Testing theories with Bayes factors

Zoltan Dienes

University of Sussex

E-mail: dienes@sussex.ac.uk

To appear in the *Cambridge Handbook of Research Methods and Statistics for the Social and Behavioral Sciences*, edited by Austin Lee Nichols & John E. Edlund.

Abstract

Bayes factors, a measure of evidence for one model versus another, are a useful tool in the behavioral and social sciences, partly because they can provide evidence for no effect relative to the sort of effect expected. By contrast, a non-significant result does not provide evidence for the H_0 tested. If non-significance does not in itself count against any theory predicting an effect, how could a theory fail a test? Bayes factors provide a measure of evidence from first principles. A severe test is one that is likely to obtain evidence against a theory if it were false—to obtain an extreme Bayes factor against the theory. Bayes factors show why cherry picking degrades evidence, how to deal with multiple testing, and how optional stopping is consistent with severe testing. Further, informed Bayes factors can be used to link theory tightly to how that theory is tested, so that the measured evidence does relate to the theory.

Key words:

Bayes factor, severe test, evidence, multiple testing, optional stopping, cherry picking, priors

An integral part of science is testing theories, and many journal articles are phrased as attempts to test theories (though see McPhetres et al., 2020). A key inferential tool used for this purpose is statistics—often null hypothesis significance testing (NHST). However, a non-significant result using the H_0 of no effect does not in itself provide evidence against a theory that predicted an effect. That is, a theory that predicts an effect is not necessarily disconfirmed just because the outcome was non-significant (cf. Edlund et al., 2021); paradoxically, it may even receive confirmation (Dienes & McLatchie, 2018). So, how can we put our theories to the test—ideally, to a “severe” test—so there is a chance they may fail if they are false? After all, passing a severe test would intuitively seem to provide evidence for a claim. Consistent with this, a Bayesian approach is to first define evidence; the strength of evidence of data for one hypothesis versus another is the amount by which the strength of belief in the one hypothesis versus the other should change in the light of the data. This definition, combined with the claim the strength of belief should ideally be consistent with the axioms of probability, leads to the Bayes factor. This chapter will go through how the Bayes factor as a measure of strength of evidence can be used to severely test theories.

Bayes factor as evidence

How much should strength of beliefs change in the light of data? Let $P(H_1)$ be the probability of H_1 —the strength of belief in H_1 —where H_1 is (for example) a hypothesis that something exists (e.g., a difference between conditions, a relationship between two variables, etc.). Let $P(H_0)$ be the probability of H_0 —the strength of belief in H_0 —where H_0 is the hypothesis that (for example) something does not exist (there is not a difference between conditions; there is no relationship, etc.). For example, H_1 could be the hypothesis that there is subliminal perception in a certain paradigm; H_0 that it does not exist. $P(H_1)/P(H_0)$ is called

the prior odds in favor of H_1 rather than H_0 —the relative strength of belief in H_1 rather than H_0 before data are collected. Data D are collected. Then $P(H_1|D)/P(H_0|D)$ is the posterior odds in favor of H_1 rather than H_0 —the relative strength in belief in H_1 rather than H_0 in the light of data D . With some simple rearranging of the axioms of probability, one obtains

$$P(H_1|D) / P(H_0|D) = P(D|H_1) / P(D|H_0) \times P(H_1)/P(H_0)$$

$$\text{Posterior odds} = \text{Bayes factor} \times \text{prior odds}$$

Thus, the Bayes factor is the amount by which one should normatively change confidence in one hypothesis versus another. That is, the Bayes factor is a measure of strength of evidence from first principles, namely the axioms of probability and a definition of evidence (Jeffreys, 1939; Morey et al., 2016; Rouder et al., 2009).

Prior odds may be purely personal. I may believe a priori that subliminal perception is quite likely. Someone else may believe subliminal perception is scarcely credible. Yet, if we both agree on the predictions made by the theory that perception is subliminal under stated conditions, and if these predictions can be made for objective reasons (as we discuss later), then the $P(D|H_1)$ and $P(D|H_0)$ are relatively objective. The Bayes factor would tell each of us how much to change our beliefs in the light of data, in the same direction for both of us, even if we began at different starting points. We started this section by interpreting probabilities as subjective (i.e., as strengths of belief). By doing so, we can separate out what may be purely personal (i.e., prior odds) from the relatively objective message of the data—the Bayes factor.

If the data were impossible on H_1 and possible on H_0 , the Bayes factor would be zero—the data would provide zero evidence for H_1 rather than H_0 (or conversely, infinite

evidence for H_0 rather than H_1). H_1 would be falsified. Conversely, if the data were impossible on H_0 and possible on H_1 , the Bayes factor would be infinite and the evidence would be infinite for H_1 relative to H_0 . H_0 would be falsified. If the data were equally probable on H_1 versus H_0 , the Bayes factor would be 1, and the data would not be evidence for either hypothesis relative to the other. In sum, the Bayes factor varies between 0 and infinity, with 1 as the neutral point of no evidence.

How far from 1 would a Bayes factor have to go before the evidence is good enough for a decision? If there are clear costs and benefits, the answer can be rationally determined (Lindley, 2014). However, in many cases in the social and behavioral sciences, where a decision needs to be made, determining the costs and benefits would be a difficult and unilluminating extra step. Has a possible confound been ruled out clearly enough that we can move on and apply the main theory to another problem? Has a version of the main theory been corroborated well enough that for now we continue to work on it rather than postulate another variant? Should a variable be dropped from the model? In most cases, and in most areas of the social and behavioral sciences, the prior probabilities and costs and benefits are sufficiently similar for the options considered that a convention for good enough evidence would be useful. A conventional degree of good enough evidence would also help stop special pleading for different strengths of evidence for different decisions based on how well it suited a researcher. However, a Bayes factor is a continuous measure of strength of evidence without special bumps at particular values, and it is always open to argue against a convention in particular cases.

In terms of the history of statistical inference, Edgeworth in 1885 suggested two standard errors difference was evidence for a difference between means “just worth taking note of” (Stigler, 1999, p 103). According to Cowles and Davis (1982), in the early years of the twentieth century, two standard errors difference had become a frequent convention for

significance in a variety of sciences. Thus, Fisher was following an already existent tradition in explicitly proposing $p = .05$ as a suitable significance level—the p value that roughly corresponds to two standard errors difference. Jeffreys (1939) developed Bayes factors as a system of inference. He found that, when he addressed the same problems as Fisher, a Bayes factor of three roughly corresponded to 5% significance. This is true, provided the obtained difference (or sample parameter value more generally) is roughly that expected based on H_1 . Thus, using a Bayes factor of roughly 3 as a convention for evidence just worth taking note of, is not arbitrary; it is roughly the amount of evidence scientists have been using for 100 years. That does not mean it is useful to take the convention as a black and white boundary. Stigler’s phrase “just worth taking note of” is often the right attitude. And one may wish to revisit conventions no matter how ancient they are. *Cortex* (2021), for example, now uses a convention of 6 for a Bayes factor indicating good enough evidence (and correspondingly a significance level of 2%).

If a Bayes factor simply re-expressed the evidence already represented by a given significance level, not much would be gained by using Bayes factors. However, there is no monotonic relationship between a Bayes factor and a p value (Jeffreys, 1939; Lindley, 1957, Morey, 2018). The same p value can correspond to a wide range of different Bayes factors, depending on how one models H_1 (i.e., depending on the predictions of the theory being tested). Further, a Bayes factor can distinguish evidence for H_0 (where H_0 is the hypothesis of no effect) from not enough evidence to distinguish H_0 from H_1 . A Bayes factor of about 1/3 is the same amount of evidence for H_0 relative to H_1 as is a Bayes factor of 3 for H_1 relative to H_0 . A Bayes factor closer to 1 is only evidence “not worth more than bare mention” in Jeffreys’ (1939) terms (i.e., not enough evidence to distinguish H_0 from H_1). Thus, a Bayes factor makes a vital distinction not available to those who look only at a p value (using the H_0 of no effect)—it can distinguish evidence for no effect from not enough evidence to say

whether there is an effect. For example, if one obtained a Bayes factor of 0.6, one would not conclude anything about H_1 vs H_0 ; if such a result corresponded to a non-significant result, that non-significant result does not count against a theory that predicted a difference. This is a fact that should be reflected in the discussion section of a paper—the result should be treated as non-evidential.

The Bayes factor depends on how probable the data are given H_1 . Thus, a model is needed of the plausibility of different population values given H_1 (e.g., population mean differences). Call this the model of H_1 . The model of H_1 is often called the *prior* of the Bayes factor, but the term “prior” also refers to the prior odds we just referred to. Thus, the word “prior” has two referents in the context of Bayes factors. For example, one might say that “prior odds are purely personal and subjective, and as Bayes factors depend on priors, that makes Bayes factors arbitrary”. But this argument conflates the two referents of “prior”. Thus, I use two different terms; the term “model of H_1 ” refers to the mathematical way the predictions of a theory are represented; the “prior odds” refers to the relative probability of the two hypotheses being contrasted. Typically, in results sections of papers (see Dienes, 2021a, for examples), I give only the Bayes factor and not the prior odds (though the full Bayesian machinery, including prior odds, can be useful).

Theories, hypotheses and models

Let a substantial theory be a theory that could be tested by the study—a theory for which the possible data could count against. A theory could be, for example, that “cultivating kindness, intensely to a person who has wronged you, involves suppressing anger and will produce a rebound feeling of anger over the next few hours.” From the substantial theory, background assumptions are used to generate predictions. Predictions can be expressed as

specific hypotheses with defined independent and dependent variables. For example, one out of several hypotheses relevant to testing the above theory could be: If people practice kindness meditation (with script X) to a person who has wronged them rather than a stranger (the control), a subsequent bump by a confederate as they leave the room will produce more anger-relevant facial muscle activation (from 1=trace to 5=maximum).

A model is a probability distribution of effects given a hypothesis. For example, the null hypothesis, H_0 , states that there will be no difference between groups in anger on the 1-5 scale. One can represent this as a plot of plausibility (probability density) against different population group differences, with a spike at zero; the only possible difference on this hypothesis is zero (See Figure 1.) The alternative hypothesis, H_1 , states that there is a difference on the anger scale in population means between the two groups. The model of H_1 is the probability distribution for various group differences. What range of differences could be expected by the theory? Are some differences more plausible than others? Note that the theory states that the mechanism is rebound after suppression. Thus, a norming study could be run where one group is asked to consider a person who has wronged them and suppress any anger they feel; the other group does likewise for a stranger. The estimated amount of difference in rebound anger can inform the model of H_1 for the main experiment. For example, the posterior distribution of the difference in rebound anger in the norming experiment (i.e., the distribution of the uncertainty in the population mean difference) could be the model of H_1 for the main experiment. Note how theory is used to inform the model of H_1 ; indeed, how else could we test the theory than testing its actual predictions? Note also that the model of H_1 represents the information we have, to date. Thus, the Bayes factor is in that sense provisional; with better information (e.g. a larger norming study), the model of H_1 could be better informed (cf. Popper, 1959, p 275).

Figure 1 (1/4 page). Common models of H_1 and H_0 . Each graph is a plot of the plausibility (probability density) of different possible population parameter values (e.g., slopes, mean differences). Let a positive value be in the direction predicted by the theory (e.g., “compassion to enemy” group will show greater anger than “compassion to stranger” group). H_0 is a point H_0 —there is a spike of probability for just one value and no difference between the means (or zero slope). Note that H_1 allows a range of population values consistent with the theory, with smaller values more likely than larger ones.

A Bayes factor indicates the amount of evidence for the model of H_1 relative to the model of H_0 . If the background assumptions used in generating the predictions are safe, then evidence counting against H_1 relative to H_0 also counts against the theory that predicts H_1 . In this way, the theory itself can be tested. Sometimes different theories lead to the same predictions. In this case, the Bayes factor comparing those theories would be comparing the model of H_1 against the same model of H_1 , and there would not be evidence distinguishing those two theories. For example, one theory could be that the greater the kindness one has felt for a period, the less anger one would feel in the next few hours. The reason the groups in the

main study differ in anger, after a bump, is then because they cultivated different degrees of kindness. One could deal with this by asking the two groups cultivate the same specific degree of kindness - which one would have to assure by finding evidence for H_0 on a measure of kindness. Thus, the evidence by a $B > 3$ (or $> k$, whatever was good enough) for the H_1 (predicted by the substantial theory) versus H_0 supports the substantial theory relative to other theories that predict H_0 . Or more generally, a $B > 3$ for H_1 versus a H_2 (predicted by another theory) supports the theory predicting H_1 over the theory predicting H_2 .

The evidence provided by a Bayes factor comparing H_1 with H_0 supports the substantial theory over others that predict H_0 , but only to the extent the model of H_1 represents the predictions of the theory. The theory should do work in generating those predictions; without the theory, ideally other background knowledge would not make the same predictions (Popper, 1963). In this example, the theory was itself used in generating predictions; the model of H_1 was based on a norming experiment estimating rebound anger because that is the mechanism that the theory specifically postulates. The rebound norming experiment would not be relevant to a theory that postulated a difference in anger in the main study because the groups differed in amount of kindness cultivated. For example, the predictions about anger after a bump may turn out much the same in a particular study for the rebound theory and the kindness theory, but that would have to be shown.

That is, one should not think there are two entirely separate phases to theory testing—working out the predictions of a theory and statistical hypothesis testing without reference to theory. One could not get evidence for no effect unless there were a theory to specify the range of effects predicted. There is no such thing as evidence for nothing being there independent of what size thing there could be. One cannot test the absence of something in the absence of theory. Conversely, the evidence for nothing being there can be different relative to different theories. A theory that says the thing looked for can only be very small (is

there a flea bite on the arm?) will need a smaller standard error to get evidence for it not being there compared with a theory that claims it could be very large (is there a dog bite on the arm?). The theory does not have to be mathematical or computational to make relevant predictions; psychological theories predict more than researchers typically realize.

Just as several theories can make the same or similar predictions, so the prediction of a single theory can be modeled in several ways, each model just as reasonable a representation of the prediction as the others. What is needed is to show robustness of the conclusions over different models of H_1 that are roughly equally reasonable as representations of the prediction. We will discuss a simple way of doing this in the next section.

How to model H_1

Predictions can often be reasonably fixed by scientific context. Some simplifying assumptions are also useful. The precise shape of the model of H_1 is rarely relevant to a theory. A simple approach is to use a distribution with a mode of zero (e.g., a normal distribution—presuming zero is the value predicted by H_0). This means the shape puts most probability around the same value as H_0 ; thus, the shape can make it slightly harder to discriminate H_1 from H_0 . That is, when the Bayes factor does give good evidence for the one hypothesis rather than the other, it was despite rather than because of the precise model of H_1 . If the theory predicts a direction of the effect, the distribution below zero can be removed, leaving a half-normal distribution (see Figure 1). The standard deviation of the normal or half-normal distribution scales the rate at which the curve drops towards zero. Thus, one sets the standard deviation to the approximate scale of effect predicted. Only 5% of the area of distribution is beyond two standard deviations out; thus, twice the scale of effect is

also the rough maximum. Another way of putting this is, if one has information to specify a rough maximum, set the standard deviation of a half-normal to half that maximum.

A rough maximum can often be estimated by determining the possible “room to move” (Dienes, 2019). Klaschinski et al. (2017) investigated whether briefly assuming an expansive rather than contracting pose would increase one’s performance in a subsequent interview. Participants were interviewed, and raters evaluated how well the subject performed on a 1 (“awful”) to 7 (“amazing”) scale. The theory tested was that assuming a posture typical of being confident rather than insecure would make one perform better in a demanding situation. People in the power posing condition were rated on average 4.23 Likert units; those in the control condition were 4.17 Likert units. The mean difference is 0.06 Likert units (SE = 0.21 Likert units) and is non-significant at the 5% level ($t = .06/0.21 = 0.28$, $p = .78$). A non-significant difference by itself allows no conclusion about whether there is a population difference. To see which way the evidence points, the possible size of the population differences needs to be determined. The control condition had a mean of about 4; thus, the room to move—the biggest difference if the theory were true—would be if the power pose group scored the maximum on the scale (7), giving a maximum possible difference of $(7-4) = 3$ Likert units.

Given this estimated maximum, the plausibility of different possible population differences, if the theory were true, can be modeled as a half-normal with a mode of zero and an SD of $\max/2 = 3/2 = 1.5$ Likert units. The “half” of the half-normal indicates the theory predicts in a certain direction; performance is predicted to be higher in the expansive rather than contracting pose condition. The SD of the half-normal indicates the rough scale of effect expected (1.5 Likert units). The Bayes factor can be calculated and written as $B_{\text{HN}(0,1.5)} = 0.18$, where the “HN” indicates a half-normal was used to model H_1 , the “0” refers to the mode and the “1.5” to its SD. To obtain this Bayes factor, go to www.bayesfactor.info for a ShinyApp.

Enter “0.06” for the “mean difference”, “0.21” for the “standard error”, “1.5” for the “hypothesized mean difference” and click “positive 1-tailed” to make it half-normal. Here, the work done by the theory is in indicating the direction of the effect; the constraints in other aspects of the data themselves then limit how big the effect could be. 0.18 is conventionally considered as moderate evidence for no effect, as it is less than 1/3.

In fact, in this case, there is more information. Klaschinski et al (2017) was a replication of Cuddy et al. (2015). In the original study, Cuddy et al. found the power posing group were rated as 4.63 Likert units; the control group 3.81 units. Thus, the difference was 0.8 Likert units with a standard error of 0.28 Likert units. Using 0.8 Likert units as a scale factor (instead of the previous 1.5), $B_{HN(0,0.8)} = 0.32$. Since this is less than a third, the same qualitative conclusion results—moderate evidence for no effect. In fact, for this second Bayes factor, the theory is slightly different. Using the effect from a previous study to inform a replication attempt tests the theory that the methods given in the Method section of the original study lead to the sort of results given in the Results section (call this the replication hypothesis; cf. Popper, 1959, p. 66).

As the original study was an attempt to test the same substantial theory, the test of the replication hypothesis is relevant to the substantial theory (for interpreting evidence against a replication hypothesis, see Edlund et al., 2021). In any case, the use of 1.5 for the scale factor in the first Bayes factor is somewhat arbitrary. How robust is the conclusion to different scale factors? One can report a Robustness Region—the set of scale factors that lead to the same qualitative conclusion (e.g., $B < 1/3$; Dienes, 2019); in this case, it can be notated $RR_{B < 1/3} [0.8, >6]$. The “0.8” indicates that 0.8 is the lowest the scale factor can be and there still be moderate evidence for H_0 ; the scale factor can exceed the length of the scale (1-7), and there would still be evidence for H_0 . The lowest scale factor implies the effect could plausibly be

between 0 and 1.6 Likert units (i.e., from 0 to twice the scale factor). If there is no compelling reason why the effect must be more constrained than this, the conclusion is robust.

How does one construct a model of H_1 ? One uses assumptions that are either simple or scientifically motivated to derive predictions. The predictions can then be tested against the observed effect. One might be tempted to think “I do not know what effect sizes my theory predicts; the observed effect is the best information I have about its size, so why don’t I use this to model H_1 ?” Then, the predictions can never clash with the data, so the data can never count against the theory. To test the theory, we need the model of H_1 to possibly clash with the data. In the power posing example, if we used the observed effect of 0.06 Likert units as the scale factor, we obtain $B_{HN(0, 0.06)} = 1.02$; this does not count against the theory. When we use the room-to-move heuristic, by contrast, even though we are drawing on an aspect of the same data, this does not compromise the ability of the test to find evidence against the theory—as indeed it did in this example (Dienes, 2019; cf. Devezet et al., 2020).

Severe testing

Popper (1963) defined a severe test as one in which, if the theory were false, there is a high probability the theory will be found false. Popper (p. 526) formalized the notion of a severe test by representing it as the ratio of the probability of the predicted outcome given the theory to the probability of the outcome assuming the theory were false (and assuming the rest of background knowledge). That is, a severe test is one that can generate an extreme Bayes factor (i.e., one that is very large or very small) depending on whether the theory is true or false. Bayes factors, as a conceptual tool, therefore go hand in hand with determining the conditions for a severe test (van Dongen et al., 2020; Vanpaemel, 2020).

Assume that all modelling assumptions are satisfactory (e.g., the distribution of the data). If H_0 is true, a Bayes factor is, with increasing subjects, eventually driven to 0. That is, if a theory is false (H_0 is true), with enough data, a Bayes factor is driven to show the theory false. Likewise, if there is a difference, then the Bayes factor is, with increasing subjects, eventually driven to infinity. That is, if H_0 is false, with enough data, H_0 is eventually found false, but this only occurs if our models represent everything relevant in an approximately good enough way. One crucial aspect of the modelling is how subjects' data are generated by the world for a given population mean difference (or value of whichever parameter is being tested). If a sample mean difference of 0.8 is found, then how plausible is this mean for different possible population mean differences? This is the likelihood function.

This function does not take on a simple form when researchers cherry pick and hack. For example, assume the researcher determines how the sample mean difference can be made as close to a predicted value as possible, by playing with different ways of excluding outliers, removing participants according to different equally sensible exclusion criteria, adding or taking away covariates, and so on. “*B*-hacking” is this process of trying different analyses to push *B* in the direction one wants. With this new likelihood function, the probability of obtaining the “data”—represented as the finally cited mean difference—may be very similar given H_0 as given H_1 .

Consider a case, after some *B*-hacking, where an uncorrected Bayes factor (i.e., one that does not consider hacking) shows evidence for H_1 , even when H_0 is true. The Bayes factor, using the appropriate likelihood function (i.e., considers hacking, how cited mean differences are generated given a population parameter value), should be close to 1, and the whole process thereby is uninformative. The test would in no way be severe, as a false theory would not often be found false. Thus, the concept of a Bayes factor, by measuring evidence,

shows why cherry picking and hacking render data non-evidential. The “data” are just about as probable given H_0 as H_1 .

Cherry picking. To get the appropriate Bayes factor, hacking and cherry picking must be considered, as they influence the likelihood function. Modelling the effects of hacking may be difficult in a real world case. So, rather than try to model the effects of generic hacking, it is better to set up safeguards to protect against their effect or, in the absence of safeguards, to acknowledge the potential role of *B*-hacking. In terms of safeguards, one can pre-register the analytic protocol, including the modelling of H_1 , or use a fully-fledged Registered Report (Chambers, 2019; Dienes, 2021b). Alternatively, one can use blind analyses, in which the analyst is given several data sets and does not know which is the real one (MacCoun & Perlmutter, 2015). One may be able to rely on simplicity and strong theoretical arguments (Szollosi et al., 2020), though, given the ease of hacking in even simple situations, this is unlikely to be a complete solution (Wagenmakers, 2019). In terms of acknowledging the issue, one may treat the result as more tentative than the Bayes factor indicates at face value and then seek to replicate it; the original result could, for example, be published as an Exploratory Report (McIntosh, 2017).

Stopping rule. The formula for the Bayes factor given at the beginning of this chapter showed that is a complete summary of the evidence—how much one should change one’s beliefs, assuming that the axioms of probability represent idealized constraints on strength of belief. Thus, all that matter for evidence is the probability of the data given the hypotheses contrasted. In particular, it is not relevant to interpreting a Bayes factor as evidence when one stops collecting data. That is, one need not state, in advance, a stopping rule when using Bayes factors (Dienes, 2016; Hendriksen et al., 2020; Rouder & Haff, 2020). Optional stopping, a questionable research practice for significance testing, is not one for Bayes factors. In this way, the evidence indicated by Bayes factors has the same property as

evidence does in all other arenas - apart from significance testers doing research (Wagenmakers et al., 2019). In no other case is evidence adjusted according to a stopping rule, be it detectives searching for evidence for the murderer, children learning to speak, animals learning where the best food is, or significance testers learning a foreign language in their spare time. Optional stopping with Bayes factors (e.g., stopping when the Bayes factor is greater than 10 or less than 1/10 to ensure strong evidence) allows severe testing of either H_1 or H_0 .

Multiple testing. Multiple testing of different hypotheses, all relevant to a theory, also shows the way in which Bayes factors reflect severe testing requirements. Consider the substantial theory that extrinsic motivation reduces intrinsic motivation because, if people see that they behave a certain way to get an extrinsic reward like money, they conclude they do not want to behave that way for its own sake. An experiment puts children in a room and rewards them for playing with some toys rather than others—randomly selected for different groups. The children are then observed surreptitiously for how much time they play with each toy when they think they are alone. It is found that children play less with the rewarded than the control toys but only for females in the afternoon. The Bayes factor for that specific contrast is greater than 3. The authors conclude that the theory is supported, perhaps especially for people more sensitive to the sorts of rewards used. This is a case of multiple testing.

Multiple testing can result in increasingly implausible hypotheses being tested, including hypotheses that do not follow simply from the substantial theory. The test of the substantial theory is the overall difference in time playing with rewarded rather than control toys. All data relevant to the theory must be included in evaluating the theory. Researchers may be vaguely aware that significant results carry more evidential value than non-significant ones, so they may think they can focus on significant results and safely ignore the non-

significant results. However, cherry picking is wrong in any school of statistical inference. From a Bayesian perspective, it is apparent that all data relevant to testing the theory provide some degree of evidence. Thus, the test of the theory is provided by all the data (see Dienes, 2008, pp 108-114, and Dienes, 2016, for different examples). If the theory, as stated, is wrong, the data as a whole, reflecting the simplest test of the theory, is likely to show it wrong, given enough data.

In this example, the Bayes factor for the main effect of reward may be less than 1/3, in which case the evidence counts against the theory. There is also evidence for a peculiar combination of conditions leading to a reward effect, but this cannot be used in isolation as support for the theory. What often makes the results of multiple testing seem unsatisfactory is the low prior probability of the hypotheses that get evidence, given the other hypotheses that had evidence against them. Alternatively, multiple testing often finds support for hypotheses that do not follow in a simple way from a simple theory. One might explicitly consider the prior odds of the different hypotheses and thereby demand more evidence (i.e., a higher Bayes factor threshold). In the example considered here, the simplest approach, rather than adjusting thresholds, may be to acknowledge that the substantial theory has evidence against it and leave it at that. Alternatively, one could conjecture a simple and bold theory that might explain why females in the afternoon and not other combinations show the effect, derive predictions from such a theory, and test them in a further study.

Notice that no general Bonferroni correction (or other familywise correction procedures) quite gets the point right in dealing with multiple comparisons. The Bonferroni correction used by significance testers says one should use a significance threshold of $.05/k$ if one conducts k tests in a family of tests. What is a family of tests (cf Kruschke, 2011)? A temptation might be to say that a family is all the tests relevant to a theory, but theories come in hierarchies; there is the most general theory (e.g. dissonance theory), there is an

application of that theory to a particular problem (intrinsic motivation), and there are more specific applications (children playing with toys). There is no single answer to what k should be for theory testing because there are simultaneously several theories at play (Dienes, 2016). Imagine a paper that looked at children with toys and adults with artwork, but significance testing demands a single once-and-for-all answer. A Bayes factor is a measure of evidence relative to theories—if you want to know the evidence for theory X , consider all data relevant to theory X . One can simultaneously consider the evidence for different theories in a hierarchy with different Bayes factors. Because significance testing considers tests must be controlled relative to a person, it fails to be able to measure evidence relative to theories.

Consider a theory tested by asking five questions on a Likert scale. Each question alone may provide some evidence for the theory (e.g., $B > 3$ and $p < .05$ for each question). With a Bonferroni correction (requiring $p < .05/5$, or $B > 3 \times 5$), each question may fail to provide good enough evidence for the theory. However, when the questions are combined together in a single measure, the evidence would be $B > 3$, $p < .05$. That is, forming an overall conclusion by applying Bonferroni to the individual questions is inappropriate in this case; the automatic application of familywise error correction does not solve the problem of multiple testing.

Where the prior probability of one hypothesis being true can be roughly assessed by the number of hypotheses tested, a Bonferroni correction can be useful (cf. Westfall et al., 1997). For example, when looking at 20,000 different genes to see if any correlates with a phenotype, it makes sense to take into account how improbable each H_1 is—approximately $1/20,000$. In the context of such implausible individual hypotheses, a more severe than typical test is called for—a Bayes factor threshold of $3 \times 20,000$ can be used. In general, when data dredging, one should consider prior probability because the whole point of dredging is most hypotheses are wrong.

So how can there be evidence for theories that go beyond the data?

Scientific theories are interesting because they are bold and go beyond the data (Popper, 1963, p 330). After finding that his subjects were slower at naming colors when the words were incongruent, but were not slower at naming words when the colors were incongruent, Stroop (1935) did not say “My theory is that the students at George Peabody College for Teachers in 1936 ...”; instead he said that, in general, “the associations that have been formed between the word stimuli and the reading response are evidently more effective than those that have been formed between the color stimuli and the naming response” (p 660). According to Popper, he was not inducing the claim (and hence having to stay close to the data) but rather testing it. As the claim is meant to be general, if it fails for students at George Peabody College for Teachers, it fails; Stroop’s study constituted a test of the claim. The more general the claim, the more opportunities there are for finding Bayes factors that go against it. Hence bold general claims are easier to severely test; more situations can constitute tests of general claims than claims that “stay close to the data.”

When testing a theory, one conjectures a world in which the theory may be true or false. There is a probability structure that arises because of the way that conjectured world generates data. Within that world, probabilities change in the light of data. In that sense, contrary at face value to Popper’s claims that induction does not exist, induction can occur in the sense of changing probabilities of theories in the light of data (Jeffreys, 1939). That is what evidence is—something that normatively changes strength of belief. But Popper is right. The induction does not occur in any absolute sense; it occurs in the conjectured world. In that sense, every statistical test is a thought experiment (Greenland, 2017). It is a conjectured world, but part of the conjecture is that the conjectured world is a good enough

approximation of the real world so that conclusions in the conjectural world are relevant to the real world. Ultimately, the only use of the conjectural world is to test claims about the real world.

How can we check if the conjectured world is relevant to the real world? The only option is to test its assumptions (Morey et al., 2013; Notturmo, 1999). We embed the conjectured world in a larger also-conjectured world (Kruschke, 2013a). For example, the larger world may assume different degrees of skew in the likelihood distribution while the world in which we conducted our test may assume a normal likelihood; both worlds assume a certain family of distributions. In this way, each assumption can, in principle, be tested one by one. Significance tests are often used by researchers for checking assumptions, yet a significant violation does not mean that the violation was extreme enough to make any difference to conclusions; a non-significant result does not mean there was not a violation that makes a huge difference to conclusions (Kruschke, 2013a). That is, as in any area of testing, to test if there is nothing there, one must consider how big the thing is that is relevant. Was the violation large enough to alter conclusions?

More work is needed on how well calibrated Bayes factors remain depending on degrees of violation of assumptions (Rouder & Haaf, 2020). One rule of thumb is that the same conditions that lead to error rates changing markedly for significance tests are likely to be the same that lead to the equivalent Bayes factors to be badly calibrated. For example, in conditions where one should use adjusted degrees of freedom for a t-test, because of variance inequality, use those same adjusted degrees of freedom for the corresponding Bayes factor. This is because, if one makes simplifying assumptions, Bayes factors become monotonic with p values (Benjamin et al., 2017); thus, whatever conditions influence the correctness of those p values will influence the corresponding Bayes factors. If those Bayes factors are so affected, so are the Bayes factors one would use in real situations. Piece meal testing of

assumptions (Mayo, 2018) can help determine where one's conjectured world ceases to be relevant to the real world.

In deriving predictions from a theory, and thereby setting up a conjectural world, the probabilities in that world (e.g., in the form of the model of H_1) are idealized subjective probabilities—strengths of belief assigned to an idealized person confronting the relevant background knowledge. Once postulated, the model becomes objective, in the sense that its consequences can be discovered and the assumptions criticized by anyone (Popper, 1972). A model of H_1 is meant to follow from theories using simple and otherwise well tested assumptions; what matters are the reasons why the assumptions are made, not what any specific individual believes. In that sense, one need not follow a purely subjective Bayesian approach that treats predictions (as models of H_1) as purely subjective feelings of what size an effect may be. One does need to make a judgment that the model is satisfactory enough to use it, but any judgment against an assumption is simply a promissory note that reasons can be found for criticizing that assumption; what will matter in the end are those reasons that can be made public and criticized by anyone else (Miller, 1999; Notturmo, 1999). Similarly, one need not follow a purely objective Bayesian approach that uses default models of H_1 —a single model of H_1 for any theory (e.g. a Cauchy with scale factor 1, as used by Rouder et al, 2009). Testing theories means representing the predictions of specifically those theories and confronting those predictions with data.

Criticisms of Bayes factors

1. Different models of H_1 give different answers.

This is sometimes called the problem of prior sensitivity; if you change the model of H_1 , you get different Bayes factors—so, how can you know what the evidence is? (Kruschke,

2013b). However, this is not a fault of Bayes factors but a virtue. Different theories, or different assumptions connecting theory to predictions, make different predictions. It is useful and necessary that different predictions lead to different degrees of evidence for different theories. That is, the model of H_1 needs to represent the predictions of the theory put to test. The problem then becomes one of knowing what your theory predicts. Dienes (2019, 2021a) discuss different heuristics for modelling H_1 given different theories and scientific contexts. Inferential ambiguity can arise when the same prediction can be equally well modelled in different ways; one way of partially addressing this is with the Robustness Region mentioned above (Dienes, 2019). Also, as mentioned above, prior sensitivity provides extra analytic flexibility and opportunities for *B*-hacking. This is partly addressed by Robustness Regions and should also be addressed by other methods for dealing with analytic flexibility (e.g., pre-registration; see previous discussion; Dienes, 2021b).

2. A default Bayes factor is not relevant to your theory (e.g., Tendeiro & Kiers, 2019).

Make sure the model of H_1 represents the predictions of your theory. Always explicitly state what your model of H_1 is, and give an objective reason why you set any parameter value. Do not use a default scale factor just because it is already there in the software; treat it as an invitation for you to consider whether the default settings are relevant.

The urge to use defaults may be based on regarding inferential statistics as insulated from theory: Whatever the theory, one conducts the standard inferential statistics, and they indicate whether there is an effect there or not. On this insulating approach, conclusions about whether or there is an effect can then be used to count for or against theoretical predictions, no matter the theory. This simplistic separation of statistics from theory may be one reason why we have been having problems properly testing our theories.

3. The point H_0 is never true; therefore, the only thing to do is to estimate

Meehl (1967) argued that, in the case of real world correlations, all point H_0 's are false (i.e., the H_0 of there being exactly no difference or no slope); surely, all point H_0 's are false to some decimal point for any hypothesis in the social and behavioral sciences if only because modeling assumptions are never exactly true. This claim can be used to argue for the futility of performing any hypothesis test against the point H_0 , whether using significance testing or Bayes factors. On this view, one should always just reject the point H_0 and be done with it (see Baguley, 2012, p 368; Wagenmakers, 2017, for arguments and replies).

Bayes factors can compare any two models. The point H_0 is often used as one of the models to test the claim that there is something there rather than nothing. Baguley (2012) illustrates how accurate a point H_0 can be with Wiseman and Greening's (2002) Online experiment with 27,856 participants, testing the existence of ESP with a chance baseline of 50%. The 95% CI was [49.6%, 50.2%]. That is, even though the statistical assumptions that lead to an estimated predicted 50% for no ESP must be approximate, whatever the true H_0 is, it lies within an interval of roughly [-.3%, +.3%] around 50%. That means we can trust the same modeling assumptions in other similar contexts to be accurate within this interval. Now, imagine the same modeling assumptions with a similar experiment, investigating a more mundane effect, with a roughly expected effect size of about 5% above baseline. Further, one might decide that a minimally interesting effect would be any value 0.5% above baseline. Results just give support for H_1 using a point H_0 : 55% correct $SE = 2.5\%$, $B_{H(0,5)} = 4.27$ against a point H_0 . A different web Bayes factor calculator will be illustrated this time to obtain the Bayes factor, as this calculator allows interval H_0 's. Go to <https://bayesplay.colling.net.nz/>. For likelihood, click on "normal"; enter "5" for mean, and

“2.5” for SD. For alternative prior, click on “normal”. Enter “0” for mean and “5” for SD. To make it a half-normal, click on “lower limit” and enter “0”. For null prior, enter “point”. Enter “0” for point. Click “calculate”.

A Bayes factor compares any two models, so we can compare a similar model of H_1 (but with a mode of 0.5 above 50%) against an interval H_0 around 50% $[-0.5\%, +0.5\%]$. This gives $B_{H(0.5,5)} = 4.44$ —virtually the same answer. To obtain this Bayes factor, proceed in the same way as for the previous paragraph; go to null prior, click “uniform”, and enter “-0.5” for the lower limit and “0.5” for the upper limit. Finally, for the lower limit for H_1 , change to “0.5”. Click calculate. In sum, the point H_0 was a perfectly adequate approximation of the true interval H_0 . This will be generally true whenever the standard error is large compared to the null interval. The point null will often be a perfectly adequate approximation in well-controlled experimental research; this is similar to assuming any particular distribution, or any particular statistical model at all, is a useful approximation in any research. If we are going to reject all point H_0 s because they are not exactly true, we must similarly reject all models whatsoever. When big data are involved, the standard error may be very small; then, it will be important to be clear about how H_0 should be modeled—maybe with a uniform—and the plausibility of H_1 may be zero below a minimally interesting value (see Palfi & Dienes, 2019, for Bayes factors with interval H_0 ’s and an example case; see Skora et al., 2020, for another type of example using interval H_0 ’s).

Hypothesis testing is needed whenever a researcher wants to ask whether something exists—Should a term be in the model? Is there an interaction? Are people at chance? If the question of something existing can be taken for granted, then all one need do is estimate. However, if one only estimates, do not infer from the estimation that there is no effect nor treat the estimation as grounds for asserting that something exists rather than not existing at all (existence has been presumed by estimation, not tested). With these provisos, a paper with

only estimation may be a useful option. Can you draw all the conclusions you want from saying whatever the effect is, it is plausibly between such and such bounds? It is an option I have followed for several papers (e.g., Palfi et al., 2020).

4. A Bayes factor does not control error rates.

Mayo (2018) criticized Bayes factors for not controlling error rates. A Bayes factor measures strength of evidence; what you do with evidence (over many cases) determines error rates. Fixing one (evidence or error rates) at a desired value does not fix the other at any particular value. The intuition behind criticizing Bayes factors for not precisely controlling error rates is presumably that given error rates are not fixed, and one should be more cautious in updating beliefs from Bayes factors. In sum, the critic says: Don't take Bayes factors seriously. But this is backwards. Evidence, not error rates, is how much you should change your confidence in H_1 versus H_0 . The critic's intuition presumably arises because the more evidence one has in the long run, the lower error rates tend to be; one may mistakenly feel that, unless error rates are precisely fixed, evidence cannot be "true" evidence.

Evidence constrains error rates without fixing them. In fact, Bayes factors can minimize the weighted sum of Type I and II errors (DeGroot, 1986, p. 444 for likelihood ratios; Pericchia & Pereira, 2016 for Bayes factors). Evidence and error rates are related, and strong evidence will be associated with small (but not fixed) error rates (e.g. see Tables in Dienes, 2016). Still, if one's interest is in evidence, it is Bayes factors that measure it. Error rates may be useful in understanding what a certain quantity of evidence means (cf. Hendriksen et al., 2020); if error rates seem too high, then that level of evidence may be too low for your needs. It is instructive to consider how, with evidence less than infinite, there must always be a probability that the evidence reverses as more data comes in; with evidence

as low $B = 3$ (roughly, $p=.05$), the reversals are reasonably common (see Tables in Dienes, 2016).

In sum, there are no convincing arguments against Bayes factors as such, only against their misuse or misunderstanding.

Conclusion

I personally follow a policy of a “ B for every p ”. When hypothesis testing, I report both p values and Bayes factors, whose model of H_1 is informed by the theories tested. Inferences are always with respect to the Bayes factors, but p values are given so people can see their relation to Bayes factors. In fact, there is a Bayesian interpretation of p values. Imagine we are using half-normal models of H_1 in a paper. If a result is significant $p<.05$, then there is some standard deviation for the half-normal for which $B > 3$; likewise, if $p<.01$, then there is some standard deviation for which $B > 10$. That standard deviation may not be relevant for the theory, or it may not be the only relevant one; only the Bayes factors with scientifically motivated standard deviations will measure evidence relevant to the theory. Still the p value carries information, so it is useful to quote it.

If one decides a threshold for regarding evidence as adequate (e.g., 3), then one can say that when $B > 3$, “There was evidence for a main effect of ...”; when $B < 1/3$, “There was evidence for no main effect ...”; and when B is between those values (i.e., close to 1), “There was no evidence one way or the other for a main effect...” In the latter case, one would draw no theoretical conclusions in the discussion about whether there is an effect (see Dienes, 2021a for examples). However, when $B < 1/3$, one could use that result to count against a theory that predicted the H_1 tested. That is, theories can actually be tested.

Bayes factors are also useful when more participants are needed. If one has already collected data, and a reviewer asks for more, significance testing is now illegitimate. For the same reason, combining studies together for significance testing is also illegitimate. The stopping rule must be respected for significance testing. Thus, evidence is wasted. How can a theory be severely tested if the data fall short of severely testing it and it is forbidden to collect more? Because Bayes factors measure evidence, and one can always accumulate evidence until one has enough, Bayes factors must be used for hypothesis testing once data is to be combined.

In sum, Bayes factors are an invaluable practical and conceptual tool for understanding, and practically engaging in, the severe testing of theories.

REFERENCES

- Baguley, T. S. (2012). *Serious Stats: A Guide to Advanced Statistics for the behavioral Sciences*. London: Palgrave Macmillian.
- Benjamin, D. J., Berger, J. O., Johannesson, M., et al. (2017). Redefine statistical significance. *Nature Human Behaviour*, doi: 10.1038/s41562-017-0189-
- Chambers, C. D. (2019). What's next for registered reports? *Nature*, 573 (7773), 187-189
- Cortex. (2021).
http://cdn.elsevier.com/promis_misc/PROMIS%20pub_idt_CORTEX%20Guidelines_RR_29_04_2013.pdf Downloaded 19 March 2021.
- Cowles, M., & Davis, C. (1982). On the Origins of the .05 Level of Statistical Significance. *American Psychologist*, 37, 553-558.
- Cuddy, A. C., Wilmuth, C. A., Yap, A. J., & Carney, D. R. (2015). Preparatory power posing affects nonverbal presence and job interview performance. *Journal of Applied Psychology*, 100, 1286–1295.
- DeGroot, M. H. (1986). *Probability and statistics, 2nd edition*. Massachusetts: Addison-Wesley.
- Devezer, B., Navarro, D. J., Vandekerckhove, J., & Buzbas, E. O. (2020). The case for formal methodology in scientific reform. bioRxiv 2020.04.26.048306; doi: <https://doi.org/10.1101/2020.04.26.048306>
- Dienes, Z. (2016). How Bayes factors change scientific practice. *Journal of Mathematical Psychology*, 72, 78-89.

- Dienes, Z. (2019). How do I know what my theory predicts? *Advances in Methods and Practices in Psychological Science*, 2, 364-377.
- Dienes, Z. (2021a). How to use and report Bayesian hypothesis tests. *Psychology of Consciousness: Theory, Research, and Practice*, 8, 9–26
- Dienes, Z. (2021b; first draft). The inner workings of Registered Reports. To appear in Austin Lee Nichols & John E. Edlund (Eds), *Cambridge Handbook of Research Methods and Statistics for the Social and Behavioral Sciences*, <https://doi.org/10.31234/osf.io/yhp2a>
- Dienes, Z., & McLatchie, N. (2018). Four reasons to prefer Bayesian over significance testing. *Psychonomic Bulletin & Review*, 25, 207-218.
- Edlund, J., Cuccolo, K., Irgens, M. S., Wagge, J. R., & Zlokovich, M. S. (2021). Saving Science Through Replication Studies. *Perspectives on Psychological Science*, <https://doi.org/10.1177/1745691620984385>
- Fisher, R. A. (1925). *Statistical methods for research workers*. Edinburgh: Oliver & Boyd.
- Greenland, S. (2017). The Need for Cognitive Science in Methodology. *American Journal of Epidemiology*, 186, 639–645.
- Hendriksen, A., de Heide, R., & Grünwald, P. (2020). Optional Stopping with Bayes Factors. *Bayesian Analysis*, (available at <https://arxiv.org/pdf/1807.09077.pdf>)
- Jeffreys, H. (1939). *The theory of probability*. Oxford, England: Oxford University Press.
- Klaschinski, L., Schnabel, K., & Schröder-Abé, M. (2017) Benefits of power posing: effects on dominance and social sensitivity, *Comprehensive Results in Social Psychology*, 2, 55-67.
- Kruschke, J. K. (2011). Bayesian Assessment of Null Values Via Parameter Estimation and Model Comparison. *Perspectives on Psychological Science*, 6, 299-312.

- Kruschke, J. K. (2013a). Posterior predictive checks can and should be Bayesian. *British Journal of Mathematical and Statistical Psychology*, 66, 45–56.
- Kruschke, J. K. (2013b). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, 142, 573–603
- Lindley, D. V. (1957). A statistical paradox. *Biometrika*, 44, 187–192.
- Lindley, D. V. (2014). *Understanding uncertainty, revised edition*. New Jersey: Wiley.
- MacCoun, R., & Perlmutter, S. (2015). Hide results to seek the truth. *Nature*, 526, 187–189.
- Mayo, D. G. (2018). *Statistical inference as severe testing: How to get beyond the statistics wars*. Cambridge: Cambridge University press.
- Meehl, P. E. (1967). Theory testing in psychology and physics: A methodological paradox. *Philosophy of Science*, 34, 103–115.
- McIntosh, R. D. (2017). Exploratory reports: A new article type for Cortex. *Cortex*, 96, A1–A4.
- McPhetres, J., Albayrak-Aydemir, N., Barbosa Mendes, A., Chow, E. C., Gonzalez-Marquez, P., Loukras, E., ... Volodko, K. (2020). A decade of theory as reflected in Psychological Science (2009-2019). *PLOS One*, 10.1371/journal.pone.0247986
- Miller, D. (1999). *Critical Rationalism: A Restatement and Defence*. Illinois: Open Court.
- Morey, R. (2018). <https://medium.com/@richarddmorey/redefining-statistical-significance-the-statistical-arguments-ac9007bc1f91>
- Morey, R. D., Romeijn J. W., & Rouder J. N. (2013). The humble Bayesian: Model checking from a fully Bayesian perspective. *British Journal of Mathematical and Statistical Psychology*. 66, 68--75

Morey, R. D., Romeijn, J.-W., & Rouder, J. N. (2016). The philosophy of Bayes factors and the quantification of statistical evidence. *Journal of Mathematical Psychology*, 72, 6–18.

Notturmo, M. A. (1999). *Science and the Open Society*. Budapest: Central European University Press.

Palfi, B., & Dienes, Z. (2019, December 10). When and how to calculate the Bayes factor with an interval null hypothesis. <https://doi.org/10.31234/osf.io/9chmw>

Palfi, B., Moga, G., Lush, P., Scott, R. B., & Dienes, Z. (2020). Can Hypnotic Suggestibility Be Measured Online? *Psychological Research*, 84, 1460–1471.
<https://doi.org/10.1007/s00426-019-01162-w>

Pericchia, L., & Pereira, C. (2016). Adaptive significance levels using optimal decision rules. *Brazilian Journal of Probability and Statistics*, 30, 70–90.

Popper, K. R. (1959). *The logic of scientific discovery*. London: Hutchinson

Popper, K. R. (1963). *Conjectures and Refutations: The growth of scientific knowledge*. London: Routledge..

Popper, K. R. (1972). *Objective Knowledge: An Evolutionary Approach*. Oxford: Oxford University Press

Rouder, J., & Haaf, J. M. (2020). Optional Stopping and the Interpretation of the Bayes Factor. <https://doi.org/10.31234/osf.io/m6dhwR>

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225–237.

- Skora, L., Livermore, J. J. A., Dienes, Z., Seth, A., & Scott, R. B. (2020). Feasibility of Unconscious Instrumental Conditioning: A Registered Replication. *Cortex*, <https://doi.org/10.31234/osf.io/p9dgn>
- Stigler, S. M. (1999). *Statistics on the table: The history of statistical concepts and methods*. London: Harvard University Press.
- Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*, 643–662
- Szollosi, A., Kellen, D., Navarro, .. & Donkin, C. (2020). Is preregistration worthwhile? *Trends in Cognitive Sciences*, *24*, 94-95.
- Tendeiro, J. N., & Kiers, H. A. L. (2019). A review of issues about null hypothesis Bayesian testing. *Psychological Methods*, *24*(6), 774–795.
- van Dongen, N. N. N., Wagenmakers, E., & Sprenger, J. (2020, December 16). A Bayesian Perspective on Severity: Risky Predictions and Specific Hypotheses. <https://doi.org/10.31234/osf.io/4et65>
- Vanpaemel, W. (2020). Strong Theory Testing Using the Prior Predictive and the Data Prior. *Psychological Review*, *127*, 136–145, <http://dx.doi.org/10.1037/rev0000167>
- Wagenmakers, E. (2017). <https://jasp-stats.org/2017/10/25/test-interval-null-hypotheses-jasp/> Downloaded 23 March 2021.
- Wagenmakers, E. (2019). <https://www.bayesianspectacles.org/a-breakdown-of-preregistration-is-redundant-at-best/> Downloaded 20 March 2021
- Wagenmakers, E., Gronau, Q. F., & Vandekerckhove, J. (2019, March 7). Five Bayesian Intuitions for the Stopping Rule Principle. <https://doi.org/10.31234/osf.io/5ntkd>

Westfall, P. H., Johnson, W. O., & Utts, J. M. (1997). A Bayesian perspective on the Bonferroni adjustment. *Biometrika*, 84, 419–427.

Wiseman, R., & Greening, E. (2002) The Mind Machine: A Mass Participation Experiment into the Possible Existence of Extrasensory Perception. *British Journal of Psychology*, 93, 487–99.