Draft version 1.2, 07/07/20. This manuscript has not been peer reviewed. Please do not copy or cite without author's permission. A version of this manuscript will be part of a paper “Applications of Bayes factor in Psychological Research”.

Interval Hypotheses: When and how to calculate the Bayes factor with an interval null hypothesis

Palfi, B.¹, & Dienes, Z.¹

¹School of Psychology, University of Sussex, Brighton, UK

Correspondence regarding this vignette can be sent to Bence Palfi (palfibence@gmail.com).

Existing calculators of the Bayes factor typically represent the prediction of the null-hypothesis with a point null model (e.g., Dienes & McIlatchie, 2018; Morey & Rouder, 2015; Wagenmakers et al., 2018). The point null model remains a good approximation of the null hypothesis as long as the standard error of the estimate of interest is large relative to the range of theoretically negligible effect sizes. However, this assumption may be violated, for instance, when one has a large dataset and the standard error becomes comparable to or smaller than the range of effect sizes that is meaningless according to the theory. In this case, the conclusion that was based on the point null model may not be scientifically valid. Here, we introduce a case study to demonstrate why it is critical to calculate the Bayes factor with an interval null hypothesis rather than a point null hypothesis when the smallest meaningful effect size is not approximately zero relative to the standard error. The example illustrates how a Bayes factor can be used to test whether one intervention is superior to another (e.g. placebo) when a smallest meaningful effect size is a relevant consideration.

Data. Utilising data submitted to the U.S. Food and Drug Administration, Kirsch et al. (2002) evaluated in a meta-analysis whether or not antidepressant drugs are more efficient than placebo. In their crucial test, they incorporated the results of 19 randomised placebo controlled trials covering 3 of the most widely used antidepressants: fluoxetine, venlafaxine and nefazodone. Their analysis included the data of 3708 patients with moderate to severe depression, from which 2401 were assigned to the drug and 1307 to the placebo group. The clinical trials typically lasted 6 weeks and the crucial test was conducted on the data registered on the last visit of the patients. The severity of depression was measured by the Hamilton Depression Scale (HAM-D, Hamilton, 1960) in each trial.

Model. The two entertained hypotheses are H_0 , which asserts that antidepressants are equally effective to placebo, and H_1 , which postulates that antidepressants are more efficient than placebo. We modelled the predictions of H_0 on the raw HAM-D scores with a uniform distribution symmetric around zero and with an upper bound of 3 units, that is $M_0 : U[-3, 3]$. The upper bound of M_0 was derived from the criterion of clinical meaningfulness defined by the National Institute for Health and Care Excellence (NICE) that sets the threshold of a meaningful intervention at a change of 3 HAM-D units. Moreover, an evaluation of the correspondence between HAM-D scores and clinicians' ratings of the patients' improvement revealed that a 3 units change on HAM-D corresponds to "no change" on the scale assessing the clinicians' ratings (Moncrieff & Kirsch, 2015). Hence, any change smaller than 3 HAM-D units can be regarded as clinically irrelevant and thus defining an *interval* rather than point H_0 is necessary. To represent the scientific situation, M_0 should represent this interval H_0 . To model the predictions of H_1 , we applied a half-normal distribution with a mode of 3, which is the upper bound of the null interval, and an SD of 4, that is $M_1 : H(3, 4)$. The SD of the distribution was based on the finding that the change of 7 HAM-D units is equivalent to the clinicians' rating of "minimally improved" (Moncrieff & Kirsch, 2015). H_1 represents a plausible range of effects from 3 to about 11 HAM-D units, i.e. 7 is in the middle. We report the result of the corresponding Bayes factor analysis as $BF_{10} M_0:U[-3, 3], M_1:H(3, 4)$. We consider an alternative scenario in which we neglect the NICE recommendation regarding the smallest meaningful effect size on the HAM-D. In this scenario, we use a point-null model to represent

the predictions of H0 and a half-normal distribution with a mode of zero and an SD of 7 to model the predictions of H1. Again, the SD of the distribution is derived from the finding that an increment of 7 HAM-D units is deemed “minimally improved” by clinicians (Moncrieff & Kirsch, 2015). The corresponding Bayes factor is reported as $BF_{10 \text{ M1:H}(0,7)}$. Note that clinical meaningfulness is best represented in raw not standardized units.

Bayes factor, Prior and Posterior Model Probabilities. The meta-analysis comparing the efficiency of antidepressants to placebo found a mean difference of 1.8 HAM-D units ($SE = 0.36$) between the groups. These results translate into strong evidence for H0 over H1 as these data are about 480 times more likely under M0 than under M1 ($BF_{10 \text{ M0:U}[-3, 3], \text{M1:H}(3, 4)} = 2.08 \times 10^{-3}$).¹ Interestingly, if we ignore the NICE criterion about the smallest meaningful effect size, and use a point null model for H0 then our results would favour H1 over the point H0 ($BF_{10 \text{ M1:H}(0,7)} = 395.37$; cf. Monden et al., 2018). This apparent (but not actual) contradiction between the conclusion of the two analyses highlights why one needs to carefully consider the model of H0 (as well as the model of H1). Misrepresenting H0 by not defining in the model that a range of effect sizes is predicted by H0 (and hence representing any positive effect size however small as consistent with H1) provides unreasonable support for H1, i.e. support based on incorrect modelling of the scientific context. In the presented case study, the most reasonable model of the predictions of H0 is a symmetric interval around zero with an upper bound of 3 HAM-D units. The corresponding analysis reveals strong evidence that three of the most widely used antidepressant drugs are clinically equivalent to placebo. That is, three common antidepressants deliver no meaningful advantage over placebo in terms of alleviating depression.

Key References and Software. The R script of the Bayes factor calculator is based on the calculator of Dienes and Mclatchie (2018). For an introduction on how to specify the parameters of the model representing the predictions of H1 see Dienes (2019). The R script of the Bayes factor calculator and of the analysis can be found on the OSF website².

Dienes, Z. (2019). How Do I Know What My Theory Predicts? *Advances in Methods and Practices in Psychological Science*. 2(4), 364-377.

Dienes, Z., & Mclatchie, N. (2018). Four reasons to prefer Bayesian analyses over significance testing. *Psychonomic bulletin & review*, 25(1), 207-218.

Hamilton, M. (1960). A rating scale for depression. *Journal of neurology, neurosurgery, and psychiatry*, 23(1), 56-61.

¹ To ascertain the robustness of the conclusion to the parameters of M0 and M1, we identified the smallest and largest values of the boundaries of M0 and of the SD of M1 with which we gain good enough evidence for H0 (for more information on Robustness Regions [RR] see Dienes, 2019). Keeping the parameters of M1 constant while modifying the boundaries of M0 yielded $RR_{\text{M0:U}[1.99, \text{Inf}]}$; the evidence for lack of effectiveness holds for any plausible model of H0. Note that the Bayes factor provides good enough evidence for H0 even when the 95% CI of the effect [1.09, 2.51] extends far beyond the upper boundary of the interval null model (contrast with equivalence testing; Schuirmann, 1987). Fixing the parameters of M0 and changing the SD of M1 resulted in $RR_{\text{M1:H}[0, \text{Inf}]}$; the evidence for H0 holds for any plausible model of H1.

² <https://osf.io/k9c5q/>

Kirsch, I., Moore, T. J., Scoboria, A., & Nicholls, S. S. (2002). The emperor's new drugs: An analysis of antidepressant medication data submitted to the US Food and Drug Administration. *Prevention & Treatment*, 5(1), 23a.

Moncrieff, J., & Kirsch, I. (2015). Empirically derived criteria cast doubt on the clinical significance of antidepressant-placebo differences. *Contemporary Clinical Trials*, 43, 60-62.

Monden, R., Roest, A. M., van Ravenzwaaij, D., Wagenmakers, E. J., Morey, R., Wardenaar, K. J., & de Jonge, P. (2018). The comparative evidence basis for the efficacy of second-generation antidepressants in the treatment of depression in the US: A Bayesian meta-analysis of Food and Drug Administration reviews. *Journal of affective disorders*, 235, 393-398.

Morey, R. D., & Rouder, J. N. (2015). *BayesFactor 0.9.11-1*. Comprehensive R Archive Network.

Schirmann, D. J. (1987). A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. *Journal of pharmacokinetics and biopharmaceutics*, 15(6), 657-680.

Wagenmakers, E. J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., ... & Meerhoff, F. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic bulletin & review*, 25(1), 58-76.