Bayes Factor Model Comparisons Across Parameter Values for Mixed Models

Maximilian Linde*

Don van Ravenzwaaij

Unit of Psychometrics and Statistics, Department of Psychology, Faculty of Behavioural
and Social Sciences, University of Groningen

Special issue in Computational Brain and Behavior

*Correspondence concerning this article should be addressed to Maximilian Linde, University of Groningen, Department of Psychology, Unit of Psychometrics and Statistics, Grote Kruisstraat 2/1, Heymans building, Room 217, 9712 TS Groningen, The Netherlands, Phone: +31 503632702, Email: m.linde@rug.nl.

Abstract

Nested data structures, in which conditions include multiple trials and are fully crossed with participants, are often analyzed using repeated-measures analysis of variance or mixed effects models. Typically, researchers are interested in determining whether there is an effect of the experimental manipulation. These kinds of analyses have different appropriate specifications for the null and alternative models, and a discussion on which is to be preferred and when is sorely lacking. van Doorn et al. (2021) performed three types of Bayes factor model comparisons on a simulated data set in order to examine which model comparison is most suitable for quantifying evidence for or against the presence of an effect of the experimental manipulation. Here we extend their results by simulating multiple data sets for various scenarios and by using different prior specifications. We demonstrate how three different Bayes factor model comparison types behave under changes in different parameters, and we make concrete recommendations on which model comparison is most appropriate for different scenarios.

Keywords: Bayes factor, mixed effects models, model comparison, simulation

1 Introduction

The presence of nested or hierarchical data structures is common in various fields of science. van Doorn et al. (2021, henceforth vDAHSW) describe a typical example of a response time experiment where multiple trials are nested within conditions, and conditions are fully crossed with participants (i.e., condition is a within-subjects variable). In such experiments researchers are typically interested in the effect of the experimental manipulation (i.e., the condition) on the dependent variable (i.e., some measure of the response time). Critically, there are multiple statistical approaches towards that end and, as alluded to by vDAHSW, the choice among them depends on many factors.

Traditionally, the observations in all trials of each participant within each condition are aggregated by calculating the mean, leaving one value of the dependent variable per participant and condition. Subsequently, a repeated-measures analysis of variance (RM-ANOVA) is conducted on this aggregated data set and information on the effect of the experimental manipulation can be extracted. Such an effect is typically labeled a *fixed effect* and it represents a shift in the population mean between conditions (see, e.g., the top-row panels of Figure 1 in vDAHSW).

Alternatively, mixed effects models (also called multilevel, hierarchical, varying effects, or random effects models; cf. McElreath, 2020) have gained popularity and are considered to be one of the most important developments in the statistical sciences in recent years (Gelman & Vehtari, 2021). These models are able to capitalize on the full unaggregated data structure. In addition to fixed effects, these models provide information on potential baseline differences between participants (random intercepts, see e.g., the middle-row panels of Figure 1 in vDAHSW) and the potential differential impact of the experimental manipulation on participants (random slopes, see, e.g., the bottom-row panels of Figure 1 in vDAHSW) in addition to information on the effect of the experimental manipulation. Random intercepts and random slopes vary at the participant level (i.e., one per participant) and can characterize differences in effects on the level of condition, but also on the level of individual trials. For the remainder of this paper, random intercepts and slopes strictly refer to individual effects on the level of condition. We refer the reader to, for instance, Baayen et al. (2008), Gelman and Hill (2006), Kruschke (2015), McElreath (2020), Singmann and Kellen (2019), and Snijders and Bosker (2011) for general introductions to mixed effects models.

Independent of whether RM-ANOVA or mixed effects models are used, testing for the presence or absence of an effect of the experimental manipulation requires the specification of a null model $\mathcal{M}_0$ that *does not* contain the effect of interest and an alternative model $\mathcal{M}_1$ that *does* contain the effect of interest. In the Bayesian inference framework, a Bayes factor can be calculated that quantifies the likelihood of the data y under $\mathcal{M}_1$ relative to the likelihood of y under $\mathcal{M}_0$ (or vice versa):

$$\text{BF}_{10} = \frac{P(y \mid \mathcal{M}_1)}{P(y \mid \mathcal{M}_0)}. \quad (1)$$

Bayes factors can range from 0 to ∞ , where $\text{BF}_{10} > 1$ indicates evidence in favor of $\mathcal{M}_1$ and $\text{BF}_{10} < 1$ evidence in favor of $\mathcal{M}_0$ (see, e.g., Etz & Vandekerckhove, 2018; Heck et al., 2020; Jeffreys, 1961; Kass & Raftery, 1995; van Ravenzwaaij & Etz, 2021).

As emphasized by vDAHSW, there is no consensus yet on which models should be used as $\mathcal{M}_0$ and $\mathcal{M}_1$ for the Bayes factor model comparison regarding the effect of the experimental manipulation. vDAHSW list two possibilities for selecting $\mathcal{M}_0$, two possibilities for selecting $\mathcal{M}_1$, and three comparisons between those: RM-ANOVA, comparing $\mathcal{M}_{0-RM}$ to $\mathcal{M}_{1-RM}$, Balanced null, comparing $\mathcal{M}_{0-BN}$ to $\mathcal{M}_{1-BN}$, and Fixed null, comparing $\mathcal{M}_{0-RM}$ to $\mathcal{M}_{1-BN}$ (see vDAHSW for more details). The two versions of $\mathcal{M}_0$ under consideration are:

- $\mathcal{M}_{0-RM}$, containing random intercepts for participants α_i but no random slopes for participants θ_i and no fixed effect of condition ν :

$$\begin{aligned} Y_{ijm} &\sim \text{N}(\alpha_i, \sigma_\epsilon^2) \\ \alpha_i &\sim \text{N}(\mu, \sigma_\alpha^2) \end{aligned} \quad (2)$$

- $\mathcal{M}_{0-BN}$, containing random intercepts for participants α_i and random slopes for participants θ_i but no fixed effect of condition ν :

$$\begin{aligned} Y_{ijm} &\sim \text{N}(\alpha_i + x_j\theta_i, \sigma_\epsilon^2) \\ \alpha_i &\sim \text{N}(\mu, \sigma_\alpha^2) \\ \theta_i &\sim \text{N}(0, \sigma_\theta^2) \end{aligned} \quad (3)$$

Thus, the difference between $\mathcal{M}_{0-RM}$ and $\mathcal{M}_{0-BN}$ lies in the inclusion of random slopes

for participants. The two versions of $\mathcal{M}_1$ under consideration are:

- $\mathcal{M}_{1-RM}$, containing the fixed effect of condition ν and random intercepts for participants α_i but no random slopes for participants θ_i :

$$\begin{aligned} Y_{ijm} &\sim N(\alpha_i + x_j\nu, \sigma_\epsilon^2) \\ \alpha_i &\sim N(\mu, \sigma_\alpha^2) \end{aligned} \tag{4}$$

- $\mathcal{M}_{1-BN}$, containing the fixed effect of condition ν , random intercepts for participants α_i , and random slopes for participants θ_i :

$$\begin{aligned} Y_{ijm} &\sim N(\alpha_i + x_j\theta_i, \sigma_\epsilon^2) \\ \alpha_i &\sim N(\mu, \sigma_\alpha^2) \\ \theta_i &\sim N(\nu, \sigma_\theta^2) \end{aligned} \tag{5}$$

Again, the difference between $\mathcal{M}_{1-RM}$ and $\mathcal{M}_{1-BN}$ lies in the inclusion of random slopes for participants. The three comparisons between these, as described in vDAHSW, are:

- RM-ANOVA: $\mathcal{M}_{0-RM}$ versus $\mathcal{M}_{1-RM}$ using the aggregated data. Both models allow for different intercepts between participants ($\sigma_\alpha^2 \neq 0$): The only difference between the two models being compared is that in $\mathcal{M}_{0-RM}$ the fixed effect of condition is $\nu = 0$, whereas in $\mathcal{M}_{1-RM}$ the fixed effect of condition is $\nu \neq 0$.
- Balanced null: $\mathcal{M}_{0-BN}$ versus $\mathcal{M}_{1-BN}$ using the full data. Both models allow for different intercepts and slopes between participants ($\sigma_\alpha^2 \neq 0$ and $\sigma_\theta^2 \neq 0$): The only difference between the two models being compared is that in $\mathcal{M}_{0-BN}$ the fixed effect of condition is $\nu = 0$, whereas in $\mathcal{M}_{1-BN}$ the fixed effect of condition is $\nu \neq 0$.
- Strict null: $\mathcal{M}_{0-RM}$ versus $\mathcal{M}_{1-BN}$ using the full data. Both models allow for different intercepts between participants ($\sigma_\alpha^2 \neq 0$): The difference between the two models being compared is that in $\mathcal{M}_{0-RM}$ the fixed effect of condition is $\nu = 0$ and the variance of random slopes for participants is $\sigma_\theta^2 = 0$, whereas in $\mathcal{M}_{1-BN}$ the fixed effect of condition is $\nu \neq 0$ and the variance of random slopes for participants is $\sigma_\theta^2 \neq 0$.

Therefore, these three Bayes factor model comparison methods ask different questions. The essence of the RM-ANOVA and the Balanced null methods is to examine whether

there is an average effect of the experimental manipulation across participants. In contrast, the Strict null method examines whether there is an average effect of the experimental manipulation across participants or whether there is variation of the effect of the experimental manipulation for individual participants.

1.1 The Current Study

The results of Example 1 in vDAHSW offer important insights on the usability of the different Bayes factor model comparison methods for nested data. However, the Bayes factors obtained by vDAHSW only apply to a single simulated data set and a unique prior specification. Since the data simulation entails sampling variability, it is worthwhile to examine the robustness of the Bayes factors applied to multiple simulated data sets with the same parameter values. Moreover, although the choice of parameter values for the data simulation of vDAHSW is plausible, it is unclear whether the behavior of the three Bayes factor model comparison methods is similar for data sets that were generated using different parameter values. Lastly, vDAHSW used one specific setting for the priors. It is possible that different settings for the priors affect the three types of Bayes factor comparisons differently. In the current study, we explore the scenario of Example 1 by vDAHSW more thoroughly in order to assess the robustness of the Bayes factor comparison results for data generated using different parameter values and for different choices of priors.

2 Methods

We calculated Bayes factors corresponding to the RM-ANOVA, the Balanced null, and the Strict null model comparison types for various scenarios of data sets and prior specifications. As a reference point, we print here the parameter values vDAHSW used in their Example 1:

- $M = 15$ trials,
- $I = 20$ participants,
- $J = 2$ conditions,
- $\nu = 0.5$, the fixed effect of condition,

- $\sigma_\alpha^2 = 0.5$, the variance of the random intercepts for participants,
- $\sigma_\theta^2 = 1$, the variance of the random slopes for participants,
- $\rho_{\alpha\theta} = 0$, the correlation between random intercepts and random slopes for participants.

In both the original simulation and our simulations, trials, participants and conditions are fully crossed. The authors used multivariate Cauchy priors with scales of $r_{fixed} = 0.5$ and $r_{random} = 1$ for fixed and random effects, respectively.

We simulated data sets using the parameter values of vDAHSW, while varying one parameter at a time. Doing this allowed us to inspect the effect of variations in each of the parameters individually on the results reported by vDAHSW (see the Appendix for two-way interactions between parameters and the Supplementary Material, available at <https://osf.io/gz57a/>, for higher-way interactions). The parameter values we used are (parameter values shown in bold font correspond to the default values employed by vDAHSW):

- $\nu = \{0, 0.1, \dots, \mathbf{0.5}, \dots, 0.8\}$,
- $\sigma_\alpha^2 = \{0, 0.25, \mathbf{0.5}, \dots, 2\}$,
- $\sigma_\theta^2 = \{0, 0.25, \dots, \mathbf{1}, \dots, 2\}$,
- $\rho_{\alpha\theta} = \{-0.8, -0.6, \dots, \mathbf{0}, \dots, 0.8\}$,
- $I = \{5, 10, \dots, \mathbf{20}, \dots, 50\}$,
- $M = \{5, 10, \mathbf{15}, \dots, 50\}$,
- $r_{fixed} = \{0.2, 0.3, \dots, \mathbf{0.5}, \dots, 0.8\}$,
- $r_{random} = \{0.4, 0.6, \dots, \mathbf{1}, \dots, 1.6\}$,

yielding 70 different combinations of parameter values.

The ranges of parameter values were chosen in a way such that they entail a broad spectrum of values that we consider to be realistic. For instance, the values for ν are oriented at rules of thumb for effect sizes by Cohen (1988), where $\nu = 0.2$, $\nu = 0.5$, and $\nu = 0.8$ are considered small, medium, and large effects, respectively; $\nu = 0$ was

included to be able to compare Bayes factor model comparison methods when no effect of the experimental manipulation is present and the Bayes factors should thus indicate pro $\mathcal{M}_0$ evidence. As another example, the range of mathematically possible correlations is $[-1, 1]$, which gave us a natural restriction for $\rho_{\alpha\theta}$; since we deemed correlations of -1 or 1 highly unlikely, we limited the range for $\rho_{\alpha\theta}$ to be $[-0.8, 0.8]$.

Note that in our simulations, we limited the number of conditions to $J = 2$. Presently, the data generation function provided by vDAHSW (available at <https://osf.io/tjgc8/>) only works for $J = 2$ conditions (personal communication with vDAHSW). Moreover, equivalent to the methods of vDAHSW, we keep the fixed effect of trial, the variance of random intercepts of trial, the variance of random slopes of trial, and the correlation of random intercepts and random slopes of trial constant at 0. Because of sampling variability, we created 50 data sets for each of the 70 possible parameter value combinations, which allowed us to assess the robustness of the resulting Bayes factors to sampling fluctuations.

Of all 3500 full unaggregated data sets, we calculated an additional aggregated data set by taking the mean of the scores on the dependent variable of the trials for each participant and each condition, yielding 3500 full data sets and 3500 corresponding aggregated data sets. Analogous to the procedure followed by vDAHSW, we calculated the RM-ANOVA, the Balanced null, and the Strict null Bayes factors with the BayesFactor package (Morey & Rouder, 2018) written in R (R Core Team, 2021). The relevant files can be accessed at <https://osf.io/gz57a/>, which strongly rely on the files from vDAHSW (accessible at <https://osf.io/p74wk/>).

3 Results

The results of the comparisons between the RM-ANOVA, the Balanced null, and the Strict null Bayes factor model comparisons for (1) the fixed effect of condition ν ; (2) the variance of random intercepts for participants σ_α^2 ; (3) the variance of random slopes for participants σ_θ^2 ; (4) the correlation of random intercepts and random slopes for participants $\rho_{\alpha\theta}$; (5) the number of subjects I ; (6) the number of trials M ; (7) the scale of the multivariate Cauchy prior for fixed effects r_{fixed} ; and (8) the scale of the multivariate Cauchy prior for random effects r_{random} are shown in Figures 1, 2, 3, 4, 5, 6, 7, 8,

respectively. Importantly, in each Figure, only the parameter shown on the x -axis is varied; all other parameters are kept constant at the values used in vDAHSW. Panel A of each Figure shows the obtained Bayes factors; Bayes factors higher than the highest 75% quantile (i.e., the upper edge of the highest boxplot) are not shown. In order to enhance the visibility of the differences in the Bayes factors for the RM-ANOVA and Balanced null comparisons, Panel B of each Figure displays a zoomed-in version of Panel A. Moreover, results for the RM-ANOVA Bayes factors were obtained on the aggregated data and results for the Balanced null and Strict null Bayes factors on the full data. In the Appendix and Supplementary Material, we present simulations that vary multiple parameters at once. The Appendix shows potential two-way interactions between parameters. Higher-way interactions can be inspected in the Supplementary Material, available at <https://osf.io/gz57a/>.

Figure 1 shows that for the full data sets the Strict null method always provides overwhelming evidence in favor of $\mathcal{M}_1$, even when the fixed effect of condition is $\nu = 0$. This strongly suggests that the Strict null Bayes factor is largely driven by the presence of random slopes for participants, which are part of $\mathcal{M}_1$ but not $\mathcal{M}_0$ (see also vDAHSW). This behavior is likely to be undesirable in most practical applications: When the fixed effect of condition is of interest, pro $\mathcal{M}_0$ evidence should be obtained when the true fixed effect of condition is (close to) $\nu = 0$. In contrast, the Balanced null method produces Bayes factors that vary meaningfully with the fixed effect of condition; that is, the higher the fixed effect of condition, the higher the Bayes factors. Nevertheless, the Balanced null Bayes factors are conservative, only providing weak evidence towards $\mathcal{M}_1$ even for large fixed effects of condition (i.e., $\nu > 0.5$). For the aggregated data, the RM-ANOVA method varies meaningfully with the fixed effect of condition. When $\nu = 0$ evidence in favor of $\mathcal{M}_0$ is usually obtained. For moderate and large fixed effects of condition (i.e., $\nu > 0.5$), the evidence is relatively weak but still quite clearly points towards $\mathcal{M}_1$.

Figure 2 indicates that for the full data sets the Bayes factors for the Balanced null and Strict null methods are unaffected by the variance of random intercepts for participants. For all variances of random intercepts for participants, the Strict null method provides overwhelming evidence for $\mathcal{M}_1$, whereas the Balanced null method displays ambiguous evidence for $\mathcal{M}_0$ versus $\mathcal{M}_1$. Similarly, for the aggregated data sets, the RM-ANOVA method also produces Bayes factors that do not depend on the

Figure 1: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the fixed effect of the independent variable ‘condition’. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

variance of random intercepts for participants (or at least, not strongly). Generally, the RM-ANOVA method provides relatively weak evidence in favor of $\mathcal{M}_1$. The invariance of all three types of Bayes factor model comparison methods to the variance of random intercepts for participants is as it should be because this term is included in all null models and alternative models used in the corresponding comparisons.

As shown in Figure 3, for the full data sets the Strict null method still provides extreme evidence in favor of $\mathcal{M}_1$ for all variances of random slopes for participants. It is further visible that the Bayes factors obtained with the Strict null method sharply increase as the variance of random slopes for participants increases. In contrast, the

Figure 2: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the variance of the random intercepts for participants. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

Balanced null method yields Bayes factors providing ambiguous evidence for relatively high variances of random slopes for participants (i.e., $\sigma_\theta^2 > 0.5$). As the variance of random slopes for participants decreases, the Balanced null Bayes factors indicate clear evidence towards $\mathcal{M}_1$. When the aggregated data are used, the RM-ANOVA method also yields Bayes factors providing ambiguous evidence for high variances of random slopes for participants (i.e., $\sigma_\theta^2 > 0.5$) and clear pro $\mathcal{M}_1$ evidence for lower variances of random slopes for participants. We believe that the Balanced null and RM-ANOVA Bayes factors are higher for low variances of the random slopes for participants due to the posterior being more narrow. The reason for this is that the fixed effect of condition can be estimated

Figure 3: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the variance of the random slopes for participants. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

more precisely in the absence of noise (i.e., variance of random slopes for participants). This effect should be even more pronounced for RM-ANOVA, as the models no longer have to compensate for the existence of variance of random slopes for participants which is not present in the model specification.

The Balanced null and Strict null Bayes factors on the full data sets as well as the RM-ANOVA Bayes factors on the aggregated data sets are largely unaffected by the correlation of the random intercepts and random slopes for participants (see Figure 4). While the Strict null Bayes factors provide overwhelming evidence in favor of $\mathcal{M}_1$, the Balanced null and RM-ANOVA Bayes factors indicate ambiguous evidence.

Figure 4: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the correlation between the random intercepts and random slopes for participants. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

The results displayed in Figure 5 suggest that the evidence towards $\mathcal{M}_1$ for the Balanced null, Strict null, and RM-ANOVA model comparisons is bolstered by an increasing number of participants. This is as it should be, because having more participants reduces the uncertainty around the estimate of the fixed effect of condition ν . In contrast, Figure 6 shows that an increasing number of trials leads to stronger evidence for $\mathcal{M}_1$ in case of the Strict null Bayes factors on the full data sets but not in case of the Balanced null Bayes factors on the full data sets and the RM-ANOVA Bayes factors on the aggregated data sets. Related to the conclusions of vDAHSW, we suspect that the Strict null Bayes factors increase in magnitude with an increasing number of trials

Figure 5: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the number of subjects. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

because more trials ensure a reduced uncertainty around the estimate of the variance of the random slopes for participants, demonstrating more clearly that $\sigma_\theta^2 \neq 0$.

Figures 7 and 8 demonstrate that the Bayes factors resulting from the RM-ANOVA, Balanced null, and Strict null model comparisons are quite robust against variations of the multivariate Cauchy priors for fixed and random effects, respectively. That is, none of the three Bayes factor model comparison types change substantially as a function of the multivariate Cauchy prior scales.

In terms of two-way interactions (see Figure 9), we found effects for some of the dyads of fixed effect, variance of slopes for participants, variance of intercepts for participants, and number of participants, depending on the specific Bayes factor model comparison.

Figure 6: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the number of trials. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

More details can be found in the Appendix.

4 Discussion

In Example 1 of vDAHSW, the RM-ANOVA, Balanced null, and Strict null Bayes factor model comparison types were compared on a nested data set, where multiple trials were nested within conditions, and conditions and participants were fully crossed. The authors concluded that the RM-ANOVA and Strict null Bayes factors were unreasonably large for a modest effect of condition of $\nu = 0.5$ and only $I = 20$ participants due to the presence of variance of random slopes for participants. Only the Balanced null Bayes

Figure 7: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the scale of the multivariate Cauchy prior for fixed effects. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

factor was appropriately conservative, due to the fact that the variance of random slopes for participants is present in both the null and alternative model being compared.

We extended their findings by comparing the three Bayes factor types for various scenarios. We found that the RM-ANOVA Bayes factor on aggregated data sets and the Balanced null Bayes factor on full data sets varied meaningfully with changes in the fixed effect of condition ν , showing pro $\mathcal{M}_0$ evidence when appropriate. The Strict null Bayes factor on the full data sets is largely driven by the presence of variance of random slopes for participants σ_θ^2 . A larger sample size benefits all three Bayes factors, whereas a larger number of trials increases the Strict null Bayes factor, but not the Balanced null

Figure 8: Bayes factors for RM-ANOVA (purple) on the aggregated data, Balanced null (green) on the full data, and Strict null (yellow) on the full data as a function of the scale of the multivariate Cauchy prior for random effects. Panel B displays a zoomed-in version of Panel A. The empty dots show the Bayes factors for 50 repetitions with the same parameter values.

and RM-ANOVA Bayes factors. The variance of random intercepts for participants, the correlation between random intercepts and slopes for participants, and the settings of the priors had no qualitative effect on the results of vDAHSW. We also observed considerable variation of the Bayes factors for all three Bayes factor model comparison methods for multiple data sets that were sampled with identical population parameters, which in some cases led to qualitatively different conclusions. This variability mainly stems from sampling fluctuations. Clearly then, any single Bayes factor resulting from a single data set should be judged with caution (as is the case for all measures that are data-dependent and thus subject to sampling variability; e.g., p -values, confidence intervals).

In general, we agree with the conclusions of vDAHSW. More specifically, with the settings we used in our simulations, the Strict null Bayes factor always provides overwhelming evidence in favor of $\mathcal{M}_1$, even when the fixed effect of condition is around $\nu = 0$. In contrast, the Balanced null Bayes factor varies meaningfully with the fixed effect of condition, providing evidence towards $\mathcal{M}_0$ when the fixed effect of condition is around $\nu = 0$ and at least weak evidence towards $\mathcal{M}_1$ when the fixed effect of condition is approximately $\nu > 0.5$. The flipside is that the Balanced null Bayes factor appears to be quite conservative, habitually producing Bayes factors of a magnitude of about 10 even when the fixed effect of condition is $\nu = 0.5$. We should keep in mind though, that these results are for the scenario when the sample size is $I = 20$. Arguably it should be difficult to obtain decisive evidence for such a small sample size in case of all but the most extravagant effect sizes.

The Strict null Bayes factor clearly asks a different question from the Balanced null Bayes factor. Where the latter asks ‘Is there a mean effect of condition?’ the former asks ‘Is there any effect of condition?’. Traditionally, the focus of testing and quantifying effects has been on the mean, so in that sense the Balanced null comparison is more in line with tradition, but it is certainly conceivable that the variance of random slopes for participants is the focal interest of the researcher. In that case, we recommend a more gradual comparison of models, comparing $\mathcal{M}_{0-RM}$ to $\mathcal{M}_{0-BN}$ (i.e., the models that differ only in their inclusion of variance of random slopes for participants), so that there is no conflation of effects in the Bayes factor.

Another approach is to calculate an inclusion Bayes factor. Inclusion Bayes factors allow one to compare all models that include a predictor (e.g., fixed effect of condition or variance of random slopes for participants) with all those that do not, to determine the relative strength of evidence for each factor on the dependent variable based on the data (van den Bergh et al., 2020). For instance, the inclusion Bayes factor for the fixed effect of condition for the four previously discussed models would be:

$$\text{BF}_{\nu \neq 0, \nu = 0} = \frac{\left(\frac{P(\mathcal{M}_{1-RM} | y) + P(\mathcal{M}_{1-BN} | y)}{P(\mathcal{M}_{0-RM} | y) + P(\mathcal{M}_{0-BN} | y)} \right)}{\left(\frac{P(\mathcal{M}_{1-RM}) + P(\mathcal{M}_{1-BN})}{P(\mathcal{M}_{0-RM}) + P(\mathcal{M}_{0-BN})} \right)}; \quad (6)$$

and the inclusion Bayes factor for the variance of the random slopes for participants for

the four previously discussed models would be:

$$\text{BF}_{\sigma_{\theta}^2 \neq 0, \sigma_{\theta}^2 = 0} = \frac{\left(\frac{P(\mathcal{M}_{0-BN} | y) + P(\mathcal{M}_{1-BN} | y)}{P(\mathcal{M}_{0-RM} | y) + P(\mathcal{M}_{1-RM} | y)} \right)}{\left(\frac{P(\mathcal{M}_{0-BN}) + P(\mathcal{M}_{1-BN})}{P(\mathcal{M}_{0-RM}) + P(\mathcal{M}_{1-RM})} \right)}. \quad (7)$$

For a more extensive discussion on inclusion Bayes factors, see Heck and Bockting (2021).

4.1 Conclusion

Bayes factors are not magic. They can answer the question researchers ask. In that sense none of the three comparisons is strictly superior to any other. In practice, our intuition is that researchers are most often interested in testing for the presence of a fixed effect of condition. In such a case, the Balanced null comparison using full unaggregated data or possibly the RM-ANOVA comparison using aggregated data are to be preferred. If the focus is on the variance of random slopes for participants alone, we believe an entirely different model comparison is in order. Regardless, we echo the sentiment by vDAHSW that ‘a discussion on best practices in Bayes factor model comparison in mixed models’ is in order and we hope that our simulation results contribute to this discussion.

5 Funding

This research was supported by a Dutch scientific organization VIDI fellowship grant awarded to Don van Ravenzwaaij (016.Vidi.188.001).

6 Conflict of interest/Competing interests

The authors have no conflicts of interest and no competing interests to declare.

7 Availability of data and material

The simulated data is publicly available at <https://osf.io/gz57a/>.

8 Code availability

The simulation code is publicly available at <https://osf.io/gz57a/>.

9 Author's contributions

ML: methodology, formal analysis, writing - original draft. DvR: conceptualization, writing - review and editing, supervision, funding acquisition.

10 Ethics approval

Not applicable.

11 Consent to participate

Not applicable.

12 Consent for publication

Not applicable.

References

- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59, 390–412. <https://doi.org/10.1016/j.jml.2007.12.005>
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd). Lawrence Erlbaum Associates.
- Etz, A., & Vandekerckhove, J. (2018). Introduction to Bayesian inference for psychology. *Psychonomic Bulletin and Review*, 25(1), 5–34. <https://doi.org/10.3758/s13423-017-1262-3>
- Gelman, A., & Hill, J. (2006). *Data analysis using regression and multilevel/hierarchical models*. Cambridge University Press.
- Gelman, A., & Vehtari, A. (2021). What are the most important statistical ideas of the past 50 years?
- Heck, D. W., & Bockting, F. (2021). Benefits of Bayesian model selection and averaging for mixed-effects modeling. <https://doi.org/10.31234/osf.io/zusd2>
- Heck, D. W., Boehm, U., Böing-Messing, F., Bürkner, P.-C., Derks, K., Dienes, Z., Fu, Q., Gu, X., Karimova, D., Kiers, H. A. L., Klugkist, I., Kuiper, R. M., Lee, M. D., Leenders, R., Leplaa, H. J., Linde, M., Ly, A., Meijerink-Bosman, M., Moerbeek, M., ... Hoijsink, H. (2020). A review of applications of the Bayes factor in psychological research. <https://doi.org/10.31234/osf.io/cu43g>
- Jeffreys, H. (1961). *Theory of probability* (3rd). Oxford University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 773–795. <https://doi.org/10.2307/2291091>
- Kruschke, J. K. (2015). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (2nd). Academic Press.
- McElreath, R. (2020). *Statistical rethinking: A Bayesian course with examples in R and Stan* (2nd). Chapman & Hall/CRC.
- Morey, R. D., & Rouder, J. N. (2018). *BayesFactor: Computation of Bayes factors for common designs* [R package version 0.9.12-4.2]. <https://CRAN.R-project.org/package=BayesFactor>

- R Core Team. (2021). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria. <https://www.R-project.org/>
- Singmann, H., & Kellen, D. (2019). An introduction to mixed models for experimental psychology. *New Methods in Cognitive Psychology*, 28, 4–31.
- Snijders, T. A. B., & Bosker, R. J. (2011). *An introduction to basic and advanced multilevel modeling* (2nd). Sage.
- van den Bergh, D., van Doorn, J., Marsman, M., Draws, T., van Kesteren, J., Derks, K., Dablander, F., Gronau, Q. F., Kucharský, Š., Komarlu Narendra Gupta, A. R., Sarafoglou, A., Voelkel, J. G., Stefan, A., Ly, A., Hinne, M., Matzke, D., & Wagenmakers, E.-J. (2020). A tutorial on conducting and interpreting a Bayesian ANOVA in JASP. *L'Année Psychologique*, 120(1), 73–96. <https://doi.org/doi.org/10.3917/anpsy1.201.0073>
- van Doorn, J., Aust, F., Haaf, J. M., Stefan, A., & Wagenmakers, E.-J. (2021). Bayes factors for mixed models. <https://doi.org/10.31234/osf.io/y65h8>
- van Ravenzwaaij, D., & Etz, A. (2021). Simulation studies as a tool to understand Bayes factors. *Advances in Methods and Practices in Psychological Science*. <https://doi.org/10.1177/2515245920972624>

A Appendix: Two-way Interactions Between Parameters

Here, we present results from our extended parameter simulation in terms of pairs of parameters. This should allow for a qualitative inspection of the presence of two-way interactions for each of the three Bayes factor model comparisons. Figure 9 presents the results for RM-ANOVA Bayes factors, Figure 10 presents the results for Balanced null Bayes factors, and Figure 11 presents the results for Strict null Bayes factors.

Turning first to the results of the RM-ANOVA Bayes factors (Figure 9), we see interactions for (1) fixed effect of condition versus variance of random intercepts for participants (highest for high fixed effect of condition and zero variance of random intercepts for participants); (2) fixed effect of condition versus variance of random slopes for participants (highest for high fixed effect of condition and zero variance of random slopes for participants); (3) fixed effect of condition versus number of participants (highest for high fixed effect of condition and high number of participants); and (4) variance of random slopes for participants and number of participants (highest for high number of participants and zero variance of random slopes for participants). Generally, these results make sense: Since the relevant comparison of RM-ANOVA focuses on the fixed effect of condition, any reduction in variance of nuisance parameters will bolster the Bayes factor. Its susceptibility to zero variance of random intercepts for participants mirrors the change that can be observed for RM-ANOVA specifically for the main effects (see leftmost result in Figure 2).

For the Balanced null Bayes factors (Figure 10), we see interactions for (1) fixed effect of condition versus variance of random slopes for participants (highest for high fixed effect of condition and zero variance of random slopes for participants); (2) fixed effect of condition versus number of participants (highest for high fixed effect of condition and high number of participants); and (3) variance of random slopes for participants and number of participants (highest for high number of participants and zero variance of random slopes for participants). These results mirror the results for RM-ANOVA, except that there are no interactions involving the variance of random intercepts for participants.

For the Strict null Bayes factors (Figure 11), we see interactions for (1) fixed effect of condition versus number of participants (highest for high fixed effect of condition and high

number of participants); (2) number of participants versus variance of random slopes for participants (highest for high number of participants and high variance of random slopes for participants); (3) fixed effect of condition versus number of trials (highest for high fixed effect of condition and high number of trials); (4) number of trials versus variance of random slopes for participants (highest for high number of participants and high variance of random slopes for participants); and (5) number of trials versus number of participants (highest for high number of participants and high number of trials). Interestingly, two-way interactions for the Strict null Bayes factor always feature the number of participants or trials. The pattern is always the same: High values for both parameters lead to the largest Bayes factor. As for the previous two model comparisons, generally two-way interactions were observed for parameters that also showed main effects (see Figures 5 and 6). As a result, we caution against over-interpreting these two-way interactions.

Figure 9: Median RM-ANOVA Bayes factors on aggregated data for all combinations of values of two respective parameters. The median is calculated for the Bayes factors over all values of the six remaining parameters. Note that the y -axes in the panels display different ranges of Bayes factors.

Figure 10: Median Balanced null Bayes factors on aggregated data for all combinations of values of two respective parameters. The median is calculated for the Bayes factors over all values of the six remaining parameters. Note that the y -axes in the panels display different ranges of Bayes factors.

Figure 11: Median Strict null Bayes factors on aggregated data for all combinations of values of two respective parameters. The median is calculated for the Bayes factors over all values of the six remaining parameters. Note that the y -axes in the panels display different ranges of Bayes factors.