

1 Teacher's Corner: Evaluating Informative Hypotheses Using the Bayes Factor in Structural
2 Equation Models

Abstract

3
4 This Teacher's Corner paper introduces Bayesian evaluation of informative hypotheses for
5 structural equation models, using the free open-source R packages `bain`, for Bayesian
6 informative hypothesis testing, and `lavaan`, a widely used SEM package. The introduction
7 provides a brief non-technical explanation of informative hypotheses, the statistical
8 underpinnings of Bayesian hypothesis evaluation, and the `bain` algorithm. Three tutorial
9 examples demonstrate informative hypothesis evaluation in the context of common types of
10 structural equation models: 1) confirmatory factor analysis, 2) latent variable regression, and
11 3) multiple group analysis. We discuss hypothesis formulation, the interpretation of Bayes
12 factors and posterior model probabilities, and sensitivity analysis.

13 *Keywords:* `bain`; Bayes Factor; Informative Hypotheses; Structural Equation Modeling

14 Word count: 5948

15 Teacher’s Corner: Evaluating Informative Hypotheses Using the Bayes Factor in Structural
16 Equation Models

17 Hypotheses play a central role in deductive, theory-driven, research. A hypothesis
18 allows a researcher to draw inferences about a population, based on data sampled from that
19 population. In the context of structural equation modeling, there are two commonly used
20 approaches to hypothesis evaluation. Firstly, researchers can construct a set of competing
21 models, where each model represents several theoretically-derived substantive hypotheses.
22 Researchers can then use information criteria to select the best model in the set. Commonly
23 used information criteria include Akaike’s information criterion (AIC, Akaike, 1974), the
24 Bayesian information criterion (BIC, Schwarz, 1978), and the deviance information criterion
25 (DIC, Spiegelhalter, Best, Carlin, & Linde, 2002). Secondly, hypotheses about specific
26 parameters within a model can be tested by comparing a null hypothesis against an
27 alternative hypothesis using the likelihood ratio test (Wilks, 1938) or the Wald test (Buse,
28 1982).

29 A third approach is informative hypothesis evaluation (Hojtink, 2011). Informative
30 hypotheses are theoretically-derived statements about directional differences and equality
31 constraints between model parameters of interest. Informative hypotheses address an
32 important limitation of classical null-hypothesis significance testing: The null-hypothesis
33 that a parameter is equal to zero is often a “straw man” hypothesis. It holds little credibility,
34 and exists purely for the purpose of being rejected. The researcher’s actual theory, on the
35 other hand, is subsumed under a very broad alternative hypothesis, and is not directly
36 tested. The paradox inherent in this approach is that rejecting the straw man
37 null-hypothesis cannot be interpreted as evidence in support of the researcher’s theory, but
38 merely as evidence against the null. Informative hypotheses overcome this counter-intuitive
39 limitation, by explicitly testing a researcher’s theoretical beliefs.

40 Evaluating informative hypotheses is particularly straightforward from a Bayesian

41 perspective. Bayesian inference is already widely applied in the context of multivariate
42 normal linear models (see, for example: Well, Kolk, & Klugkist, 2008; Braeken, Mulder, &
43 Wood, 2015; Jong, Rigotti, & Mulder, 2017; Zondervan-Zwijnenburg et al., 2019). Methods
44 for Bayesian hypothesis evaluation within the structural equation modeling framework are
45 also available (Gu et al., 2019a; Van De Schoot, Hoijtink, Hallquist, & Boelen, 2012).
46 However, they are less frequently applied (but see Van Lissa, Hawk, Branje, Koot, & Meeus,
47 2016). This might be, in part, because user-friendly software was not available. In this
48 Teacher’s Corner paper, we show how Bayesian tests of informative hypotheses about
49 parameters in structural equation models can easily be conducted in **R**, using the **bain**
50 **package** (Gu et al., 2019b; Gu, Mulder, & Hoijtink, 2018; Hoijtink, Gu, & Mulder, 2018;
51 Mulder, 2014). From version 0.2.3 on, the package can evaluate informative hypotheses
52 about structural equation models estimated with the free, open-source SEM-package **lavaan**
53 (Loehlin & Beaujean, 2016; Rosseel, 2012). For a tutorial and technical details, see Hoijtink
54 et al. (2019b).

55 **Formulating Informative Hypotheses**

56 Informative hypotheses are formulated in terms of equality (=) and inequality (<, >)
57 constraints between target parameters. For example, one might hypothesize that one
58 regression coefficient is greater than the another, $H_1 : \beta_1 > \beta_2$, or that both are equal to a
59 specific value, $H_2 : (\beta_1, \beta_2) = 0.6$, or that one is greater than the other, which in turn is equal
60 to zero, $H_3 : \beta_1 > \beta_2 = 0$. The **bain** package uses a simple syntax to specify such hypotheses,
61 which is explained in detail in the **package vignette**. Here, we provide a brief overview of the
62 syntactical elements that are relevant in the context of structural equation models:

- 63 • **s1**, ..., **s6**: Refers to the target parameters **s1** up to **s6**. Substitute these with the
64 names of parameters in your model.
- 65 • **s1 = c**: An equality constraint, indicating that parameter **s1** is equal to constant **c**
- 66 • **s1 > c**: An inequality constraint, indicating that parameter **s1** is larger than constant

67 c

- 68 • $s1 = s2 = s3$: Three parameters have equal values.
- 69 • $(s1, s2, s3) > 0$: Three parameters, grouped by parentheses, are greater than zero.
- 70 • $c1 * s1 + c2 < c3 * s2 + c4$: A linear transformation of $s1$ (where a constant is
- 71 added to, or multiplied with, $s1$) is smaller than with a linear transformation of $s2$.
- 72 • $\dots \& \dots$: Within one hypothesis, the ampersand connects two constraints.
- 73 • $\dots ; \dots$: The `;` separates two distinct informative hypotheses.

74 When writing informative hypotheses about parameters of a `lavaan` model,
 75 parameters can be referenced by name. These names should be (unique abbreviations of) the
 76 parameter names used by `lavaan`. For example, `lavaan` labels the factor loading of the
 77 indicator `Ab` on the latent variable `A` as `A=~Ab`. This label, "`A=~Ab`", can be referenced
 78 verbatim in `bain` syntax, as in "`A=~Ab > .6`".

79 Note that comparing parameters (usually) makes sense only if they are on the same
 80 scale. For example, imagine that income is predicted by IQ and SES, where IQ is measured
 81 using a normed test ($M = 100, SD = 15$), and SES is rated on a 10-point ordinal scale
 82 which we treat as continuous. The regression coefficients for these predictors are β_{IQ} and
 83 β_{SES} , respectively. Since IQ and SES are measured on different scales, the hypothesis that
 84 $\beta_{IQ} < \beta_{SES}$ is meaningless. The unstandardized coefficients reflect both the strength of the
 85 relation of the predictors with income, and the scale with which the predictors were
 86 measured. The hypothesis does make sense with regard to the standardized model estimates,
 87 however. As a counterexample, if family income is predicted by maternal and paternal
 88 working hours, then the regression coefficients are on the same scale (dollars per hour of
 89 work), and can be directly compared. These examples illustrate that, except when
 90 comparing predictors measured on the same scale, or in other exceptional situations, it is
 91 usually safer to apply `bain` only to standardized model parameters.

92 Bayesian Hypothesis Evaluation

93 One of the key features of the Bayesian approach is that p-values, common to
94 null-hypothesis significance testing, are dispensed with. Instead, hypotheses are evaluated
95 using the *Bayes factor* (Kass & Raftery, 1995). The Bayes factor quantifies the relative
96 support provided by the data for two competing hypotheses. For example, let H_i be an
97 informative hypothesis that describes some (in)equality constraints among model parameters.
98 Let H_u be an unconstrained hypothesis that places no constraints on these model
99 parameters. The Bayes factor BF_{iu} , quantifies the support in favor of H_i relative to H_u . If
100 this Bayes factor BF_{iu} is larger than 1, the data provide more support for H_i than for H_u . If
101 it is smaller than 1, the data provide more support for H_u than for H_i . A Bayes factor near 1
102 is indecisive; both hypotheses are equally supported. The Bayes factor can be inverted to
103 express support in favor of H_u , relative to H_i . To this end, one can compute BF_{iu} as $1/\text{BF}_{ui}$.
104 Thus, if $\text{BF}_{iu} = 8.11$, then we can conclude that the data provide 8.11 times more support
105 for H_i than for H_u . Conversely, BF_{ui} (note that the order of the indices has changed) would
106 be $1/8.11 = .12$.

107 Since the Bayes factor is a relative measure of support, it should *not* be compared to a
108 threshold value. If, for example, $\text{BF}_{iu} = 102.75$ it is clear that the data provides
109 overwhelming support for H_i over H_u . With smaller values, such as $\text{BF}_{iu} = 7.34$, a
110 preference for H_i can still be defended, but other researchers might debate this preference,
111 and with even smaller values, such as $\text{BF}_{iu} = 3$, there is a preference for H_i , but H_u is
112 definitely not disqualified. Thus, the Bayes factor can, and should, be interpreted on a
113 continuous scale. This also sets it apart from the dichotomous decision making imposed by
114 the p-value. It is up to the scientific community to decide when “enough” evidence is
115 obtained to completely rule out a hypothesis. For a more elaborate discussion of Bayesian
116 hypothesis evaluation using `bain`, not specific to structural equation modeling, see the
117 tutorial by (Hojtink et al., 2019b).

118 **Statistical underpinnings.** The Bayes factor BF_{iu} can be written as a ratio of two
 119 marginal likelihoods of the hypotheses given the data (m), or alternatively, as the ratio of
 120 “fit” (f_i) and “complexity” (c_i , Gu, Mulder, & Hoijtink, 2018):

$$BF_{iu} = \frac{m(H_i | \text{data})}{m(H_u | \text{data})} = \frac{f_i}{c_i}.$$

121 The notion of fit reflects the extent to which the data is in agreement with the restrictions
 122 specified in the hypothesis, and its complexity reflects how specific the hypothesis is (Gu,
 123 Mulder, & Hoijtink, 2018). This ratio of fit and complexity is a concept that is also reflected
 124 in information criteria such as the AIC (Akaike, 1974) and the DIC (Spiegelhalter et al.,
 125 2002).

126 The **bain** algorithm estimates fit and complexity based on normal approximations of
 127 the prior and posterior distributions for the target parameters of the hypothesis. These
 128 distributions have a known mean and covariance matrix (Gu et al., 2018; Hoijtink et al.,
 129 2018). The posterior is defined by the observed parameter estimate and their asymptotic
 130 covariance matrix. For hypotheses with only inequality constraints, the fit (f_i) is then given
 131 by the proportion of this posterior distribution that is in agreement with the hypothesis (Gu
 132 et al., 2018; Hoijtink et al., 2018). For hypotheses with equality constraints, the fit is defined
 133 in terms of the posterior density at the constraints.

134 The prior distribution is constructed to provide an adequate quantification of
 135 complexity (see Gu, Mulder, & Hoijtink, 2018; Hoijtink, Gu, & Mulder, 2019). This is
 136 achieved by setting the prior mean along the boundary of the hypotheses under
 137 consideration. The prior covariance matrix is a scaling transformation of the posterior
 138 covariance matrix. Scaling increases the variances, leading to a flatter distribution. By
 139 default, **bain** scales the covariance matrix to be as flat as it would have been if it were based
 140 on the smallest possible sample required to estimate the target parameters. This is based on
 141 the concept of a minimal training sample (Berger & Pericchi, 2004; Mulder, 2014; O’Hagan,
 142 1995). Thus, the prior covariance matrix is much flatter, and therefore less informative, than

143 the posterior. The complexity (c_i) is given by the proportion (for inequality constrained
 144 hypotheses) or density (for equality constrained hypotheses) for the region of the prior
 145 distribution that is in agreement with the hypothesis.

146 **Evaluating a single informative hypothesis.** One way to evaluate a single
 147 informative hypothesis is to compare it with an “unconstrained” hypothesis, as
 148 in the preceding paragraphs. Let H_i signify any informative hypothesis that describes some
 149 (in)equality constraints among model parameters, such as $H_i : \beta_1 > \beta_2$, or $H_i : \beta_1 > \beta_2 = 0.6$.
 150 The unconstrained hypothesis H_u places no constraints on these model parameters:
 151 $H_u : \beta_1, \beta_2$. The Bayes factor BF_{iu} then quantifies the relative support provided by the data
 152 in favor of the informative hypothesis, relative to the unconstrained hypothesis - or in other
 153 words, how likely is it that the specified parameter constraints are true, relative to any other
 154 ordering of parameters. Throughout this paper, we use the notation $\text{BF}_{.u}$ to refer to Bayes
 155 factors of this type in the general sense, where $.$ signifies any informative hypothesis.

156 A second way to evaluate support in favor of an informative hypothesis, is to compare
 157 it to its complement. The complement is an alternative hypothesis that covers every ordering
 158 of parameter values that is *not* in line with the original hypothesis. If the informative
 159 hypothesis H_i expresses “the researcher’s theory”, and $!$ represents logical negation (“not”),
 160 then the complement $H_c : !H_i$ means “not the researcher’s theory”. Comparing against the
 161 complement allows researchers to investigate whether their expectation is, or is not,
 162 supported by the data. Bayes factor of the type $\text{BF}_{.c}$ indicate whether the data provide more
 163 support in favor of, or against, an informative hypothesis. In principle, the complement is
 164 defined by reference to a specific informative hypothesis, such that the complement of H_1 is
 165 $!H_1$, and the complement of H_2 is $!H_2$. For hypotheses with at least one equality constraint,
 166 however, the unconstrained hypothesis and the complement are the same. Since version
 167 0.2.4, `bain` reports both $\text{BF}_{.u}$ and $\text{BF}_{.c}$ by default.

168 It is worth pointing out that alternative, non-Bayesian methods exist that compare
 169 informative hypotheses against the null-hypothesis (Vanbrabant, Van De Schoot, Van Loey,

170 & Rosseel, 2017; Van De Schoot, Hoijtink, & Deković, 2010). When using `bain`, it is also
 171 possible to evaluate the null-hypothesis by specifying it as an informative hypothesis (i.e., a
 172 hypothesis that constrains all parameters to be equal to zero, or to be equal to one another),
 173 and comparing it with other informative hypotheses using the approach elaborated in the
 174 next paragraph.

175 **Comparing two informative hypotheses.** A second question researchers might
 176 want to address, is which of two informative hypotheses, H_1 and H_2 , is most supported by
 177 the data. The Bayes factor BF_{12} reflects the amount of support provided by the data in
 178 favor of H_1 , relative to H_2 . It is computed by taking a ratio of two other Bayes factors:

$$BF_{12} = \frac{BF_{1u}}{BF_{2u}}$$

179 This approach is valid because Bayes factors for any two informative hypotheses can be
 180 compared if both have the same denominator. In the previous section, we explained that it is
 181 not possible to compare Bayes factors of the type $BF_{.c}$, because the complement of H_1 is not
 182 the same as that of H_2 . However, Bayes factors of the type $BF_{.u}$ are comparable, because
 183 the unconstrained hypothesis is identical for all informative hypotheses. Thus, BF_{12} can be
 184 computed to contrast a pair of user-specified informative hypotheses.

185 By default, `bain` will compute Bayes factors to contrast all informative hypotheses.
 186 Thus, given three hypotheses, $H_1 : \beta_1 = \beta_2 = \beta_3 = 0$, $H_2 : \beta_1 > 0 \ \& \ \beta_2 > 0 \ \& \ \beta_3 > 0$, and
 187 $H_3 : \beta_1 > \beta_2 > \beta_3 > 0$, `bain` will compute BF_{12} , BF_{13} , and BF_{23} . These Bayes factors are
 188 stored in the `$BFmatrix` element of the output.

189 **Comparing more than two hypotheses.** Any two informative hypotheses can be
 190 straightforwardly compared using the method outlined above. When there are more than
 191 two candidate hypotheses, however, comparing all of their mutual Bayes factors quickly
 192 becomes cumbersome. In this case, it is easier to compare the so-called posterior model
 193 probabilities for each hypothesis H_i , that is, $P(H_i|\text{data})$. Each posterior model probability
 194 has a value between 0 and 1, and the posterior model probabilities for a set of hypotheses

195 sum to 1.0. Under the assumption that, a priori (before observing the data), each hypothesis
196 is equally likely, the posterior model probabilities contain the same information as the Bayes
197 factors upon which they are based. If, for example, $\text{BF}_{12} = 3.5$, $\text{BF}_{13} = 7.0$ and $\text{BF}_{23} = 2.0$,
198 the corresponding posterior model probabilities are $P(H_1|\text{data}) = .7$, $P(H_2|\text{data}) = .2$, and
199 $P(H_3|\text{data}) = .1$, respectively. Note that, $\text{BF}_{12} = \frac{P(H_1|\text{data})}{P(H_2|\text{data})} = \frac{.7}{.2} = 3.5$. Posterior model
200 probabilities can also be interpreted as Bayesian error probabilities. If the set of hypotheses
201 under consideration contains H_1 , H_2 , and H_3 , and the corresponding posterior model
202 probabilities are .7, .2, and .1, respectively, then the Bayesian error probability associated
203 with a preference for H_1 is equal to $.2 + .1 = .3$.

204 **A fail-safe hypothesis.** It is important to emphasize that posterior model
205 probabilities only indicate which of the hypotheses *in the set* receives the most support from
206 the data. Consequently, if all of the hypotheses in the set misrepresent the true relationship
207 among parameters in the population, then researchers risk selecting the best of a set of “bad”
208 hypotheses. Two approaches can be used to mitigate this risk. The first approach uses the
209 unconstrained hypothesis H_u as a “fail-safe” hypothesis. Recall that H_u places no
210 constraints on the parameters. If the best hypothesis in the set receives more support than
211 the unconstrained hypothesis, we are reassured that it is not just the best of a set of bad
212 hypotheses. This approach is currently implemented in `bain`. The second approach would be
213 to include a hypothesis that is the complement of the union of all informative hypotheses in
214 the set. A nice feature of this second approach is that, whereas H_u overlaps with each of the
215 hypotheses under consideration, the complement of the union does not. However, as to yet
216 this option is not implemented in `bain`.

217 Structural Equation Modeling Using `lavaan`

218 In this paper, we present a subset of the (multiple group) structural equation models
219 that can be specified using the `lavaan` function `sem`, and for which informative hypotheses
220 can be formulated and processed with `bain`. The interested reader is advised to visit

221 <http://lavaan.org/>, where mini-tutorials and examples are used to explain all the functions
222 and options available in the `lavaan` package. For a general introduction to structural
223 equation modeling, the interested reader is referred to Loehlin and Beaujean (2016). As will
224 be elaborated upon in the discussion, it is relatively easy to use `bain` for the evaluation of
225 hypotheses for all models that can be specified in `lavaan`.

226 When used in conjunction with `lavaan`, `bain` extracts the (standardized or
227 unstandardized) target parameter estimates (per group), the covariance matrix of the
228 estimates (per group) and the sample size (per group) from the `lavaan` output object.
229 Target parameters are defined as model parameters about which informative hypotheses are
230 formulated. By contrast, nuisance parameters are parameters not involved in the hypotheses
231 of interest. `bain` is validated for use with target parameters that are either 1) regression
232 coefficients, 2) intercepts, or 3) factor loadings. Thus, by default, all (residual) (co)variances
233 are treated as nuisance parameters, along with any remaining parameters not involved in the
234 hypotheses.

235 A final note regarding assumptions: As explained earlier, `bain` constructs a default
236 prior distribution for the target parameters (per group), and derives a normal approximation
237 of the posterior. Asymptotically, the posterior distribution is indeed normal (see, for
238 example, Gelman et al., 2013, Chapter 4). However, `bain` should only be used if
239 approximate normality can be assumed, given the sample size. Rosseel (2020) provides
240 references that validate the use of structural equation modeling when the sample size is at
241 least 200. This approximate prior and posterior form the basis for the computation of Bayes
242 factors for the informative hypotheses. A more detailed accessible introduction is presented
243 in Hoijsink et al. (2019b), and the statistical underpinnings of the method are substantiated
244 in Gu et al. (2018) and Hoijsink et al. (2018).

Tutorial examples

We present tutorial examples for three commonly used types of structural equation models: 1) confirmatory factor analysis, 2) latent variable regression, and 3) multiple group analysis. Each example follows a three-step workflow. In the first step, `lavaan` is used to estimate the parameters of a structural equation model. In the second step, one or more informative hypotheses are formulated. In the third step, the results of the `lavaan` analysis and the hypotheses are fed into `bain`, which renders a Bayesian evaluation of the hypotheses, returning Bayes factors and posterior model probabilities.

All examples use the synthetic data set “`sesamesim`”, which is included with the `bain` package. These data are generated to have similar distributional characteristics and covariances to the Sesame Street data provided by Stevens (2012). These data concern the effect of watching the tv-series “Sesame Street” for one year on the knowledge of numbers of 240 children aged between 34 to 69 months. We will use the following variables: Age in months (`age`), the Peabody test, which measures the “mental age” of children (`peabody`; score range 15 to 89), and `sex`, with boys coded as 1, and girls as 2. Several variables were measured both before- and after watching Sesame Street for one year: Knowledge of numbers (`Bn`: before, and `An`: after); knowledge of body parts (`Bb` and `Ab`, respectively), letters (`Bl` and `Al`), forms (`Bf` and `Af`), relationships (`Br` and `Ar`), and classifications (`Bc` and `Ac`). Models are fit using `lavaan`, and Figures are plotted using `tidySEM` (Van Lissa, 2020).

Example 1: Confirmatory Factor Analysis

A two-factor confirmatory factor analysis is specified using the syntax below, in which the A(fter) measurements of all subtests load on the factor A, and the B(efore) measurements load on the factor B (see Figure 1).

```
model1 <- 'A =~ Ab + Al + Af + An + Ar + Ac
          B =~ Bb + Bl + Bf + Bn + Br + Bc'
```

```
fit1 <- sem(model1, data = sesamesim, std.lv = TRUE)
```

268 The argument `std.lv = TRUE` implies that the model is identified by standardizing
 269 the latent variables B and A. This allows the formulation of informative hypotheses with
 270 respect to each of the factor loadings, including the first.

271 **Specifying informative hypotheses.** One plausible hypothesis for this
 272 confirmatory factor analysis might be that indicators are strongly related to the factors to
 273 which they are assigned. This is reflected by the following hypothesis, which states that all
 274 (standardized) factor loadings are larger than .6:

```
hypotheses1 <- "(A=~Ab, A=~Al, A=~Af, A=~An, A=~Ar, A=~Ac) >.6 &
                (B=~Bb, B=~Bl, B=~Bf, B=~Bn, B=~Br, B=~Bc) >.6"
```

275 This example consists of one hypothesis about two groups of parameters, enclosed by
 276 parentheses, which are chained by the ampersand symbol. Note that, although we could
 277 group all loadings between brackets, before and after are separated for clarity. In this
 278 example, the target parameters are factor loadings, the sample size is $N = 240$, and therefore,
 279 we assume that the posterior distribution of the target parameters is approximately normal.

280 **Evaluating hypotheses.** Now, we will evaluate the informative hypotheses for this
 281 example using `bain()`. As input to the function, we use the `lavaan` output object `fit1` and
 282 the hypotheses `hypotheses1` that were specified above. The argument `standardize =`
 283 `TRUE` ensures that the hypotheses are evaluated in terms of standardized model parameters.

284 Before calling `bain()`, we set a seed for the random number generator using
 285 `set.seed()`. This is necessary to ensure computational replicability, because `bain` draws
 286 random samples from the prior and posterior distributions of the target parameters. If
 287 another seed is used, a different random sample will be drawn, which could lead to
 288 differences in the resulting Bayes factors and posterior model probabilities. These differences

289 should be negligible, and it is good practice to conduct a sensitivity analysis for Monte Carlo
290 error (the variability due to different random seeds) by changing the seed to ensure that the
291 results are replicated.

```
set.seed(100)
results1 <- bain(fit1, hypotheses1, standardize = TRUE)
results1
```

292 The resulting `bain()` output is presented in Table 1. The Bayes factor BF_{1c} , which
293 compares H_1 to its complement, is found on the row for H_1 , in column BF.c. As can be seen,
294 $BF_{1c} = 93.33$, that is, the data offers overwhelming support in favor of H_1 . This is not
295 surprising when we examine the parameter estimates and their 95% central credible intervals
296 using the `summary()` function (see Table 2).

```
summary(results1)
```

297 In agreement with H_1 , all observed standardized loadings are larger than .6. Note that,
298 a preference for H_1 compared to H_u comes with a Bayesian error probability of .01: A 1%
299 probability that the choice for H_1 is incorrect, conditional on the set of models (see Table 1).

300 Example 2: Latent Regression

301 A latent regression model is specified using the code below. The measurement model
302 for the factors B and A is the same as in Example 1. In this example, however, the
303 correlation from the preceding example is replaced by a regression coefficient. Moreover, `age`
304 and `peabody` are included as observed covariates. This analysis thus allows us to investigate
305 whether children's knowledge after watching Sesame Street for a year is predicted by their
306 knowledge one year before, as well as by their biological- and mental age.

```

model2 <- 'A =~ Ab + Al + Af + An + Ar + Ac
          B =~ Bb + Bl + Bf + Bn + Br + Bc
          A ~ B + age + peabody'

fit2 <- sem(model2, data = sesamesim, std.lv = TRUE)

```

307 **Specifying informative hypotheses.** This example contains three hypotheses,
 308 separated by semicolons, regarding the relative importance of `B`, `age`, and `peabody` when
 309 predicting `A`:

```

hypotheses2 <- "A~B > A~peabody = A~age = 0;
               A~B > A~peabody > A~age = 0;
               A~B > A~peabody > A~age > 0"

```

310 H_1 specifies that the regression coefficient of `B` on `A` is greater than zero, and that the
 311 coefficients of `age` and `peabody` on `A` are equal to zero. H_2 specifies that the regression
 312 coefficient of `B` on `A` is greater than that of `peabody` on `A`, which in turn is bigger than that
 313 of `age` on `A`, which is equal to zero. H_3 specifies that the coefficient of `B` on `A` is greater than
 314 that of `peabody` on `A`, which, in turn, is greater than that of `age` on `A`, which is greater than
 315 zero.

316 **Evaluating hypotheses.** The code below evaluates the hypotheses specified for the
 317 latent regression example:

```

set.seed(748)

results2 <- bain(fit2, hypotheses2, standardize = TRUE)

```

318 The results are reported in Table 3. When H_1 , H_2 , and H_3 are compared to their
 319 respective complements, there is substantial support for H_1 , somewhat less for H_2 , and
 320 substantially less support for H_3 . The posterior model probabilities, PMPb, help determine

321 which of the three informative hypotheses is the best of the set, and whether the
322 unconstrained hypothesis H_u holds any credulity. Supported by a posterior model
323 probability of .79, H_1 appears to be the best of the set of hypotheses. However, a choice for
324 H_1 implies a Bayesian error probability of $.17 + .03 + .01 = .21$, that is, it would be unwise
325 to ignore the possibility that another hypothesis (especially H_2) might also be a good
326 candidate. It is clear that the regression coefficient of **B** is larger than zero, but maybe the
327 regression coefficient of **peabody** is also larger than zero. We can see how these findings
328 relate to the model parameters by calling `summary()` on the `bain` object (see Table 4).

329 **Example 3: Multiple Group Analysis**

330 This example demonstrates how to evaluate informative hypotheses about freely
331 estimated parameters across groups in a multi-group structural equation model. It is
332 important to emphasize that the Bayes factor implemented in `bain` is only valid for multiple
333 group models without any between-group parameter constraints. The reason is that `bain`
334 requires a separate asymptotic covariance matrix for the parameters of each group. This is
335 only possible when no between-group constraints are imposed, because then (and only then)
336 is the asymptotic covariance matrix block-diagonal, and can we extract a covariance matrix
337 per group. For more information, see Hoijsink, Gu, and Mulder (2018). A multiple group
338 model can be estimated by specifying a grouping variable in the call to `sem`. The code below
339 runs an analysis in which the parameters of a regression model are estimated separately for
340 boys and girls. The model predicts knowledge of numbers after watching Sesame Street for a
341 year based on prior knowledge of numbers, and the peabody mental age test (see Figure 2).

```
model3 <- ' postnumb ~ prenumb + peabody '  
# Assign labels to the groups to be used when formulating hypotheses  
sesamesim$sex <- factor(sesamesim$sex, labels = c("boy", "girl"))  
# Fit the multiple group structural equation model  
fit3 <- sem(model3, data = sesamesim, group = "sex")
```


342 **Specifying informative hypotheses.** For the multiple group (boys versus girls)
 343 structural equation model, we evaluate two hypotheses: That standardized regression
 344 coefficients are equal for boys and girls (H_1), or that they are smaller for boys as compared
 345 to girls (H_2). In other words, are number knowledge before and the peabody test better
 346 predictors of number knowledge after for girls than for boys?

```
hypotheses3 <- "postnumb~prenumb.boy = postnumb~prenumb.girl &
                postnumb~peabody.boy = postnumb~peabody.girl;
                postnumb~prenumb.boy < postnumb~prenumb.girl &
                postnumb~peabody.boy < postnumb~peabody.girl"
```

347 **Evaluating hypotheses.** The results, displayed in Table 5, indicate that H_1
 348 receives 41.20 times more support from the data than its complement. Conversely, H_2
 349 received $1/.16 = 6.25$ times less support than its complement. These results indicate that the
 350 predictability of “postnumb” does not depend on gender. This is also reflected by the
 351 posterior model probabilities that show that a decision in favor of H_1 comes with a Bayesian
 352 error probability of only 0.02.

```
set.seed(235)
results3 <- bain(fit3, hypotheses3, standardize = TRUE)
```

353 This conclusion is corroborated by the model coefficients, obtained by running
 354 `summary(results3)`. As seen in Table 6, the credible intervals for the regression coefficients
 355 for boys and girls show substantial overlap.

356 Further extensions

357 Sensitivity Analysis

358 Bayes factors for hypotheses containing at least one equality constraint are sensitive to
 359 the scaling factor used to construct the prior distribution. Recall that the default scaling

360 factor in `bain` is based on the notion of a minimal training sample; the smallest sample size
361 required to estimate the target parameters. This default scaling factor is set by the default
362 argument `fraction = 1` in the call to `bain()`. A default argument does not need to be
363 specified, but can be changed manually by specifying a different value. The smallest possible
364 scaling factor is the default, 1. Larger scaling factors increase confidence in the prior, making
365 it more concentrated and less spread out. Thus, specifying `fraction = 2` raises the scaling
366 factor to twice the size of the minimal training sample, and `fraction = 3` to thrice the size.

367 The reason hypotheses containing at least one equality constraint are sensitive to the
368 scaling factor is that equality constraints are represented as a fixed-width “slice” of the
369 parameter space around the constraint value (in technical terms, the point density at this
370 value). If the width of the prior changes, the ratio of the fixed-width slice to the overall width
371 of the prior changes. Hypotheses specified using only inequality constraints are not sensitive
372 to the scaling factor, because these constraints divide the parameter space (like cutting the
373 distribution into two halves). As the width of the prior changes, the space on both sides of
374 the constraint decreases commensurately, so their ratio remains the same (see Hoijtink et al.,
375 2019b for a full explanation).

376 It is possible to conduct a sensitivity analysis to examine how sensitive the Bayes
377 factors are to the scaling factor. The convenience function `bain_sensitivity()` accepts a
378 vector argument called `fractions = ...`, and returns a list of `bain` objects. The
379 `summary()` function for this sensitivity analysis accepts an argument `which_stat`, that can
380 be used to request a sensitivity analysis table for a specific statistic (by default, this is the
381 “BF”). Below, we demonstrate how to conduct a sensitivity analysis, based on Example 2:

```
set.seed(753)

results_sens <- bain_sensitivity(fit2, hypotheses2, fractions = c(1, 2, 3),
                               standardize = TRUE)

summary(results_sens)
```

382 The results are presented in Table 7. It shows that the value of BF_{3c} is invariant,
383 whereas BF_{1u} and BF_{2u} decrease as the scaling factor increases. The posterior model
384 probabilities change accordingly, as can be seen in Table 8.

```
summary(results_sens, which_stat = "PMPb")
```

385 The remaining question is how to deal with the sensitivity of the Bayes factor to the
386 scale factor. There are three potential courses of action. Firstly, if all hypotheses under
387 consideration are formulated using only inequality constraints, the Bayes factors are
388 invariant, as can be seen from BF_{3c} in Table 8. Secondly, if the hypotheses contain equality
389 constraints, researchers can rely on the default scaling factor implemented in `bain`. The
390 resulting Bayes factors tend to favor hypotheses with equality constraints over their
391 complement. This approach ensures that the evidence in the data has to be compelling
392 before it is concluded that the constraints do not hold. When applied to null-hypotheses (i.e.,
393 an equality constrained hypothesis stating that a parameter is equal to zero), this
394 conservative approach curtails the false positive rate. This is appropriate, especially in the
395 context of the replication crisis (see, for example, Open Science Collaboration, 2015).
396 Thirdly, researchers can execute a sensitivity analysis, as in the preceding example:
397 Empirically investigate the sensitivity of the Bayes factors to the scaling factor, and report
398 the results. In our experience, conclusions are usually robust with respect to different values
399 of the scaling factor. This can also be seen in Table 7: Although the Bayes factor for H_1
400 decreases from 150.87 to 50.29, the conclusion remains that H_1 is substantially more
401 supported than H_c . Furthermore, in terms of posterior model probabilities, the conclusion
402 remains that H_1 is the best hypothesis, and that H_2 cannot be ruled out.

403 **Experimental applications**

404 The examples above all use the standard interface of the `bain()` function, which
405 requires two arguments: A model object, and a hypothesis. This interface accepts all `lavaan`

406 model objects generated by the functions `cfa`, `sem`, and `growth`. Within these models,
407 parameters may be fixed, and data may be categorical, and hypotheses can be formulated
408 with respect to intercepts, factor loadings, and regression coefficients. Some situations that
409 cannot currently be handled by `bain` include multilevel models (specified using the `cluster`
410 argument), and defined parameters, such as indirect effects in mediation models. If a
411 researcher wishes to circumvent the standard user interface, `bain()` can be applied to a
412 named vector of parameters, instead of one of one of the model types for which methods
413 exist. This approach calls the default method of `bain`, which is less user-friendly, but more
414 flexible than the model-specific interface. Section 4.i in the `bain` package vignette illustrates
415 this approach, and demonstrates how to manually extract the target parameter estimates
416 and place them in a named vector, and how to obtain the parameter covariance matrix and
417 sample size from a `lavaan` object. This vignette can be loaded by calling
418 `vignette("bain_introduction", package = "bain")`. Note that non-standard
419 applications of `bain` that have not yet been validated should be identified as such, or
420 substantiated with a simulation study.

421

Discussion

422 This Teacher's Corner paper introduced Bayesian hypotheses evaluation for structural
423 equation models using `bain` and `lavaan`. The combination of both R packages enables the
424 free, open-source, and user friendly evaluation of informative hypotheses for structural
425 equation models. The approach elaborated in this paper uses Bayes factors, which are a
426 measure of relative support for two hypotheses. The interpretation of Bayes factors is
427 straightforward: It is a ratio of evidence in favor of one hypothesis, relative to evidence in
428 favor of another hypothesis. Bayes factors can be indecisive; the closer Bayes factors get to
429 one, the less differential support was found for either hypothesis. It is up to the scientific
430 community to decide how much evidence is sufficient evidence.

431 The advocated approach allows users to evaluate support for a single informative

432 hypothesis, either relative to its complement, or relative to an unconstrained hypothesis.
433 The Bayes factor $BF_{.c}$ compares against the complement, and expresses how much evidence
434 the data provides is in favor of “the theory”, as compared to “not the theory”. The Bayes
435 factor $BF_{.u}$ compares against the unconstrained hypothesis, and expresses how much
436 evidence the data provides is in favor of “the theory”, as compared to “any ordering of
437 parameters”. Two informative hypotheses can be compared by computing their joint Bayes
438 factor, which is a ratio of the two $BF_{.u}$ s for these hypotheses.

439 When simultaneously evaluating more than two hypotheses, it is convenient to use the
440 posterior model probabilities. These quantify the proportion of support for each hypothesis
441 in a set, conditional on the data. This was illustrated in Example 2. Bayesian error
442 probabilities additionally quantify the uncertainty of decisions about hypotheses. The
443 probability that a preference for one hypothesis in the set is incorrect, is equal to the sum of
444 posterior model probabilities for the other informative hypotheses. This is a conditional
445 probability, that is, conditional on the available data and the hypotheses in the set.

446 Structural equation models are often estimated on data that contain missing values.
447 Fortunately, the Bayes factor implemented in `bain` can also be computed if the data contain
448 missing values (Gu et al., 2019a; Hoijtink et al., 2019a). Users can use multiple imputation
449 (Van Buuren, 2018) to obtain estimates of the (standardized) target parameters, their
450 covariance matrix, and the effective sample size, and once those are available, `bain` be used
451 for the evaluation of informative hypotheses. The interested reader is referred to the vignette
452 included with the `bain` package, which includes an elaborate example.

453 Several potential limitations remain. One such limitation is the fact that `bain` utilizes
454 normal approximations of the prior and posterior distribution. This could have implications
455 for quantities whose sampling distribution is known to be non-normally distributed, such as
456 indirect effects (MacKinnon, Lockwood, & Williams, 2004). However, this problem is averted
457 by the fact that users are currently prevented from using the `lavaan` interface to `bain` for

458 derived parameters, which includes indirect effects. A second limitation is the fact that **bain**
459 cannot handle multiple group models with between-group constraints. Substantial future
460 research is required to overcome this issue. An implication of this limitation is that it is not
461 possible to impose measurement invariance in multiple group latent variable models. One
462 potential solution, that can already be applied, is to use linear transformations within the
463 **bain** hypotheses to ensure that parameters are comparable across groups. However, this
464 procedure is complicated and beyond the scope of this tutorial. Pending a future publication
465 addressing measurement invariance, researchers can contact the authors to obtain support
466 for such analyses.

467 In conclusion, **bain** enables user-friendly Bayesian evaluation of informative hypotheses
468 for structural equation models estimated in **lavaan**. The method has been validated for
469 regression coefficients, factor loadings, and intercepts, in a range of commonly specified
470 structural equation models, such as factor analyses, latent regression analyses, multi-group
471 models, and latent growth models. Its functionality will be further expanded in future
472 updates, and the default method for named vectors offers the freedom to explore applications
473 not currently covered by the standard interface.

References

- 474
- 475 Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on*
476 *Automatic Control*, 19(6), 716–723. <https://doi.org/10.1109/TAC.1974.1100705>
- 477 Berger, J. O., & Pericchi, L. R. (2004). Training samples in objective Bayesian model
478 selection. *The Annals of Statistics*, 32(3), 841–869.
479 <https://doi.org/10.1214/009053604000000229>
- 480 Braeken, J., Mulder, J., & Wood, S. (2015). Relative effects at work: Bayes factors for order
481 hypotheses. *Journal of Management*, 41(2), 544–573.
482 <https://doi.org/10.1177/0149206314525206>
- 483 Buse, A. (1982). The Likelihood Ratio, Wald, and Lagrange Multiplier Tests: An Expository
484 Note. *The American Statistician*, 36(3a), 153–157.
485 <https://doi.org/10.1080/00031305.1982.10482817>
- 486 Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013).
487 *Bayesian data analysis*. Chapman and Hall/CRC.
- 488 Gu, X., Hoijsink, H., Mulder, J., & Rosseel, Y. (2019a). Bain: A program for Bayesian
489 testing of order constrained hypotheses in structural equation models. *Journal of*
490 *Statistical Computation and Simulation*, 89(8), 1526–1553.
491 <https://doi.org/10.1080/00949655.2019.1590574>
- 492 Gu, X., Hoijsink, H., Mulder, J., & van Lissa, C. (2019b). Bain: Bayes factors for
493 informative hypotheses. (Version 0.2.3). Retrieved from
494 <https://CRAN.R-project.org/package=bain>
- 495 Gu, X., Mulder, J., & Hoijsink, H. (2018). Approximated adjusted fractional Bayes factors:
496 A general method for testing informative hypotheses. *British Journal of Mathematical*
497 *and Statistical Psychology*, 71(2), 229–261. <https://doi.org/10.1111/bmsp.12110>
- 498 Hoijsink, H. (2011). *Informative hypotheses: Theory and practice for behavioral and social*

- 499 *scientists*. Chapman and Hall/CRC.
- 500 Hoijtink, H., Gu, X., & Mulder, J. (2018). Bain, multiple group Bayesian evaluation of
501 informative hypotheses. *British Journal of Mathematical and Statistical Psychology*.
502 DOI: 10.1111/Bmsp, 12145.
- 503 Hoijtink, H., Gu, X., Mulder, J., & Rosseel, Y. (2019a). Computing Bayes factors from data
504 with missing values. *Psychological Methods*, 24(2), 253–268.
505 <https://doi.org/10.1037/met0000187>
- 506 Hoijtink, H., Mulder, J., van Lissa, C., & Gu, X. (2019b). A tutorial on testing hypotheses
507 using the Bayes factor. *Psychological Methods*, 24(5), 539–556.
508 <https://doi.org/10.1037/met0000201>
- 509 Jong, J. de, Rigotti, T., & Mulder, J. (2017). One after the other: Effects of sequence
510 patterns of breached and overfulfilled obligations. *European Journal of Work and*
511 *Organizational Psychology*, 26(3), 337–355.
512 <https://doi.org/10.1080/1359432X.2017.1287074>
- 513 Kass, R. E., & Raftery, A. E. (1995). Bayes Factors. *Journal of the American Statistical*
514 *Association*, 90(430), 773–795. <https://doi.org/10.1080/01621459.1995.10476572>
- 515 Loehlin, J. C., & Beaujean, A. A. (2016). *Latent variable models: An introduction to factor,*
516 *path, and structural equation analysis*. Taylor & Francis.
- 517 MacKinnon, D. P., Lockwood, C. M., & Williams, J. (2004). Confidence Limits for the
518 Indirect Effect: Distribution of the Product and Resampling Methods. *Multivariate*
519 *Behavioral Research*, 39(1), 99–128. https://doi.org/10.1207/s15327906mbr3901_4
- 520 Mulder, J. (2014). Prior adjusted default Bayes factors for testing (in)Equality constrained
521 hypotheses. *Computational Statistics & Data Analysis*, 71, 448–463.
522 <https://doi.org/10.1016/j.csda.2013.07.017>
- 523 O’Hagan, A. (1995). Fractional Bayes Factors for Model Comparison. *Journal of the Royal*

- 524 *Statistical Society: Series B (Methodological)*, 57(1), 99–118.
525 <https://doi.org/10.1111/j.2517-6161.1995.tb02017.x>
- 526 Rosseel, Y. (2012). Lavaan: An R package for structural equation modeling and more
527 Version 0.5-12 (BETA), 37.
- 528 Rosseel, Y. (2020). Small sample solutions for structural equation modeling. In R. Van De
529 Schoot & M. Miočević (Eds.), *Small sample size solutions: A guide for applied*
530 *researchers and practitioners*. CRC Press.
- 531 Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2),
532 461–464. <https://doi.org/10.1214/aos/1176344136>
- 533 Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Linde, A. V. D. (2002). Bayesian measures
534 of model complexity and fit. *Journal of the Royal Statistical Society: Series B*
535 *(Statistical Methodology)*, 64(4), 583–639. <https://doi.org/10.1111/1467-9868.00353>
- 536 Stevens, J. P. (2012). *Applied multivariate statistics for the social sciences*. Routledge.
- 537 Vanbrabant, L., Van De Schoot, R., Van Loey, N., & Rosseel, Y. (2017). A General
538 Procedure for Testing Inequality Constrained Hypotheses in SEM. *Methodology*,
539 13(2), 61–70. <https://doi.org/10.1027/1614-2241/a000123>
- 540 Van Buuren, S. (2018). *Flexible imputation of missing data*. Chapman and Hall/CRC.
- 541 Van De Schoot, R., Hoijsink, H., & Deković, M. (2010). Testing Inequality Constrained
542 Hypotheses in SEM Models. *Structural Equation Modeling: A Multidisciplinary*
543 *Journal*, 17(3), 443–463. <https://doi.org/10.1080/10705511.2010.489010>
- 544 Van De Schoot, R., Hoijsink, H., Hallquist, M. N., & Boelen, P. A. (2012). Bayesian
545 Evaluation of Inequality-Constrained Hypotheses in SEM Models Using Mplus.
546 *Structural Equation Modeling: A Multidisciplinary Journal*, 19(4), 593–609.
547 <https://doi.org/10.1080/10705511.2012.713267>
- 548 Van Lissa, C. J. (2020). tidySEM: Generate tidy SEM-syntax (Version 0.1.0.5). Retrieved

- 549 from www.github.com/cjvanlissa/tidySEM
- 550 Van Lissa, C. J., Hawk, S. T., Branje, S., Koot, H. M., & Meeus, W. H. J. (2016). Common
551 and unique associations of adolescents' affective and cognitive empathy development
552 with conflict behavior towards parents. *Journal of Adolescence*, *47*, 60–70.
553 <https://doi.org/10.1016/j.adolescence.2015.12.005>
- 554 Well, S. van, Kolk, A. M., & Klugkist, I. G. (2008). Effects of Sex, Gender Role
555 Identification, and Gender Relevance of Two Types of Stressors on Cardiovascular
556 and Subjective Responses: Sex and Gender Match and Mismatch Effects. *Behavior*
557 *Modification*. <https://doi.org/10.1177/0145445507309030>
- 558 Wilks, S. S. (1938). The Large-Sample Distribution of the Likelihood Ratio for Testing
559 Composite Hypotheses. *The Annals of Mathematical Statistics*, *9*(1), 60–62.
560 Retrieved from <https://www.jstor.org/stable/2957648>
- 561 Zondervan-Zwijnenburg, M. A. J., Veldkamp, S. A. M., Neumann, A., Barzeva, S. A.,
562 Nelemans, S. A., Beijsterveldt, C. E. M. van, . . . Boomsma, D. I. (2019). Parental
563 Age and Offspring Childhood Mental Health: A Multi-Cohort, Population-Based
564 Investigation. *Child Development*. <https://doi.org/10.1111/cdev.13267>

Table 1

*Bain output for the Confirmatory Factor Analysis**Model*

	Fit	Com	BF.u	BF.c	PMPa	PMPb
H1	0.88	0.01	93.33	762.05	1.00	0.99
Hu						0.01

Note. PMP: Posterior model probabilities; PMPa excludes Hu; PMPb includes Hu.

Table 2

Standardized parameter estimates for the confirmatory factor analysis

Parameter	Estimate A	CI A	Estimate B	CI B
b	0.71	[0.64, 0.78]	0.77	[0.71, 0.82]
l	0.81	[0.76, 0.86]	0.65	[0.57, 0.73]
f	0.84	[0.79, 0.88]	0.81	[0.76, 0.86]
n	0.91	[0.88, 0.94]	0.89	[0.85, 0.92]
r	0.70	[0.63, 0.77]	0.72	[0.65, 0.79]
c	0.87	[0.84, 0.91]	0.83	[0.78, 0.87]

Note. CI: Credible interval.

Table 3

Bain output for the latent regression model

	Fit	Com	BF.u	BF.c	PMPa	PMPb
H1	69.90	0.46	150.87	150.87	0.80	0.79
H2	2.96	0.09	33.16	33.16	0.17	0.17
H3	0.07	0.01	5.64	6.01	0.03	0.03
Hu						0.01

Note. PMP: Posterior model probabilities; PMPa excludes Hu; PMPb includes Hu.

Table 4

Standardized parameter estimates for latent regression

	Parameter	Estimate	CI
13	A~B	0.79	[0.73, 0.85]
14	A~age	0.00	[-0.09, 0.09]
15	A~peabody	-0.02	[-0.11, 0.08]

Note. CI: Credible interval.

Table 5

Bain output for the latent regression model

	Fit	Com	BF.u	BF.c	PMPa	PMPb
H1	7.79	0.19	41.20	41.20	1.00	0.97
H2	0.02	0.11	0.18	0.16	0.00	0.00
Hu						0.02

Note. PMP: Posterior model probabilities; PMPa excludes Hu; PMPb includes Hu.

Table 6

Parameter estimates for the multiple group model

Parameter	Estimate boy	CI boy	Estimate girl	CI girl
postnumb~prenumb	0.53	[0.38, 0.68]	0.64	[0.52, 0.76]
postnumb~peabody	0.23	[0.07, 0.40]	0.06	[-0.09, 0.22]

Note. CI: Credible interval.

Table 7

Sensitivity analysis for the Bayes factors (BF) of the latent regression model

Fraction	H1	H2	H3
1.00	150.87	36.69	6.84
2.00	75.44	25.94	6.84
3.00	50.29	21.18	6.84

Table 8

Sensitivity analysis for posterior model probabilities (PMPb) of the multiple group model

Fraction	H1	H2	H3	Hu
1.00	0.77	0.19	0.03	0.01
2.00	0.69	0.24	0.06	0.01
3.00	0.64	0.27	0.08	0.01




Figure 1. Confirmatory factor analysis

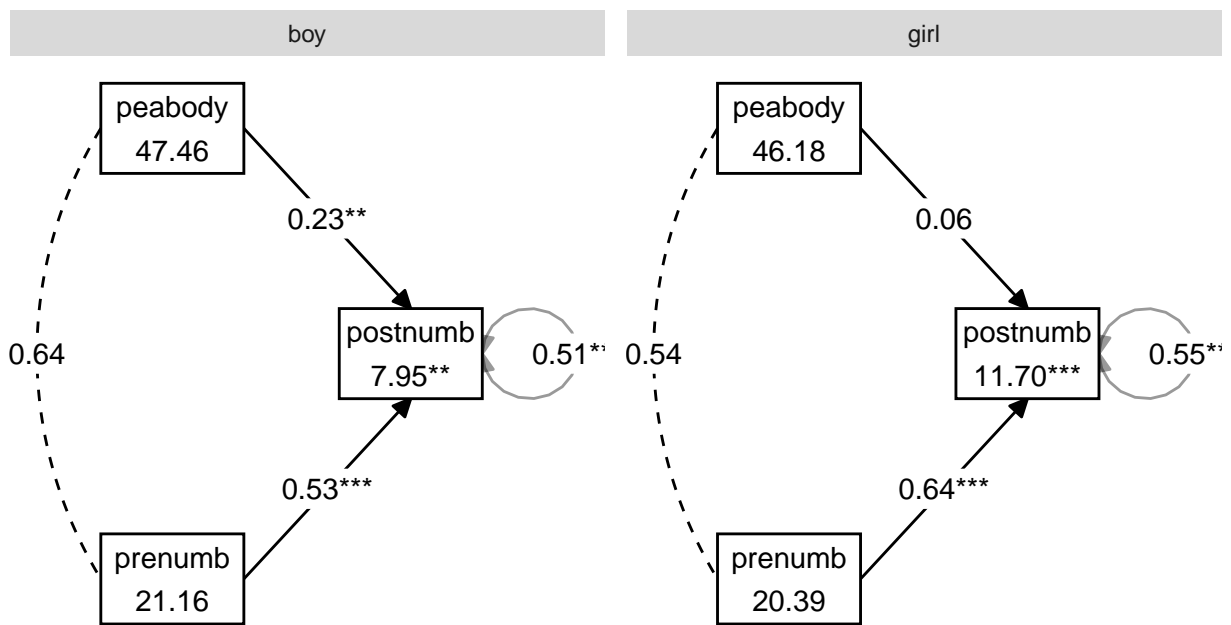


Figure 2. Multiple group analysis