shift in our confidence by observing the data: we obtain the confidence that ω_t is larger than .80 before seeing the data from the prior distribution, which equals $p(\omega_{t_{prior}} > .80) = .221$; similarly, $p(\omega_{h_{prior}} > .60) = .202$. After taking the data into account the confidence that ω_t exceeds .80 has shifted and is obtained from the posterior distribution: $p(\omega_{t_{post}} > .80) = 1$; similarly $p(\omega_{h_{post}} > .60) = .934$.

Convergence

To check if the posterior samples of ω_t and ω_h have converged to their posterior distributions, we can display the MCMC-chains in a traceplot (see, e.g., Kruschke,

¹¹The Bayesian framework allows us to deal with missing data by treating the missing values as parameters that we can estimate conditional on the remaining data and the factor model parameters; hence we include them in the Gibbs sampling algorithm of the second-order factor model and retrieve posterior samples of the missing values.

Figure 2

Posterior and Prior Distributions of ω_t and ω_h for the Impulsivity Data Set

Note. Bayesian results of the impulsivity data set with 20 items, five group factors, and a sample size of $n = 455$. The posterior distributions of ω_t and ω_h are denoted with their prior density as a dotted line and 95% credible interval bars.

2015, Chapter 7.5), and assess $\hat{R}$ and ESS. The traceplots for ω_t and ω_h display good mixing, that is, they resemble a “hairy caterpillar” (see Figure S2; cf. M. D. Lee & Wagenmakers, 2014). Successful convergence is also indicated by $\hat{R}_{\omega_t} = 1.001$, and $\hat{R}_{\omega_h} = 1.000$, meaning the within-chain variance is close to the total variance, and continuing the sampling would not significantly result in more information about the target distribution (Gelman et al., 2014, Chapter 11.5). The level of autocorrelation is acceptable, that is, $ESS_{\omega_t} = 5654$, and $ESS_{\omega_h} = 5193$, in other words, the posterior samples for coefficient ω_t and ω_h consist of more than 5000 independent samples each that contain as much information as the full samples of 6000. We also conclude that the posterior precision is satisfactory since a high ESS indicates small posterior uncertainty (Stan Development Team, 2022).

Model fit

In the previous paragraphs, we have treated coefficient ω_t and ω_h as measures of reliability. However, as noted in the introduction, the coefficients are only appropriate reliability coefficients when the corresponding factor model properly describes the data.

For the frequentist CFA model, we obtain a root mean square error of approximation (RMSEA) value of $RMSEA = .057$, with $p(RMSEA < .05) = .039$, and $RMSEA$ 90% CI [.051, .064], and a standardized root mean square residual (SRMR) of $SRMR = .057$, with $SRMR$ 90% CI [.047, .066].¹²

In the Bayesian framework, we check if the model fits the data in two different ways: 1) In a post-hoc graphical comparison called posterior predictive check (PPC) of the posterior model implied parameters and the data parameters, and 2) with indices that quantify the fit of the factor model similar to the frequentist fit measures. For the PPC, we employ the posterior distributions of the second-order factor model parameters to simulate data which we then compare to the actual data set. The PPC of the second-order factor model resembles a scree plot in that the eigenvalues of covariance matrices predicted by the model are plotted against the eigenvalues of the data covariance matrix. Inspecting the PPC plot for the impulsivity example, we assess the fit as mediocre (see Figure S3).

For the fit indices, we implemented the Bayesian versions of the RMSEA and the SRMR in the *Bayesrel*-package. Similar to the frequentist fit we obtain $BRMSEA = .057$, with $p(BRMSEA < .05) = 0$, and $BRMSEA$ 90% CI [.055, .058]; $BSRMR = .070$. The coefficients ω_t and ω_h may be interpreted as measures of reliability cautiously. The structure of the data should always be assessed when conducting a reliability analysis that makes assumptions about that structure. An a-priori check could assess the loading structure of the data set using an exploratory factor analysis (EFA) and adjust the factor model accordingly. For example, we could decide to allow an item to load on more than one factor – if we view such a constraint as contributing to the true score we intend to measure.

Robustness

In the above, we used relatively uninformative priors for the second-order factor model. To demonstrate an analysis with more informative priors, we alter some of the prior parameters to reflect a scenario with larger prior loading values and smaller residual variances than the default settings used for the simulation. Specifically, we adjust the prior primary loadings mean to $\lambda_0 = 2$, the prior secondary loadings mean to $\beta_0 = 2$, the prior manifest residuals shape and scale to $a_0 = 10$ and $b_0 = 6$, respectively, and the prior latent residuals shape and scale to $c_0 = 10$ and $d_0 = 6$, respectively.¹³ Then, we run the Bayesian analysis with the impulsivity data set again, and all other parameters stay the same. The resulting coefficients are $\hat{\omega}_t = .855$ and $\hat{\omega}_h = .630$, and credible intervals $\omega_t \in 95\% \text{ CI } [.836, .875]$ and $\omega_h \in 95\% \text{ CI } [.578, .680]$. The results are similar to the ones with relatively uninformative priors and we conclude that the posterior sampling with the given data set is robust to the prior specification.

¹²The *SRMR* and its confidence interval were computed using the unbiased method from Maydeu-Olivares (2017) implemented in *lavaan*.

¹³An inverse gamma with shape and rate parameters $a_0 = 10$ and $b_0 = 6$ yields residual variances of 2/3 on average.

Discussion

Ample research describes the shortcomings of coefficient α and other lower bounds to estimate reliability when data are multidimensional (e.g., Green & Yang, 2015; McNeish, 2018; Savalei & Reise, 2019; Sijtsma, 2009). More suited coefficients for this purpose are ω_t and ω_h which estimate different forms of reliability (e.g., Flora, 2020; Revelle & Zinbarg, 2009; Zinbarg et al., 2005). In this study, we investigated the performance of several procedures to quantify the uncertainty in reliability estimation using confidence intervals for ω_t and ω_h . The results provide a more comprehensive picture of the confidence intervals for reliability coefficients often used in research. Furthermore, we presented versions of coefficients ω_t and ω_h in the Bayesian framework, thereby also providing Bayesian credible intervals. The credible intervals are a well-performing alternative to confidence intervals that present an intuitive approach to quantifying uncertainty in reliability analysis.

Confidence intervals

In the simulation study, the *BCA* interval showed the best performance among the confidence intervals. The remaining confidence intervals performed well in most conditions, except when the sample size was small and the reliability was low. Before making any recommendations, it is important to note three issues that surfaced during the simulation study. First, all studied confidence intervals depend on a converging CFA with the original data set. Non-convergence and nonsensical parameter estimates (Heywood cases) are frequent issues with small samples and noisy data. Second, sometimes CFA with bootstrapped data sets yielded nonsensical parameter estimates. This resulted in out-of-range interval bounds and excessive coverage rates for the SE , SE_{bias} , and SE_{ln} intervals. There are multiple possible remedies to these issues. One could use EFA to obtain the ω -coefficients (see, e.g., Revelle & Condon, 2019). Alternatively, one could use bounded estimation for the CFAs involved (see De Jonckere & Rosseel, 2022). Further, within CFA, one could reject all bootstrapped parameter estimates of ω_t and ω_h that fall outside the 0 to 1 range, since these inflate the standard error of the coefficients and lead to coverage issues. Thus, one would obtain the bootstrapped intervals only based on a “bounded” empirical sampling distribution. Third, the *BCA* interval performed well but is not always available, especially when the data is noisy and the sample size is close to the number of bootstrap replications. In these cases, it may help to increase the number of bootstrap samples or – if to no avail – resort to another interval.

To sum up, if a confidence interval is desired, the CFA has no convergence issues, the model fits the data well, and the bootstrapped *BCA* interval is available, researchers may use it to accompany point estimates for ω_t and ω_h . When the reliability level is relatively high or the sample size is medium to large, all confidence intervals can be recommended. In any case, it makes sense to inspect the parameter estimates of the CFA model for any peculiarities and to compare multiple confidence interval bounds to

get a better understanding of the uncertainty around the reliability coefficients.

Credible intervals

The credible intervals reached good coverage rates (close to 95%) when the sample size was at least 500. With small sample sizes and low reliability values, the coverage rates did not reach a satisfactory level, which was also true for the confidence intervals. In general, the Bayesian coefficients performed similarly to the frequentist ones and can be recommended to any researcher who prefers to do their analysis in the Bayesian framework. The posterior distributions of coefficients ω_t and ω_h allow researchers to obtain any X% interval that contains the coefficient with an X% probability and compute the probability that the coefficient value is larger than any specified value. Such inferences are impossible to be drawn from confidence intervals.

Similar to the frequentist approach, the Bayesian versions of ω_t and ω_h require an adequately specified model structure. Often researchers estimate reliability as a prerequisite for further analysis. The prior distribution and the need for a precise model definition demand that they give thought to the prior confidence they have in different reliability and factor model parameter values before seeing the data, and that they consider the exact relations between the latent factors and manifest variables for their data.

The prior distribution is an important part of Bayesian parameter estimation; it quantifies the researcher's knowledge before the data have been observed. In settings where the sample size is reasonably large, the influence of the prior distribution is relatively small, since the data provide enough evidence to quickly change the prior beliefs ("the data overwhelm the prior"; Wrinch & Jeffreys, 1919). However, the conditions with highly informative and uninformative priors showed that when the sample size is small and the evidence in the data is insufficient (low reliability), the prior distribution affects the performance of the Bayesian reliability coefficients. In particular, when the prior parameters are highly indicative of the true factor model parameters, the Bayesian coefficients quickly converge to their population values; when the prior is not at all informative about the true parameters, the population values of the ω -coefficients are only well approximated when the sample size is large enough to overcome the prior. This is a finding in line with previous research on Bayesian SEM (MacCallum et al., 2012; Smid & Winter, 2020).

The goal of this study was to investigate relatively uninformative prior distributions for the parameters of the second-order factor model that perform well in practice. For future research on the topic, we propose three ways in which the prior distributions from this study can be altered. First, the prior distribution can incorporate prior knowledge into the estimation process. For instance, the parameters of the second-order factor model, i.e., loadings and residual variances from a previous study using the same test instrument, can be incorporated in a reliability analysis as the parameters of the prior distribution and thereby improve the precision of the Bayesian reliability

estimation. Second, we computed the prior distributions of the reliability coefficients from the prior distributions of the second-order factor model parameters. We chose the same priors for the factor model parameters for both coefficients, which resulted in differently shaped prior distributions for ω_t and ω_h (see Figure 2). The prior for ω_t distributes confidence almost evenly along the x-axis, whereas the prior for ω_h places a considerable amount of confidence on values between 0 and 0.1 (see the modest peak in Figure 2). Researchers may want to tailor the prior distributions of the factor model parameters to each coefficient separately to obtain prior distributions that better reflect their prior beliefs about the coefficients. Third, the prior distribution is affected by the number of items in a test and the number of group factors specified. In the simulation study, we chose the same prior for both model sizes ($k = 9$, and $k = 30$). Researchers may investigate how priors are affected by a varying number of items.

Practical implications for applied researchers

The key message substantive researchers should take away from this study is to always account for uncertainty in their reliability analysis. When uncertainty estimation is neglected, sampling error is dismissed, and one risks learning nothing about the reliability in the population. In practice, we recommend researchers first decide on the measurement model underlying their scale. If they decide on the second-order factor model they need to ensure that the model approximately fits their data set. When reliability point estimation is of special interest, we refer readers to a tutorial by Revelle and Condon (2019), although it is hard to imagine a scenario where one would not want to account for sampling error. If frequentist confidence intervals are desired, a second-order CFA model and the reliability coefficients may be estimated, and, if the sample size is at least 500 or the reliability is relatively high, any of the confidence intervals investigated in the simulation study will suffice to reflect uncertainty about the reliability estimates. When the sample size is small and the reliability low, the *BCA* interval should be used.¹⁴

If Bayesian statistics are favoured, we recommend using the R-package *Bayesrel* for the reliability analysis. The package allows to estimate the Bayesian ω_t and ω_h and to adjust the prior hyperparameters of the second-order factor model. Given a specific data set, the first step in a Bayesian reliability analysis should be the decision on how much prior confidence one has in the parameters of the second-order factor model, and thus in the reliability values. We distinguish two scenarios: prior information about the scale in question is available, and no prior information is available. When prior information is available, researchers may incorporate this information by specifying more informative priors for the second-order factor model (see subsection “Priors” pp. 14-16). The prior variances of all parameters should be reasonably large regardless so that the posterior is not bound by the prior distribution. From the simulation study one might learn

¹⁴Interested readers may obtain bootstrapped confidence intervals for the ω -coefficients by bootstrapping the second-order factor model in *lavaan* while specifying the ω -coefficients as estimands in the *lavaan*-syntax.

that when the reliability is low and the sample size small, a poor performance by the reliability coefficients might be counteracted by highly informative priors with small variances. We advise against such a procedure since one risks learning nothing about the data that goes beyond the information already contained in the prior. Even when prior information is available, we advise always estimating the Bayesian coefficients with at least one more set of priors that is less informative, for example, the relatively uninformative set from the simulation study. Naturally, one should present the results with both sets of priors in a transparent manner. This way, one is safeguarded from making inferences about the reliability coefficients that are too dependent on prior information.

When no prior information is available it makes sense to start the reliability analysis with a relatively uninformative set of priors. Subsequently, the reliability coefficients should also be estimated with multiple sets of priors differing in informativeness (Depaoli et al., 2021; Van Erp et al., 2018). The resulting coefficients and their inferences should be reported for each prior setting. Ideally, the sample size or the reliability is large enough so the prior settings have little to no influence on the reliability estimates (see the data example). In any case, a reported reliability analysis should illuminate the reasoning for the choice of prior and contain results for more than one set of prior parameters.

Limitations

We chose the second-order factor model to estimate coefficients ω_t and ω_h . The second-order model suits many psychological constructs that are measured by multiple group factors that relate to a general factor (e.g., Moshagen, in press). For those constructs that do not fit the second-order model, other factor models are appropriate. For instance, a general factor for the impulsivity-scale might explain variance in the manifest variables not through group factors but directly; hence, the group factors explain the variance in the manifest variables that cannot be explained by the general factor. Such a bi-factor model is a general factor model that includes the second-order factor model. Fitting a second-order factor model and transforming it into a bi-factor model – as done in this study – is not equivalent to fitting a bi-factor model, unless the second-order model fits the data. In other words, all second-order models are bi-factor models but not all bi-factor models are second-order models. The procedure to estimate coefficients ω_t and ω_h from a bi-factor model is mostly similar to the procedure described in this study and differs in the way the loadings and residual variances are obtained to compute the reliability coefficients (see, e.g., Flora, 2020; Revelle & Condon, 2019).

The model structure we implemented in the simulation study was simple, meaning there were no cross-loadings or error covariances specified. Real data may contain cross-loadings or correlating errors. In both the frequentist and Bayesian framework such aspects can be addressed through the model specification before running the analysis (see, e.g., Flora, 2020). Whereas the confirmatory approach is inherently subject

to identification issues, an EFA can avoid these. The EFA approach to estimating coefficients ω_t and ω_h is well known in the frequentist framework (see Revelle, 2020) and could also be applied in the Bayesian framework by fitting a Bayesian EFA (Conti et al., 2014). Within an EFA model, all possible cross-loadings are estimated. Arguably, it is a matter of a given research question whether or not to model all possible loadings between group factors and items. On the one hand, it can maximize the estimated true score variance, on the other hand, it can lead to capitalization on chance.

Conclusion

In practice, measurement instruments are often multidimensional, as are the psychological constructs they attempt to measure. The coefficients ω_t and ω_h estimate the total and general factor reliability, respectively, for such instruments. We assume it is undisputed that an estimate of uncertainty around the point estimates of these coefficients is needed, but it is an often neglected practice. To provide researchers with a comprehensive picture of the available uncertainty intervals around ω_t and ω_h , we compared the confidence intervals for the coefficients in a simulation study. Furthermore, we described a method to obtain Bayesian versions of ω_t and ω_h and their posterior distributions. The information in this study allows researchers to make informed choices about the confidence intervals they wish to report for reliability coefficients or even go beyond standard frequentist practice by expressing their practical beliefs about reliability estimates by employing the Bayesian framework with posterior distributions of ω_t and ω_h .

References

- Asparouhov, T., & Muthén, B. (2021). *Bayesian analysis of latent variable models using Mplus*. <https://www.statmodel.com/download/BayesAdvantages18.pdf>
- Bader, M., & Moshagen, M. (in press). No probifactor model fit index bias, but a propensity toward selecting the best model. *Journal of Abnormal Psychology*. <https://doi.org/10.1037/abn0000685>
- Bollen, K. A. (1989). *Structural equations with latent variables*. John Wiley & Sons, Inc. <https://doi.org/10.1002/9781118619179>
- Box, G. E. P., & Tiao, G. C. (1973). *Bayesian inference in statistical analysis*. Addison-Wesley. <https://doi.org/10.1002/9781118033197>
- Canty, A., & Ripley, B. D. (2021). *Boot: Bootstrap R (S-Plus) functions*. (R package version 1.3-27).
- Cho, E. (2022). The accuracy of reliability coefficients: A reanalysis of existing simulations. *Psychological Methods*. <https://doi.org/10.1037/met0000475>
- Conti, G., Frühwirth-Schnatter, S., Heckman, J. J., & Piatek, R. (2014). Bayesian exploratory factor analysis. *Journal of Econometrics*, 183(1), 31–57. <https://doi.org/10.1016/j.jeconom.2014.06.008>

- Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, 16(3), 297–334. <https://doi.org/10.1007/BF02310555>
- Cumming, G. (2014). The new statistics: Why and how. *Psychological Science*, 25(1), 7–29. <https://doi.org/10.1177/0956797613504966>
- Cyders, M. A., Littlefield, A. K., Coffey, S., & Karyadi, K. A. (2014). Examination of a short English version of the UPPS-P Impulsive Behavior Scale. *Addictive Behaviors*, 39(9), 1372–1376. <https://doi.org/10.1016/j.addbeh.2014.02.013>
- Davison, A. C., & Hinkley, D. V. (1997). *Bootstrap methods and their application* (1st ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511802843>
- De Jonckere, J., & Rosseel, Y. (2022). Using bounded estimation to avoid nonconvergence in small sample structural equation modeling. *Structural Equation Modeling: A Multidisciplinary Journal*, 1–16. <https://doi.org/10.1080/10705511.2021.1982716>
- Depaoli, S., Liu, H., & Marvin, L. (2021). Parameter specification in Bayesian CFA: An exploration of multivariate and separation strategy priors. *Structural Equation Modeling: A Multidisciplinary Journal*, 28(5), 699–715. <https://doi.org/10.1080/10705511.2021.1894154>
- DiCiccio, T. J., & Efron, B. (1996). Bootstrap confidence intervals. *Statistical Science*, 11(3), 189–212. <https://doi.org/10.1214/ss/1032280214>
- DiStefano, C., & Hess, B. (2005). Using confirmatory factor analysis for construct validation: An empirical review. *Journal of Psychoeducational Assessment*, 23(3), 225–241. <https://doi.org/10.1177/073428290502300303>
- Dunn, T. J., Baguley, T., & Brunsden, V. (2014). From alpha to omega: A practical solution to the pervasive problem of internal consistency estimation. *British Journal of Psychology*, 105(3), 399–412. <https://doi.org/10.1111/bjop.12046>
- Efron, B. (1979). Bootstrap methods: Another look at the Jackknife. *The Annals of Statistics*, 7(1), 1–26. <https://doi.org/10.2307/2958830>
- Flora, D. B. (2020). Your coefficient alpha is probably wrong, but which coefficient omega is right? A tutorial on using R to obtain better reliability estimates. *Advances in Methods and Practices in Psychological Science*, 3(4), 484–501. <https://doi.org/10.1177/2515245920951747>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2014). *Bayesian data analysis* (3rd ed.). CRC Press.
- Geman, S., & Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (6), 721–741.
- Geyer, C. (2011). Introduction to Markov chain Monte Carlo. In S. Brooks, A. Gelman, G. Jones, & X.-L. Meng (Eds.), *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC. <https://doi.org/10.1201/b10905-2>
- Gilks, W. R., Richardson, S., & Spiegelhalter, D. (1995). *Markov chain Monte Carlo in practice*. Chapman and Hall/CRC. <https://doi.org/10.1201/b14835>
- Green, S. B., & Yang, Y. (2015). Evaluation of dimensionality in the assessment of internal consistency reliability: Coefficient alpha and omega coefficients. *Educational*

- Measurement: Issues and Practice*, 34(4), 14–20. <https://doi.org/10.1111/emip.12100>
- Jackson, D. L., Gillaspy, J. A., & Purc-Stephenson, R. (2009). Reporting practices in confirmatory factor analysis: An overview and some recommendations. *Psychological Methods*, 14(1), 6–23. <https://doi.org/10.1037/a0014694>
- Jöreskog, K. G. (1971). Statistical analysis of sets of congeneric tests. *Psychometrika*, 36(2), 109–133. <https://doi.org/10.1007/BF02291393>
- Kelley, K., & Pornprasertmanit, S. (2016). Confidence intervals for population reliability coefficients: Evaluation of methods, recommendations, and software for composite measures. *Psychological Methods*, 21(1), 69–92. <https://doi.org/10.1037/a0040086>
- Kruschke, J. K. (2015). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd ed.). Elsevier. <https://doi.org/10.1016/C2012-0-00477-2>
- Lee, M. D., & Wagenmakers, E.-J. (2014). *Bayesian cognitive modeling: A practical course*. Cambridge University Press.
- Lee, S.-Y. (2007). *Structural equation modeling: A Bayesian approach*. John Wiley & Sons Ltd. <https://doi.org/10.1002/9780470024737>
- Lord, F. M., & Novick, M. R. (1968). *Statistical theories of mental test scores*. Addison-Wesley.
- Lozano, Ó. M., Díaz-Batanero, C., Rojas, A. J., Pilatti, A., & Fernández-Calderón, F. (2018). Concordance between the original and short version of the Impulsive Behaviour Scale UPPS-P using an IRT model. *PLOS ONE*, 13(3), 1–15. <https://doi.org/10.1371/journal.pone.0194390>
- MacCallum, R. C., Edwards, M. C., & Cai, L. (2012). Hopes and cautions in implementing Bayesian structural equation modeling. *Psychological Methods*, 17(3), 340–345. <https://doi.org/10.1037/a0027131>
- Mansolf, M., & Reise, S. P. (2017). When and why the second-order and bifactor models are distinguishable. *Intelligence*, 61, 120–129. <https://doi.org/10.1016/j.intell.2017.01.012>
- Maydeu-Olivares, A. (2017). Assessing the size of model misfit in structural equation models. *Psychometrika*, 82(3), 533–558. <https://doi.org/10.1007/s11336-016-9552-7>
- McDonald, R. P. (1970). The theoretical foundations of principal factor analysis, canonical factor analysis, and alpha factor analysis. *British Journal of Mathematical and Statistical Psychology*, 23(1), 1–21. <https://doi.org/10.1111/j.2044-8317.1970.tb00432.x>
- McDonald, R. P. (1999). *Test theory: A unified treatment* (1st ed.). Psychology Press. <https://doi.org/10.4324/9781410601087>
- McNeish, D. (2018). Thanks coefficient alpha, we'll take it from here. *Psychological Methods*, 23(3), 412–433. <https://doi.org/10.1037/met0000144>
- Merkle, E. C., & Rosseel, Y. (2018). Blavaan: Bayesian structural equation models via parameter expansion. *Journal of Statistical Software*, 85(4). <https://doi.org/10.18637/jss.v085.i04>

- Morey, R. D., Hoekstra, R., Rouder, J. N., Lee, M. D., & Wagenmakers, E.-J. (2016). The fallacy of placing confidence in confidence intervals. *Psychonomic Bulletin & Review*, 23(1), 103–123. <https://doi.org/10.3758/s13423-015-0947-8>
- Moshagen, M. (in press). When a truly positive correlation turns negative: How different approaches to model hierarchically structured constructs affect estimated correlations to covariates. *European Journal of Personality*. <https://doi.org/10.1177/09500002070211050170>
- Mulaik, S. A., & Quartetti, D. A. (1997). First order or higher order general factor? *Structural Equation Modeling: A Multidisciplinary Journal*, 4(3), 193–211. <https://doi.org/10.1080/10705519709540071>
- Muthén, B., & Asparouhov, T. (2012). Bayesian structural equation modeling: A more flexible representation of substantive theory. *Psychological Methods*, 17(3), 313–335. <https://doi.org/10.1037/a0026802>
- Neyman, J. (1937). Outline of a theory of statistical estimation based on the classical theory of probability. *Philosophical Transactions of the Royal Society of London, Series A*, 236(767), 333–380. <https://doi.org/10.1098/rsta.1937.0005>
- Padilla, M. A., & Divers, J. (2016). A comparison of composite reliability estimators: Coefficient omega confidence intervals in the current literature. *Educational and Psychological Measurement*, 76(3), 436–453. <https://doi.org/10.1177/0013164415593776>
- Pfadt, J. M., Van den Bergh, D., & Goosen, J. (2022). *Bayesrel: Bayesian reliability estimation* (R package version 0.7.4.3). <https://github.com/juliuspf/Bayesrel>
- Pfadt, J. M., van den Bergh, D., Sijtsma, K., Moshagen, M., & Wagenmakers, E.-J. (2022). Bayesian estimation of single-test reliability coefficients. *Multivariate Behavioral Research*, 57(4), 620–641. <https://doi.org/10.1080/00273171.2021.1891855>
- R Core Team. (2021). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing.
- Revelle, W. (2020). *Psych: Procedures for psychological, psychometric, and personality research*. (R-package version 2.0.12). Northwestern University. <https://CRAN.R-project.org/package=psych>
- Revelle, W., & Condon, D. M. (2019). Reliability from α to ω : A tutorial. *Psychological Assessment*, 31(12), 1395–1411. <https://doi.org/10.1037/pas0000754>
- Revelle, W., & Zinbarg, R. E. (2009). Coefficients alpha, beta, omega, and the glb: Comments on Sijtsma. *Psychometrika*, 74(1), 145–154. <https://doi.org/10.1007/s11336-008-9102-z>
- Rodriguez, A., Reise, S. P., & Haviland, M. G. (2016). Evaluating bifactor models: Calculating and interpreting statistical indices. *Psychological Methods*, 21(2), 137–150. <https://doi.org/10.1037/met0000045>
- Rosseel, Y. (2012). Lavaan: An R package for structural equation modeling. *Journal of Statistical Software*, 48(2), 1–36. <https://doi.org/10.18637/jss.v048.i02>

- Savalei, V., & Reise, S. P. (2019). Don't forget the model in your model-based reliability coefficients: A reply to McNeish (2018). *Collabra: Psychology*, 5(1), 36. <https://doi.org/10.1525/collabra.247>
- Schafer, J. L. (1999). *Analysis of incomplete multivariate data* (1. ed.). Chapman & Hall/CRC. <https://doi.org/10.1201/9780367803025>
- Schmid, J., & Leiman, J. M. (1957). The development of hierarchical factor solutions. *Psychometrika*, 22(1), 53–61. <https://doi.org/10.1007/BF02289209>
- Sijtsma, K. (2009). On the use, the misuse, and the very limited usefulness of Cronbach's alpha. *Psychometrika*, 74(1), 107–120. <https://doi.org/10.1007/s11336-008-9101-0>
- Sijtsma, K., & Pfadt, J. M. (2021). Part II: On the use, the misuse, and the very limited usefulness of Cronbach's alpha: Discussing lower bounds and correlated errors. *Psychometrika*, 86(4), 843–860. <https://doi.org/10.1007/s11336-021-09789-8>
- Smid, S. C., & Winter, S. D. (2020). Dangers of the defaults: A tutorial on the impact of default priors when using Bayesian SEM with small samples. *Frontiers in Psychology*, 11, 611963. <https://doi.org/10.3389/fpsyg.2020.611963>
- Spearman, C. (1904). The proof and measurement of association between two things. *The American Journal of Psychology*, 15(1), 72–101. <https://doi.org/10.2307/1412159>
- Stan Development Team. (2022). *Stan modeling language users guide and reference manual*, v2.18. <https://mc-stan.org>
- Ten Berge, J. M. F., & Sočan, G. (2004). The greatest lower bound to the reliability of a test and the hypothesis of unidimensionality. *Psychometrika*, 69(4), 613–625. <https://doi.org/10.1007/bf02289858>
- Van Erp, S., Mulder, J., & Oberski, D. L. (2018). Prior sensitivity analysis in default Bayesian structural equation modeling. *Psychological Methods*, 23(2), 363–388. <https://doi.org/10.1037/met0000162>
- Woodhouse, B., & Jackson, P. H. (1977). Lower bounds for the reliability of the total score on a test composed of non-homogeneous items: II: A search procedure to locate the greatest lower bound. *Psychometrika*, 42(4), 579–591. <https://doi.org/10.1007/bf02295980>
- Wrinch, D., & Jeffreys, H. (1919). On some aspects of the theory of probability. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 38(228), 715–731. <https://doi.org/10.1080/14786441208636005>
- Zinbarg, R. E., Revelle, W., Yovel, I., & Li, W. (2005). Cronbach's alpha, Revelle's beta, and McDonald's omega h: Their relations with each other and two alternative conceptualizations of reliability. *Psychometrika*, 70(1), 123–133. <https://doi.org/10.1007/s11336-003-0974-7>

Supplementary Material for Manuscript “Classical and Bayesian Uncertainty Intervals for the Reliability of Multidimensional Scales” (Pfadt et al., in press)

Factor Models

To compute coefficients ω_t and ω_h the loadings of the second-order factor model are transformed into the loadings of a bi-factor model (see Figure S1), which is possible since the second-order model is nested in the bi-factor model (Mulaik & Quartetti, 1997; Schmid & Leiman, 1957). From the loadings of the bi-factor model, one may compute coefficients ω_t and ω_h (see Equations 4 and 5).

Figure S1

Exemplary Bi-Factor Model with Five Group Factors

Additional Simulation Results

Table S1

Mean Convergence Diagnostics of Coefficients ω_t and ω_h

ρ	k	n	ω_t		ω_h	
			$\widehat{R}$	ESS	$\widehat{R}$	ESS
low	9	100	1.001	3466.740	1.003	1835.555
		500	1.001	4731.534	1.003	1822.423
		2000	1.001	4927.804	1.003	1835.381
	30	100	1.003	2303.951	1.003	1729.427
		500	1.001	3707.396	1.002	2404.382
		2000	1.001	3971.618	1.002	2545.898
high	9	100	1.001	5946.222	1.001	4905.854
		500	1.001	5958.008	1.001	4836.592
		2000	1.000	5964.748	1.001	4824.499
	30	100	1.001	5880.553	1.001	5423.408
		500	1.001	5933.088	1.001	5504.919
		2000	1.000	5942.557	1.001	5507.080

Note. The $\widehat{R}$ values should lie between 1 and 1.1. The ESS is bound by the number of autocorrelated posterior samples, which is 6,000. The ρ column denotes the reliability level and represents the population values of the coefficients, low: $\omega_t = .6$, $\omega_h = .4$; high: $\omega_t = .9$, $\omega_h = .7$.

Table S2

Simulation Results for Coefficients ω_t and ω_h with Informative Priors

ρ	k	n	ω_t		ω_h	
			RMSE	Coverage	RMSE	Coverage
low	9	100	0.065	.913	0.089	.947
		500	0.027	.945	0.042	.950
		2000	0.014	.942	0.022	.954
	30	100	0.078	.881	0.109	.888
		500	0.027	.943	0.041	.951
		2000	0.013	.937	0.021	.949
high	9	100	0.019	.942	0.043	.974
		500	0.007	.958	0.021	.958
		2000	0.003	.950	0.011	.943
	30	100	0.015	.891	0.043	.919
		500	0.007	.934	0.020	.944
		2000	0.003	.939	0.010	.942

Note. The ρ column denotes the reliability level and represents the population values of the coefficients, low: $\omega_t = .6$, $\omega_h = .4$; high: $\omega_t = .9$, $\omega_h = .7$. The RMSE equals the standard deviation of the point estimates around the population value across simulation runs. The coverage is determined with $\alpha = .05$, thus a level of .95 is desirable. The informative priors have small variances and the prior loadings are the data generating loadings.

Data example

The Bayesian reliability analysis for the impulsivity questionnaire data is conducted in R. In particular, we estimate coefficients ω_t and ω_h using the Bayesrel package. The latest version can be installed by `remotes::install_github("juliuspf/Bayesrel")`. The data set together with the corresponding article can be found in the repository <https://osf.io/xyq4n/>. Here, we conduct the analysis as described in the main body of the text.

```
library(Bayesrel)
library(coda)
library(LaplacesDemon)
library(psych)
# get the data set
data <- upps

# ----- the frequentist analysis ----- #
# first the frequentist reliability of the subscales
# by means of alpha
alpha(data[, 1:4])$total$raw
alpha(data[, 5:8])$total$raw
alpha(data[, 9:12])$total$raw
alpha(data[, 13:16])$total$raw
alpha(data[, 17:20])$total$raw
```

Table S3

Simulation Results for Coefficients ω_t and ω_h with Uninformative Priors

ρ	k	n	ω_t		ω_h	
			RMSE	Coverage	RMSE	Coverage
high	9	100	0.148	.660	0.225	.357
		500	0.035	.907	0.065	.862
		2000	0.015	.933	0.025	.925
	30	100	0.169	.439	0.244	.247
		500	0.042	.816	0.070	.789
		2000	0.015	.905	0.025	.909
high	9	100	0.048	.455	0.117	.685
		500	0.011	.818	0.030	.899
		2000	0.004	.915	0.012	.942
	30	100	0.032	.759	0.092	.805
		500	0.009	.916	0.026	.914
		2000	0.004	.936	0.011	.943

Note. The ρ column denotes the reliability level and represents the population values of the coefficients, low: $\omega_t = .6$, $\omega_h = .4$; high: $\omega_t = .9$, $\omega_h = .7$. The RMSE equals the standard deviation of the point estimates around the population value across simulation runs. The coverage is determined with $\alpha = .05$, thus a level of .95 is desirable. The uninformative priors have large variances.

```
# coefficient alpha of the full scale
alpha(data)$total$raw

# frequentist omegaT and omegaH with CFA
res_freq <- omegasCFA(data, n.factors = 5, missing = "pairwise",
                      interval = .95)
res_freq

# ----- the Bayesian analysis ----- #
# for replicability set a seed:
set.seed(11)
# do the analysis
res_bayes <- bomegas(data, n.factors = 5, missing = "impute",
                    interval = .95, n.iter = 1e4, n.burnin = 2e3,
                    thin = 4)

# point estimates of omega_t and omega_h and credible intervals:
res_bayes

# prior and posterior probability for exceeding cutoffs
pOmegas(res_bayes, cutoff.t = .8, cutoff.h = .6)

# ##### convergence diagnostics ##### #
```

```

library(coda)
samp_omega_t <- res_bayes[["omega_t"]][["chains"]]
samp_omega_h <- res_bayes[["omega_h"]][["chains"]]

# restructuring the chains for traceplot:
samp_mcmc_omega_t <- lapply(as.data.frame(t(samp_omega_t)), mcmc)
samp_mcmc_omega_h <- lapply(as.data.frame(t(samp_omega_h)), mcmc)

# traceplots
traceplot(samp_mcmc_omega_t)
traceplot(samp_mcmc_omega_h)

# for R-hat or potential scale reduction factor:
gelman.diag(samp_mcmc_omega_t)$psrf[, 1]
gelman.diag(samp_mcmc_omega_h)$psrf[, 1]

# Effective sample size (the closer to 6000 the better):
sum(sapply(samp_mcmc_omega_t, ESS))
sum(sapply(samp_mcmc_omega_h, ESS))

# ##### model fit ##### #
# frequentist
res_freq # or
res_freq$model$fit.measures

# Bayesian:
# posterior predictive check for the fit of the second
# order factor model and fit indices
secoFit(res_bayes, data, ppc = TRUE, cutoff = .05)

# ##### robustness ##### #
# more informative prior parameters
set.seed(11)
res_adj_prior <- bomegas(data, n.factors = 5, missing = "impute",
                        interval = .95, n.iter = 1e4, n.burnin = 2e3, thin = 4
                        l0 = 2, beta0 = 2, a0 = 10, b0 = 6, c0 = 10, d0 = 6)
res_adj_prior

```

In R, the PPC for the second-order factor model is obtained by:

```
# posterior predictive check and fit indices for the fit of the second-order  
# factor model  
secoFit(res, data)
```

The posterior predictive check for the second-order factor model is created as follows:

1) We take the samples of the posterior second-order factor model parameters, i.e., the loadings, residual variances, and factor variances, and compute a posterior sample of the model implied covariance matrices using the equation $\mathbf{\Lambda}_0(\mathbf{I} - \mathbf{B})^{-1}\mathbf{\Psi}_\delta^+(\mathbf{I} - \mathbf{B})^{-1}\mathbf{\Lambda}_0^T + \mathbf{\Psi}_\epsilon$;

2) we draw samples of size n , n being the sample size of the data, from a multivariate normal with means $\mathbf{0}$ and the posterior model implied covariance matrices until we obtained as many simulated data sets as there are posterior model implied covariance matrices; 3) we compute the eigenvalues of the covariance matrices of the simulated data sets; 4) this way we obtain an empirical distribution for each eigenvalue; 5) we compute the 95% quantiles for each eigenvalue distribution and plot them as grey bars together with the eigenvalues of the original data set (see Figure S3).

Figure S2

Traceplots of Coefficients ω_t and ω_h for the Impulsivity Data

Note. Traceplots of the three MCMC chains of ω_t and ω_h for the impulsivity-data produced by the traceplot command from the coda-package in R. The plots are prototypical for good convergence.

Figure S3

Posterior Predictive Check of the Second-Order Factor Model for the Impulsivity Data

Note. PPC of the second-order factor model for the impulsivity-data from the Bayesrel-package. A good model fit implies that the grey bars enclose the black dots.

References

- Mulaik, S. A., & Quartetti, D. A. (1997). First order or higher order general factor? *Structural Equation Modeling: A Multidisciplinary Journal*, 4(3), 193–211. <https://doi.org/10.1080/10705519709540071>
- Pfadt, J. M., van den Bergh, D., & Moshagen, M. (in press). Classical and Bayesian uncertainty intervals for the reliability of multidimensional scales. *Structural Equation Modeling: A Multidisciplinary Journal*.
- Schmid, J., & Leiman, J. M. (1957). The development of hierarchical factor solutions. *Psychometrika*, 22(1), 53–61. <https://doi.org/10.1007/BF02289209>