© 2022, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: [10.1037/met0000505](https://doi.org/10.1037/met0000505).

On the white, the black, and the many shades of gray in between: Our reply to van
Ravenswaaij and Wagenmakers (2021)

Jorge N. Tendeiro

Hiroshima University

Henk A. L. Kiers

University of Groningen

Correspondence concerning this article should be addressed to:

Jorge N. Tendeiro

Hiroshima University

Office of Research and Academia-Government-Community Collaboration,
Education and Research Center for Artificial Intelligence and Data Innovation

1-3-2 Kagamiyama, Higashi-Hiroshima 739-8511, Japan

Email: tendeiro@hiroshima-u.ac.jp

Tlf: +81 824245263

Author Note

The contents of this paper are a direct follow-up of two other papers published at *Psychological Methods* (Tendeiro & Kiers, 2019; van Ravenzwaaij & Wagenmakers, 2021). This work has only been disseminated through sharing a preprint at PsyArXiv (Tendeiro & Kiers, 2022; <https://psyarxiv.com/tjxvz/>). This research was supported by a Japanese JSPS KAKENHI grant awarded to Jorge N. Tendeiro (grant number 21K20211).

Abstract

In 2019 we wrote a paper (Tendeiro & Kiers, 2019) in *Psychological Methods* over null hypothesis Bayesian testing and its working horse, the Bayes factor. Recently, van Ravenzwaaij and Wagenmakers (2021) offered a response to our piece, also in this journal. Although we do welcome their contribution with thought-provoking remarks on our paper, we ended up concluding that there were too many ‘issues’ in van Ravenzwaaij and Wagenmakers (2021) that warrant a rebuttal. In this paper we both defend the main premises of our original paper and we put the contribution of van Ravenzwaaij and Wagenmakers (2021) under critical appraisal. Our hope is that this exchange between scholars decisively contributes towards a better understanding among psychologists of null hypothesis Bayesian testing in general and of the Bayes factor in particular.

On the white, the black, and the many shades of gray in between: Our reply to van Ravenzwaaij and Wagenmakers (2021)

We have recently published a paper in *Psychological Methods* over null hypothesis Bayesian testing (NHBT) and its working horse, the Bayes factor (Tendeiro & Kiers, 2019, henceforth TK). Our paper offered some insight into the inner workings of NHBT, as well as a discussion of some of the properties of the Bayes factor. TK was in general well received and we are still very happy to have contributed towards a better understanding of this relatively new inferential framework. However, our work has also been criticized by some researchers who have different views about the properties and merits of the Bayes factor. Don van Ravenzwaaij and Eric-Jan Wagenmakers in particular wrote a commentary to our paper, published also here in *Psychological Methods* (van Ravenzwaaij & Wagenmakers, 2021, henceforth vRW). First and foremost, we would like to thank the authors of the commentary for offering their views over our paper. We have read this commentary with great care and we concluded that there are several points that require clarification. Below is our detailed account on our own reading of vRW.

We have organized our reply as follows. First we give a brief overview of the main points made in TK and vRW. This is necessary in order to provide a proper context to what follows. Next, we provide specific details for each of the original 11 issues discussed in TK. We draw on both the communalities between us (and there are various) as well as on what sets us apart. Importantly, we clarify some points made by vRW that, in our opinion, do not accurately reflect what we have discussed in our paper. We end this paper with some concluding remarks that highlight our current standing on this topic.

Tendeiro and Kiers (2019)

Our paper, as suggested by its title ‘A review of issues about null hypothesis Bayesian testing’, offers a broad overview over NHBT. We wrote this paper because, as we stated in TK, we still did not know a lot about NHBT and the Bayes factor before we wrote it, and we felt we were by far not the only ones. Furthermore, we realized that

NHBT and the Bayes factor were being increasingly used and cited in the literature, in spite of the fact that not a lot of *critical* literature existed. We could find a broad offer of tutorials and technical papers over the use and computation of the Bayes factor for an increasing variety of models. Also, the availability of free software allowing to compute Bayes factors was aiding more researchers to use this tool (e.g., JASP, JASP Team, 2020; the BayesFactor R package, Morey & Rouder, 2018). But, in our opinion, it is important to deeply understand the Bayes factor before one uses it (or needs to review a paper that uses it). This is especially the case for the researchers who are transitioning from the classical statistical paradigm of null hypothesis significance testing (NHST) and its working horse, the p value. The fact is that we could not easily find sources in published literature offering a critical appraisal of the Bayes factor. We therefore embarked on a journey which involved a lot of reading and discussing the many aspects and nuances of the Bayes factor. We reached a point where we concluded that we would like to share what we learned with the scientific community. And this is how TK came into being.

Our paper focused on *null hypothesis* Bayesian testing. The null hypothesis is thus of a type commonly used in social sciences research, for instance $\mathcal{H}_0 : \mu = 100$ or $\mathcal{H}_0 : \pi = 0$. The alternative is typically either a one- or a two-sided hypothesis, for example, $\mathcal{H}_1 : \mu \neq 0$ or $\mathcal{H}_1 : \pi > 0$. As it happens, we quickly realized that just defining the range of parameter values under $\mathcal{H}_1$ was not enough. In order for NHBT to operate, a prior distribution under $\mathcal{H}_1$ is required; we have called this the within-model (in this case, $\mathcal{H}_1$) prior. The Bayes factor contrasting $\mathcal{H}_0$ to $\mathcal{H}_1$ is given by (TK, Equation 2):

$$\underbrace{\frac{p(\mathcal{H}_1|\mathcal{D})}{p(\mathcal{H}_0|\mathcal{D})}}_{\text{posterior odds}} = \underbrace{\frac{p(\mathcal{H}_1)}{p(\mathcal{H}_0)}}_{\text{prior odds}} \times \underbrace{\frac{p(\mathcal{D}|\mathcal{H}_1)}{p(\mathcal{D}|\mathcal{H}_0)}}_{\text{Bayes factor, } BF_{10}}. \quad (1)$$

The prior odds represent the relative belief between both hypotheses before observing the data. The posterior odds represent the relative belief between both hypotheses after observing the data. The Bayes factor allows updating our relative belief between the two competing hypotheses or models, in light of the observed data. Observe that the Bayes factor is a ratio of probabilities of data conditional on a model. In other words,

the Bayes factor contrasts the ability of either model to predict the data that happened to be observed. The probabilities featuring in the formula of the Bayes factor are not to be confused with the posterior model probabilities, that is, the posterior probability of either model conditional on the observed data. The latter probabilities are to be found in the posterior odds.

Following an introductory illustration of the Bayes factor, the rest of TK consists of a discussion around 11 features concerning the Bayes factor. We referred to these features as ‘issues.’ We did not anticipate the fact that our choice of the term ‘issue’ would be the first point of contention of vRW. We were careful enough to define what we meant by this term: “An ‘issue’ can either be a limitation (according to us) or a feature that may (according to us) increase the risk of misuse or misinterpretation of a Bayes factor.” (TK, p. 775). Semantics aside, we discussed each of the following issues:

1. Bayes factors can be hard to compute.
2. Bayes factors are sensitive to within-model priors.
3. Use of ‘default’ Bayes factors.
4. Bayes factors are not posterior model probabilities.
5. Bayes factors do not imply a model is probably correct.
6. Qualitative interpretation of Bayes factors.
7. Bayes factors test model classes.
8. Mismatch between Bayes factors and parameter estimation.
9. Bayes factors favor the point null model.
10. Bayes factors favor the alternative.
11. Bayes factors often agree with p values.

van Ravenzwaaij and Wagenmakers (2021)

vRW did not consider any of the 11 issues above a problem or limitation of any kind. As per their Table 1, they deemed that issues 2-4, 7, and 9 are “pronounced advantages” and the remaining issues are neutral properties. We were quite surprised to learn that our views were so diametrically opposite, even though we were aware of these researchers’ inclination towards NHBT and the Bayes factor. It is imperative that we try to understand how such a sharp contrast of positions could come about.

Furthermore, vRW also disagreed with our preference for estimation over testing of null hypotheses. They suggest that their paper offers “a perspective that balances and complements” our own.

Below we revisit each issue one at a time. For each, we will briefly indicate the point(s) that we made in TK, the take given by vRW, and what we think about it. Furthermore, we will also discuss the philosophical divide between estimation and testing and, in this way, try to also offer a much necessary nuanced view over statistical inference over and beyond what is offered by the authors of the commentary.

The examples and figures in this paper can be reproduced by means of the accompanying R script available at the Open Science Framework (<https://osf.io/hxef2/>).

1. “Bayes factors can be hard to compute”

TK argued that Bayes factors are hard to compute. The reason is that each term in the Bayes factor formula is a marginal likelihood, that is, a weighted average of the likelihood of the observed data, with the weights given by the within-model priors (TK, Equation 4). However, TK also pointed out that nowadays there is software available that greatly curtails the problem. In short, the ‘issue’ of computational burden exists but is not unsurmountable.

In their commentary, vRW pointed out that some easy examples do exist that can be dealt with by hand. This is obviously true and we do agree with vRW, but it is the exception and not the rule. Afterwards, vRW *agreed* with our assessment of the

difficulty inherent to computing Bayes factors. Like TK, vRW further elaborated on available sources allowing to compute Bayes factors. Thus, interestingly, we do largely agree on this account.

vRW ended the discussion of this issue by arguing the following: “(…) perhaps the issue is not that NHBT is harder to *compute* than NHST, but that it is harder to *understand*.” We think this is an excellent observation. Indeed, in retrospect we wish we would have underlined the importance of *understanding* what is being computed. If today, we would have titled issue 1 as “Bayes factors can be hard to compute and understand”. vRW proceeded: “We find this argument unconvincing: there is ample evidence of the difficulty that scientific practitioners experience when trying to understand NHST”. We fail to understand the reasoning here: Why is the fact that NHST is hard to understand (a correct statement) an argument that dismisses the possibility that NHBT may also be hard to understand? Our point is that the concept embedded in the marginal likelihood under $\mathcal{H}_1$ is not straightforward and that, due to the marginal likelihood, the alternative hypothesis cannot be regarded in isolation from the prior as is often done in practice. To us, this is an issue on its own because it *may* prevent a deep understanding of the Bayes factor by some empirical researchers.

2. “Bayes factors are sensitive to within-model priors” and 3. “Use of ‘default’ Bayes factors”

The sensitivity of Bayesian analyses to prior assumptions is all well known and it is used as a common criticism against Bayesian statistics. It should therefore come as no surprise that TK considered the prior sensitivity too. In point 2, TK ended up concluding that, indeed, Bayes factors are sensitive to prior specification. They did so by recovering two examples from the literature (Liu & Aitkin, 2008; Berger & Pericchi, 2001). Furthermore, TK argued that Bayes factors are more sensitive in comparison to Bayesian estimation of posterior distributions. One way often used to mitigate the problem is to resort to default within-model priors; this was discussed in point 3 of TK. As argued in TK, default priors come at a price: They need not properly reflect the

researcher's belief, in spite of their optimal mathematical properties. Furthermore, default priors have changed through the years, thus adding a time dimension to the problem. As argued by TK, "A Bayes factor based on within-model prior distributions that poorly translate what we think about a phenomenon is misleading at best and egregiously wrong at times."

vRW disagree with us on both points. Concerning point 2, they argued that prior sensitivity is a desirable feature: Different priors instantiate different models, which lead to different predictions. Since Bayes factors are a means of comparing the predictions between two models, this makes all the sense to them. Concerning point 3, vRW tried to deemphasize the impact of the prior on the Bayes factor: "the adoption of reasonable non-default prior distributions has only a modest impact on the Bayes factor." They also argued that default priors are useful as a means to standardize analyses and facilitate comparisons and communication.

We do partly understand the arguments laid out by vRW. The Bayes factor is what it is: A ratio of two data probabilities under competing models. Nevertheless, this sensitivity is not trivial nor inconsequential, especially for newcomers to the Bayes factor (one main target of TK's paper). There is a reason why sensitivity to priors is one of the main complaints against Bayesians. Bayesians have gone to great lengths to deal with sensitivity. They advice users to run sensitivity analyses. Also JASP offers this possibility while conducting Bayesian tests. Furthermore, as pointed out by an anonymous reviewer, it is important to observe that prior sensitivity operates somehow differently between estimation and model comparison (for specific examples, see Gelman et al., 2013, pp. 183-184; Stone, 1997). Thus, sensitivity is to us a valid 'issue' in the exact sense of this term.

We are nevertheless puzzled by the holistic position of vRW on this matter. This is the main reason why here we decided to discuss issues #2 and #3 together. Specifically, we cannot understand how vRW can conciliate all of the following ideas simultaneously:

1. Bayes factors are sensitive to priors and that is a good thing, it is exactly as it should be (point 2).

2. We do acknowledge this sensitivity. Although not an issue to us, we advice running sensitivity analyses (point 3).
3. Nevertheless, rest assured that reasonable choices of priors do not matter that much (point 3).
4. We advise users to rely on default priors (point 3).

The global inferential position above supported by vRW seems illogical to us. If, according to vRW, the sensitivity to priors is an advantage, then sensitivity analyses and default priors should come as unideal inferential strategies. If priors adequately affect the predictive ability of competing models, resorting to default priors should *not* be encouraged, as it can be regarded as lowering the Bayes factor below its full potential. It seems to us that vRW want to have it both ways: Sensitivity is OK, but they nevertheless do look into its extent *and* at the same time recommend always using the same (default) prior anyway. Our position, on the other hand, is more coherent. We maintain that sensitivity to priors, albeit accurately reflective of distinctive model predictive abilities, is an issue to take into account while interpreting the Bayes factor (and actually makes its interpretation relatively difficult). By this we mean that it adds a new layer to the analysis that did not exist within the frequentist realm. We support that users should be aware of this feature, consider sensitivity analyses, and know that a compromise is being made by committing to default priors.

4. “Bayes factors are not posterior model probabilities”

This issue consists of underlining the distinction between the Bayes factor and the posterior odds (see Equation 1). The Bayes factor is a ratio of probabilities of *data* given the model, whereas the posterior odds is a ratio of probabilities of *models* given the data. Thus, the updated relative belief between two competing hypotheses is to be found in the posterior odds and not in the Bayes factor, which merely serves as the updating factor from prior to posterior odds. This may sound confusing at a first read, which does reinforce why we are trying to clarify the point. And, surely enough, we are

finding an increasing stream of papers that neglect the distinction we allude to (Wong, Kiers, & Tendeiro, 2021). Thus, it is evident to us that an issue exists. In TK we recommended researchers to use posterior odds, and not the Bayes factor, to draw conclusions about model preference.

vRW disagree with us. They favor the Bayes factor for quantifying the *evidence* in the data (i.e., the change in relative plausibility) over the posterior odds, which encapsulate the updated relative *belief* between both models.

Here we think that the difference between us is one of opinion. In *our* opinion, data are a vehicle for us to learn about the world. The vehicle is important, but the journey’s destination seems clearly more important. By focusing on the evidence (data), we avoid updating our views of the world, since the Bayes factor is the updating factor and not the updated state of knowledge. However, neither we nor vRW have evidence indicating what researchers are mostly interested in focusing on, and there is room for future research on this topic.

5. “Bayes factors do not imply a model is probably correct”

There is a divide between the Bayes factor and the truthfulness of either model under comparison; this is issue #5. Even a large Bayes factor cannot warrant that the ‘winning’ model is highly likely true: Both models may offer atrocious accounts of the real world and the Bayes factor may just be pointing at the least offending model. As a consequence, we argued that no *absolute* evidence for either model being compared is tenable.

vRW did not disagree with us, other than for the fact we termed this as an issue. Furthermore, they criticized our preference for looking at the posterior distribution under $\mathcal{H}_1$ and claimed that also such an approach suffers from the comparison curse. This is of course true, but we must clearly say that we did not argue otherwise. Unlike what vRW say, our approach is not more prone to model misspecification. In the realm of NHBT (and this is important to stress here), we lose *nothing* by ignoring the point null model and drawing inferences based on the posterior distribution under the

alternative model. We talked about this in TK and we will return to this topic at the end of this paper. vRW offered an example based on a $\text{beta}(500, 500)$ prior that attempts at shaking our position. That example is not well chosen, though, as it is based on very unrealistic priors. This was rather surprising to us, provided that earlier vRW have cautioned readers to our own choice of priors, viewing them as unrealistic, when we discussed the sensitivity of Bayes factors.

6. “Qualitative interpretation of Bayes factors”

Bayes factors are real values in the $(0, \infty)$ interval. As a continuous measure of evidence, it is not clear how to interpret the information conveyed by a Bayes factor. If, for example, $BF_{01} = 7.3$, what can we conclude other than ‘the observed data are 7.3 times more likely under $\mathcal{H}_0$ than under $\mathcal{H}_1$ ’? There is actually no easy answer to this *issue*. To circumvent the problem, some authors have developed qualitative grading scales of evidence (e.g., Jeffreys, 1961; Kass & Raftery, 1995). However, in TK we alerted that such qualitative labels are merely a useful convention.

vRW again did not disagree with us. They also think that the categories proposed for instance by Jeffreys (1961) are ‘useful as a rough heuristic.’ Perhaps the reason they do not regard this as an issue is that, as they mention, “any discretization of a statistical measure that is inherently continuous (e.g., p values, bounds on confidence intervals, Bayes factors, or indeed posterior probabilities) will be somewhat arbitrary and lead to a loss of information.” The Bayes factor is not unique in this sense so it may sound unfair to pinpoint this fact to the Bayes factor. True as this may be, the fact is still there. JASP for instance fosters the use of qualitative labels as it offers them by default. A less savvy user can easily jump from such an output into the realm of decision theory. However, very clearly, Bayes factors have no direct connection to decision theory (Schad, Nicenboim, Bürkner, Betancourt, & Vasishth, 2021). Utility functions, which should be tailored to the problem at hand, are needed for that. Hence we stand behind the point we made in TK and alert researchers to it.

7. “Bayes factors test model classes”

Two main ideas permeate this, to us, issue. Firstly, we argued that an alternative model of the type $\mathcal{H}_1 : \delta \neq 0$ actually instantiates an infinite *class* of models: Each non-zero value of δ identifies a model in the class. Secondly, we mentioned how it is entirely likely that the model class supported by the Bayes factor may not include the true model that generated the data (assuming such a true model exists). Especially the latter point should come as no surprise. The Bayes factor would have been the only existing inferential tool that makes no mistakes were this not the case. We offered a small example illustrating the situation.

vRW focused almost exclusively on our small example and did not comment on all the rest. In our example, $BF_{01} = 6.1$ although the true standardized effect was .1, thus different from 0. vRW went to great lengths to justify why the Bayes factor actually performs well in this case (this in spite of the glamorous fact that it is pointing at the *wrong model*). First vRW argued how hard it is to specify the within-model prior under $\mathcal{H}_1$: “The researcher is unable to specify the alternative hypothesis as a single point, because the true population value is usually highly uncertain a priori.¹” It seems strange to us that this argument does apply to all values, so also to .000001 for instance, but not to 0. Nevertheless, then they argued, backed up with lengthy quotes, that it is ideal to support the wrong, point null, model in cases of small sample size and small underlying effect. It must be noted, however, that the arguments offered are idealized and not formally justified.

Whether such behavior is fine or not must be left to each researcher’s judgement. vRW make a case for it, whereas we do not. The fact is: The Bayes factor can point at the wrong model (class). But an error is still an error, and idealizing an error as vRW do may not be at all helpful. Bayes factors cannot be equally good both when they are right and when they are wrong.

¹ This observation is valid and, as a side effect, helps putting point null models under a very bad light.

8. “Mismatch between Bayes factors and parameter estimation”

TK argued that testing as done through NHBT and Bayesian estimation are further detached than their frequentists counterparts. Recall that, in classical statistics, a two-sided 95% confidence interval includes the point null value if and only if the corresponding two-tailed significance test fails to reject that null hypothesis at 5% significance level. This is a widely known and used statistical property. In issue #8 we looked into a similar property in the Bayesian world, concluded that it need not necessarily hold, and presented it as an issue (meaning, as a possible source of misinterpretation). We did not cast the Bayes factor under a bad light for this, we just reported the facts.

vRW argued that the mismatch we mentioned is natural. After all, testing and estimation pertain to two different research questions and should therefore not be expected to lead to the same answer. We entirely agree, and in fact we have acknowledged this in TK (p. 784): “Both approaches seemingly answer different questions and therefore we should choose the approach that best suits the research question at hand.” But as their own quote from Berger (2006) indicates, “In classical statistics one frequently sees testing done by forming a confidence region for the parameter, and then rejecting a null value of the parameter if it does not lie in the confidence region. This is simply wrong if done in a Bayesian formulation.” Berger’s last remark is exactly our point: Newcomers to NHBT must cease resorting to this notion. This surely qualifies as an issue to us.

We also note that vRW set up one possible misinterpretation (they labeled it as ‘common’ but did not prove it to be the case). It is based on the wrong relation between having 96% posterior mass on positive effect sizes and concluding that the Bayes factor against the point null hypothesis would be $.96/.04 = 24$. This interpretation is clearly wrong, and vRW go as far as using this as a strawman to blame estimation for the problem. It detracts from the clear and crucial point that we (and Berger, 2006, among others), have made.

We are also aware of attempts to unify both testing and estimation, namely

through the so-called spike-and-slab prior (Mitchell & Beauchamp, 1988; Rouder, Haaf, & Vandekerckhove, 2018). We have actually also been entertaining this line of thought ourselves (Tendeiro & Kiers, 2021). Thus the point made by vRW complements rather than dismisses our own point, in our opinion.

9. “Bayes factors favor the point null model”

The p value is often blamed by its critics for being ‘violently biased against the null’ (Edwards, 1965). The question one should make at this point is: Violently biased in comparison to what? Berger and Sellke (1987) argued that posterior probabilities of $\mathcal{H}_0$ are typically (much) larger than p when testing point null models. TK exemplified this behavior. TK also argued that minimal levels of evidence required to reject $\mathcal{H}_0$ through NHST are matched with values of BF_{10} between 1 (barely worth mentioning) and 3 (substantial evidence), based on Jeffreys’s (1961) taxonomy. Thus, it appears as if NHST is more liberal (i.e., more easily leads to a conclusion of rejecting $\mathcal{H}_0$) than the Bayes factor or Bayesian posterior model probabilities. However, we also argued that the same result does not apply when the point model is dismissed (e.g., by comparing $\mu < 0$ and $\mu > 0$). This type of test is not what we termed NHBT, but it is enlightening to observe that the p value and the posterior null model probability are surprisingly similar in this situation; this point was also made salient in TK. TK proceeded by further arguing that the “culprit” for this state of affairs is the idiosyncratic nature of the point null model.

vRW were happy that we highlighted the bias of the p value against $\mathcal{H}_0$ in NHST, simply because that placed the Bayes factor under a seemingly favorable light. However, this is a too simplistic reading of our argument. Firstly, our argument is relative and based on comparing the p value to the Bayes factor. Taking the Bayes factor as the golden standard, then indeed p is too small too often. But we do not take the Bayes factor as the golden standard. The comparison of the relative behavior between the p value and the Bayes factor for directional hypotheses made very clear that other important factors are also at play (e.g., the role of the point null model).

Secondly, once we do not take the Bayes factor as the golden standard and allow ourselves to criticize it, we may wonder whether it is not the ‘default’ Bayes factor that is violently biased *in favor* of the null model after all. In fact, default priors in NHBT are often quite wide. Consider for example the default Cauchy(0.707) within-model prior for the JZS Bayesian t -test; see Figure 1. This prior places 50% of prior probability mass on effect sizes at least equal to 0.707 in either direction; we find this unrealistic in social science settings. This prior breadth implies allocating prior mass on unreasonably large effect sizes. As a consequence, the marginal likelihood for the alternative model will be severely penalized, which artificially increases the relative plausibility of the null model as the data generating mechanism. This is particularly true for small effect sizes and sample sizes (and on occasion also for large samples; Stone, 1997), thus, in many common scenarios in psychological research. In fact, this is worrisome for cases where researchers use Bayes factors to draw support *for* the null model (when for instance a frequentist significance test gave a non-significant result), for which the possibility of small effects is rather large.

We will make our point explicit by means of an example; see Figure 2. We considered the one-sample JZS Bayesian t -test for samples of varying size ($n = 10, 50, 100, 250$, and 500) and standardized effect sizes $\delta = 0(.01).3$. We entertained the Cauchy prior for δ (see Figure 1), similarly to what is done in the ‘BayesFactor’ (Morey & Rouder, 2018) and JASP (JASP Team, 2020) software packages. Four scale values for the Cauchy prior were used. The values 0.707, 1, and 1.414 correspond to what BayesFactor labels as ‘medium’, ‘wide’, and ‘ultrawide’ priors. We further considered $r = 0.5$ as an example of a more reasonable account of a priori uncertainty on commonly found effect sizes in psychology. For each combination of n , δ , and prior, we computed BF_{01} and $\mathcal{P} = P(\delta \in [-0.1, 0.1] | \text{data})$. These two quantities are not equivalent. For example, BF_{01} says nothing about the plausibility of negligibly small values of δ whereas $\mathcal{P}$ does. Also, $\mathcal{P}$ offers a direct stochastic statement about the value of δ , whereas BF_{01} is really an indirect statement about δ through the relative likelihood of the *data* observed under the two models. Nevertheless, we will use both these measures

to infer about the underlying parameter δ . For large effect sizes we expect that both BF_{01} and $\mathcal{P}$ are close to 0. As the effect size decreases towards 0 we expect that BF_{01} increases and $\mathcal{P}$ approaches 1. Each curve in a panel corresponds to a particular combination of n and δ , for effect sizes ranging from 0 (upper-right part of the curve) through 0.3 (lower-left part of the curve). For reference, we used $BF_{01} = 10$ as the reference value to draw relative support for $\mathcal{H}_0$. This level of evidence was dubbed ‘strong’ by Jeffreys (1961). We used the interval $[-0.1, 0.1]$ as a proxy for effect sizes δ that are practically equivalent to 0, in the same sense as used in equivalence testing.

Several important notes that can be derived from Figure 2 are in order:

1. As r decreases, it is harder to draw support for $\mathcal{H}_0$ through the Bayes factor. In other words, as one moves from panels D through A, the plotted lines gradually ‘move’ to the left. This is hardly new and was also highlighted in TK.
2. $BF_{01} = 10$ pairs up with relatively high posterior probabilities $\mathcal{P}$ when $r = .5$, for large sample sizes only ($\mathcal{P} = .89$ for $\delta = 0$ and $n = 250$; $\mathcal{P} = .94$ for $\delta = 0.03$ and $n = 500$). Thus, for the Cauchy(0.5) prior, which is narrower than the advised default, we could say that BF_{01} and $\mathcal{P}$ point in the same direction.
3. As r increases, the value of $\mathcal{P}$ associated to $BF_{01} = 10$ decreases a lot. This can be appreciated by focusing on the rectangles capturing the intersection of the curves with the vertical line. These rectangles ‘move’ down the vertical line as one moves from panels A through D. That is, taking $\mathcal{P}$, a direct stochastic statement of the current knowledge about δ , as the standard, we can conclude that the Bayes factor *is biased in favor of $\mathcal{H}_0$, the more so as r increases*.

The observations above are not inconsequential. For example, for the default Cauchy(0.707), it happens that $\mathcal{P} = .78$ and $BF_{01} = 10.4$ when $n = 250$ and $\delta = 0.05$. Thus, ‘strong’ evidence in favor of $\mathcal{H}_0$ that $\delta = 0$ (as suggested by the Bayes factor) is in correspondence to a posterior probability of .22 of δ being at least 0.1 in either direction. This discrepancy is worrisome. For the ‘ultrawide’ prior (panel D) it is possible to reach ‘strong’ levels of evidence for $\mathcal{H}_0$ through BF_{01} for samples of size 50

such that $\mathcal{P}$ is smaller than .50! What this shows is that both the p value and the Bayes factor suffer from a similar curse: Large values of BF_{01} , just like large p values, do not necessarily imply that there is evidence for ‘no effect’, or even that the effect is small. In other words, the often used moniker saying that “Bayes factor allow drawing support for the null model, unlike the p value”, lacks qualification. All characteristics of the model being compared – sample sizes, priors, likelihoods – must be carefully considered. We conclude that our issue # 9, which is that ‘Bayes factors favor the point null model’, is real and we urge practitioners to learn about it and to deal with it with great care.

vRW very clearly refuse *any* criticism towards the Bayes factor: “it is important that the Bayes factor is a unique and deterministic consequence of the prior distribution and the data; complaints about the behavior of the Bayes factor are therefore indirectly complaints about the prior distribution or the data.” We find this argument limitative. Deferring criticism towards the Bayes factor to either the model, prior, or data is not constructive. Following this escape route, there is also nothing wrong with the p value: The problem must lie either on the data, the model, or the test statistic. We think there is nothing to be learned by objectively refusing to understand the inner workings of a statistical tool just because the Bayes’ rule applies.

Our next disagreement is based on the value, or lack thereof, of the point null model. In TK we have argued why, according to us, too much credence is given to point null models in psychological research. vRW base their argument on a quote from Jeffreys: “variation must be taken as random until there is positive evidence to the contrary” (Jeffreys, 1939, p. 345). It is unclear to us why variation *must* be taken as random from the onset. This is too rigid a stance to our liking. There is no reason not to use the alternative model and probabilities derived from it to quantify our uncertainty about (and around) the point null. Nothing, theoretically or logically, forbids this and we are unconvinced that one must invariably follow a test that first rejects the null model before we can think of estimating our uncertainty on the true effect. vRW further argued that failure to replicate several effects in psychology seems to demonstrate that actually many parameters are zero. We consider this as invalid

proof for the existence of true zero parameter values, because they just refer to test outcomes. Hence their argument falls short and is not a reason to support the widespread use of point null models. For clarity, we do think point null models may be of value in particular instances where strong theory exists (e.g., when studying ESP). But in general, the theoretical common concept of absence of effects seems ill-devised to us. An effect's existence, magnitude, and even direction depends on a multitude of factors (population, lab, experimenters, protocols, etc.). This perennial association of point null hypotheses to actual absence perpetuates an infortunate illusion of the ability to establish invariance or consistency that is yet to be proven to generally hold in the social sciences.

10. “Bayes factors favor the alternative”

Like the p value, the Bayes factor favors the alternative model asymptotically even for arbitrarily small (but non-zero) effect sizes. This consistency property seems ideal, but one must then be reminded what it implies from a practical perspective. In practice, a rejection of $\mathcal{H}_0$ may be associated to either large or very small effect sizes. Thus it is important to not rely only on Bayes factor when one wants to quantify an effect size. Other quantities, based on posterior distributions quantifying uncertainty, are not only of value but actually essential to disambiguate between large and relevant versus small and unimportant. Furthermore, TK showed that Bayes factors accumulate evidence in favor of true alternative models much faster than they do in favor of true null models as the sample size increases. Thus, Bayes factors can indeed draw support for either model under comparison, but they do so asymmetrically and favoring the alternative model.

Concerning the first point above (consistency), vRW agreed with us that reporting effect sizes is crucial. However, they immediately imposed an analytical strategy that we do not condone and that does not follow from our argumentation. According to vRW, one “start[s] the statistical investigation with a test in order to determine whether the null hypothesis could be discarded, meaning some effect exists that is worthy of estimation (Haaf, Ly, & Wagenmakers, 2019). Only if the data provided clear

evidence against the null hypothesis would Jeffreys turn to estimation, which he then based on the alternative hypothesis.” We actually see no valid reason to test. Establishing presence by means of Bayes factors is wishful thinking. As argued before, we think effects may exist in many forms, depending on context. Furthermore, testing does not equate to establishing presence, but merely accruing evidence for existence. In other words, testing is a much more modest endeavor than vRW make believe. vRW mistakenly believe that Bayes factors allow establishing presence of effects and this should take precedence in the inference ladder: “[Estimation] leaves unaddressed the question of whether the effect exists in the first place; the statement may tempt the reader to conclude that the data support the hypothesis that the effect is present.” In fact, estimation more directly allows quantifying our uncertainty on the said effect of interest. An anonymous reviewer offered an even sharper argument downplaying the use of testing: “This is begging the question. If the researcher isn’t interested in a decision rule, but in estimation, then there is no null hypothesis. In any case, estimation assumes only that the parameter is unknown – nothing more.”

Concerning the second point (asymmetry of the Bayes factor), vRW surprisingly state that our point does not stand (“The contrast that TK allude to does not exist.”). Upon further reading, though, it is clear that vRW do not negate our statement (as it is factually true), but just choose not to give it any weight. Instead, as done to all previous issues, vRW offer a defense to the asymmetry property. Our opinion is that users should at least know that this property holds and we think it is pertinent to those considering adding Bayes factors to their toolbox. The common saying used by proponents of Bayes factors which states that Bayes factors, unlike p values, allow drawing support for either model being compared, lacks qualification to be accurate. We offered that insight in TK.

11. “Bayes factors often agree with p values”

Here, TK explored an interesting property that aligns conclusions derived from p values and Bayes factors. Following Trafimow (2003), it follows that congruence

between p values and Bayes factors occurs when the observed data are more likely under $\mathcal{H}_1$ than under $\mathcal{H}_0$. In this case, the posterior probability of $\mathcal{H}_0$ will be low. Thus, a situation in which data are less likely under $\mathcal{H}_0$ (i.e., low p value) and more likely under $\mathcal{H}_1$ (i.e., low posterior probability of $\mathcal{H}_0$) leads to congruence between the p value and the Bayes factor.

vRW claim that our remark is trivially true, but their argument implies that our point only holds for obvious, large, effects. As demonstrated above, this is not the case. Data need ‘only’ be more consistent under the alternative model than under the null model. This hardly qualifies as an example of Berkson’s interocular traumatic test (implying that for large enough effect sizes about any inferential method would work well), as suggested by vRW.

vRW also claimed that our remark is trivially false, and they owe that to the purported dichotomous nature of the p value (statistically significant or non-significant) contrasting to the trichotomized nature of the Bayes factor (favor $\mathcal{H}_0$, favor $\mathcal{H}_1$, or suspend judgement). We do not see the value in this argument, as it does not address our point. We did not claim that p values and Bayes factors are interchangeable. Instead, we argued that there are conditions under which conclusions based on either p values or Bayes factors will coincide. For the two t -test results offered by vRW, data are not largely consistent with $\mathcal{H}_1$, thus they do not counter our argument in any way (nor would that be possible, as our argument is mathematically based).

Discussion

vRW lined up two reasons that, in their understanding, are at the core of our disagreement. The first reason is that we are of the opinion that researchers should learn the basics of NHBT before attempting to use it. vRW, on the other hand, are happy to initiate a new mindless ritual. In their own words: “We believe that NHBT, even if executed as a thoughtless ritual, still markedly improves on the status quo.” We cannot fathom how this is possible nor how this could ever be sensibly advised. There are real perils inherent to NHBT and the Bayes factor, as our paper clearly shows.

Pretending that no issue exists and leading researchers to blindly adopting this tool makes us seriously worried. vRW argue that, if a thorough understanding of a statistical tool were a requirement, then “the use of p values and confidence intervals would be all but forbidden.” This argument should actually be flipped around: Precisely *because* the misunderstanding of p values and confidence intervals has led to so many serious problems, we should now use this new opportunity to do things better and *not* promote mindless Bayes factoring through an analysis. It is much more constructive to learn from our errors in the past and do better in the future than to borrow from those mistakes in the past and use them as a justification to repeat them in the future.

vRW also fear that detractors of Bayes factors will use TK as an escape goat: “we worry that in practice, researchers will refer to the TK paper mainly as an excuse for avoiding NHBT altogether and not learning anything new at all.” Obviously we clearly disagree with this take. Firstly, we do not have an agenda based on promoting an indiscriminate use of the Bayes factor. We worry little in case researchers do not adopt the Bayes factor, especially if they do not understand it sufficiently. We would like, however, that if researchers want to use the Bayes factor, they do it correctly, that is our agenda. Furthermore, we think that it is a deeper knowledge of the Bayes factor that may attract more people to this framework. In fact, focusing only on all the purported advantages without discussing rigorously any disadvantage (or issue) may sound suspicious to a newby.

The second reason of disagreement between us, according to vRW, is the fact that we claim Bayesian estimation under $\mathcal{H}_1$ to be preferred over NHBT. Clearly, we do think estimation has the upper hand. vRW suggested we defend Bayesian estimation as an ‘alternative’ to NHBT. We would rather say that the Bayes factor should not be the start nor the end of it all. If used (e.g., because the point null model is of particular interest for the research question at hand), we have no problem in looking into the relative predictive ability of two competing models (i.e., the Bayes factor). But we would never advise leaving the analysis at that. For example, unlike vRW, we do not condition fitting the alternative model upon the result of a test of ‘existence’ of the

effect at hand. Unlike what vRW repeatedly suggest in their writings, Bayes factors cannot be used to *establish* the existence of an effect. Bayes factors help informing how likely the observed data are under either model. The reasoning here is stochastic, not factual. We would be at a quandary if such a requirement needed to be met in order for us to be able to fit a model allowing for a non-null parameter value. Given the small sample sizes and miniscule effects that permeate psychological research, we would too often be left with our hands tied on inconclusive, or leaning towards the null, types of Bayes factors. Not looking at estimated effect sizes and expressing all we know about them based on an unfavourable Bayes factor would be a missed opportunity to learn more.

vRW are inconsistent about the truthfulness of the null model, depending on the argument they intend to make. For example, they made the following statement that acknowledges the falsity of all statistical models (thus, including the point null model), while actually arguing *in favor* of the value of point null hypothesis: “(. . .) it should be stressed that the Bayes factor is based on a comparison of predictive performance that is independent of the notion of absolute or relative model truth (Wagenmakers, Grünwald, & Steyvers, 2006). The idea that we should not use models that we know to be false can easily result in an inferential impasse, because all statistical models are ultimately false.” But then, to explain that one must first establish existence before estimation, they say this: “[estimation] assumes the falsity of the null hypothesis, which is often the very target of inference.” vRW cannot have it both ways. Furthermore, vRW blame us for not following Jeffreys’ primer, who advocated the ‘test first then, who knows, estimate’ route. We see no reason why we need to follow, or undercut, Jeffreys. Our reasons have been clearly outlined in TK, and we see no problem if they are not perfectly aligned with previous thinkers. We note that, unlike what vRW or Jeffreys says, not everyone thinks that “until the Bayes factor favors the alternative hypothesis, ‘the simpler hypothesis holds the field.’” In fact, there is no good reason to always favor the simpler model, especially if it erroneously leads people to associate it

to an impossible ‘absence’ desideratum. See for example Neal (1996, pp. 103-104)²:

“Sometimes a simple model may outperform a more complex model (...). Nevertheless, I believe that deliberately limiting the complexity of the model is not fruitful when the problem is evidently complex. Instead, if a simple model is found that outperforms some particular complex model, the appropriate response is to define a different complex model that captures whatever aspect of the problem led to the simple model performing well.” The position defended by this quote is misaligned with that of Jeffreys. Not a problem in our eyes, just a reflection of different views of science.

vRW offer a model selection example based on polynomial regression that is meant to undercut our preference for estimation. Unfortunately this example missed the point. vRW seem to think that we blindly adhere to an ‘estimate-only’ policy; we do not. Instead, we use judgement and cater our efforts to the situation at hand. We do include model comparison in our own toolbox, be it through Bayes factors, information criteria, or cross-validation (e.g., Bayesian leave-one-out; Vehtari, Gelman, & Gabry, 2017). What we do not encourage is blind adherence to a procedure that forces researchers to always run tests as a prerequisite.

vRW also seem to impose what question researchers are looking forward to answer, and that only Bayes factors can answer it. In their own words, the question is “to what extent do the data support the hypothesis that there is an effect?” In fact, vRW offer no empirical support defending that this is the core question that researchers seek answering. Furthermore, it is a fallacy to assume that *only* Bayes factors are correctly equipped for the job. To put it bluntly: *The Bayes factor is nothing but a number*³. Say that $BF_{10} = 10.2$. In order to answer the question above one must say that ‘the data are 10.2 times more likely were $\mathcal{H}_1$ true, compared to the null model.’ And now, what? What did we really learn about the effect? Is it large or small?; we do

² This quote was originally found at

https://statmodeling.stat.columbia.edu/2004/12/10/against_parsimo/

³ “We need to abandon the idea that there is a device that will produce a single-number decision rule” (Vehtari, Simpson, Yao, & Gelman, 2019).

not know because the Bayes factor is oblivious to that. Is it positive or negative?; we also do not know because the Bayes factor is oblivious to that (at least for two-sided alternatives). We just asked two questions (magnitude and direction) that *we* trust many researchers would surely like to answer. The question alluded to by vRW is a question of *existence*. Albeit of interest, it surely makes more sense in fields where precision is key (say, to establish the existence of Higgs boson). In psychology, with all the variants and conditionals affecting cognition and behavior, hoping to use *one number* to establish the existence of an effect may be dubbed the psychologist's fool's gold. We caution researchers to not chase fool's gold. Better embrace what one can really establish, acknowledge limitations and uncertainty, and let time take its course for evidence to accumulate and knowledge to prosper.

vRW do not explain in detail how testing through Bayes factors can be used to first establish existence (before proceeding to estimation). How is one supposed to do this? For example, above what threshold is one supposed to believe that $\mathcal{H}_0$ is incorrect? Conversely, above what threshold is one supposed to believe that $\mathcal{H}_0$ is correct? These dichotomies are unhelpful and, in reality, what vRW suggest is simply statistically impossible. We repeat: *Bayes factors cannot be used to establish the existence, or presence, of effects in psychology*. Neither can *p*-values, for that matter. We must be much more modest with respect to what we can achieve. Through Bayes factors, what we can is collect relative evidence either way. This incorrect interpretation of what can be done through the Bayes factor is what we now may refer to as issue #12 in our cautionary list of Bayes factors properties.

Conclusion

We worry that Bayes factors may start being blindly used in research. There is preliminary evidence suggesting that this is already occurring (Wong et al., 2021). Our original paper is a contribution to mitigate and prevent misconceptions. vRW fear that our contribution may be misconstrued. We believe its benefits outweigh this risk. In any case, we believe that it is best to not use the Bayes factor than to use it wrongly.

After all, we already have enough problems as it is with NHST. We wish new adopters of NHBT are aware of all the advantages, and issues, that come with it. Full awareness is the best way forward.

We do not wish to imply that Bayes factors should never be used. There may be settings for which a point null model is of value and researchers legitimately would like to compare the predictive ability of that model against other more general account of the real world. If that is the case, then by all means do use the Bayes factor. However, what we clearly want to state is that Bayes factors *must* warrant much more consideration and introspection from researchers *before* using them. We therefore frontally disagree with vRW on this account.

We think that perspectives such as those presented in TK and in this article are important, in order to provide the necessary nuance and context to the literature. We think our views contrast and balance those of vRW. Hopefully readers of these articles will be able to gather information, reason, and make an informed judgement of the merits of NHBT before they start using it. Consideration of the many issues raised in TK and vRW need not necessarily lead to clearcut black-and-white decisions. There are many shades of gray to the statistical enterprise, from data preprocessing through drawing inferential statements (e.g., Bastiaansen et al., 2020; Silberzahn et al., 2018). The Bayes factor is not superior to all other inferential strategies available across the board. We call for wise judgement also when using Bayes factors in psychological research.

References

- Bastiaansen, J. A., Kunkels, Y. K., Blaauw, F. J., Boker, S. M., Ceulemans, E., Chen, M., ... Bringmann, L. F. (2020). Time to get personal? The impact of researchers choices on the selection of treatment targets using the experience sampling methodology. *Journal of Psychosomatic Research*, *137*, 110211. doi: 10.1016/j.jpsychores.2020.110211
- Berger, J. O. (2006). Bayes factors. In S. Kotz, N. Balakrishnan, C. Read, B. Vidakovic, & N. L. Johnson (Eds.), *Encyclopedia of statistical sciences* (Second ed., Vol. 1, pp. 378–386). Hoboken, NJ: Wiley.
- Berger, J. O., & Pericchi, L. R. (2001). Objective Bayesian methods for model selection: Introduction and comparison. In *Model Selection* (pp. 135–207). Beachwood: Institute of Mathematical Statistics Lecture Notes - Monograph Series.
- Berger, J. O., & Sellke, T. (1987). Testing a point null hypothesis: The irreconcilability of p values and evidence. *Journal of the American Statistical Association*, *82*, 112–122.
- Edwards, W. (1965). Tactical note on the relation between scientific and statistical hypotheses. *Psychological Bulletin*, *63*(6), 400–402. doi: 10.1037/h0021967
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (Third edition ed.). Boca Raton: CRC Press.
- Haaf, J. M., Ly, A., & Wagenmakers, E.-J. (2019). Retire significance, but still test hypotheses. *Nature*, *567*(7749), 461–461. doi: 10.1038/d41586-019-00972-7
- JASP Team. (2020). *JASP (Version 0.14.1)[Computer software]*.
- Jeffreys, H. (1939). *Theory of probability* (First ed.). Oxford: The Clarendon Press.
- Jeffreys, H. (1961). *Theory of probability* (Third ed.). Oxford: Oxford University Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*, 773–795. doi: 10.2307/2291091
- Liu, C. C., & Aitkin, M. (2008). Bayes factors: Prior sensitivity and model generalizability. *Journal of Mathematical Psychology*, *52*, 362–375.
- Mitchell, T., & Beauchamp, J. (1988). Bayesian Variable Selection in

- Linear-Regression. *Journal of the American Statistical Association*, 83(404), 1023–1032. doi: 10.2307/2290129
- Morey, R. D., & Rouder, J. N. (2018). *BayesFactor: Computation of Bayes Factors for Common Designs*.
- Neal, R. M. (1996). *Bayesian Learning for Neural Networks* (Vol. 118; P. Bickel et al., Eds.). New York, NY: Springer New York. doi: 10.1007/978-1-4612-0745-0
- Rouder, J. N., Haaf, J. M., & Vandekerckhove, J. (2018). Bayesian inference for psychology, part IV: Parameter estimation and Bayes factors. *Psychonomic Bulletin & Review*, 25, 102–113.
- Schad, D. J., Nicenboim, B., Bürkner, P.-C., Betancourt, M., & Vasishth, S. (2021). *Workflow techniques for the robust use of bayes factors*.
- Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., . . . Nosek, B. A. (2018). Many Analysts, One Data Set: Making Transparent How Variations in Analytic Choices Affect Results. *Advances in Methods and Practices in Psychological Science*, 1(3), 337–356. doi: 10.1177/2515245917747646
- Stone, M. (1997). Discussion of papers by Dempster and Aitkin. *Statistics and Computing*, 7(4), 263–264. doi: 10.1023/A:1018502622516
- Tendeiro, J. N., & Kiers, H. (2021). *With bayesian estimation one can get all that bayes factors offer, and more*. PsyArXiv. doi: 10.31234/osf.io/zbpmv
- Tendeiro, J. N., & Kiers, H. A. L. (2019). A review of issues about null hypothesis Bayesian testing. *Psychological Methods*, 24(6), 774–795. doi: 10.1037/met0000221
- Tendeiro, J. N., & Kiers, H. A. L. (2022). On the white, the black, and the many shades of gray in between: Our reply to van Ravenzwaaij and Wagenmakers (2021). *Preprint*. doi: 10.31234/osf.io/tjxvz
- Trafimow, D. (2003). Hypothesis testing and theory evaluation at the boundaries: Surprising insights from Bayes’s theorem. *Psychological Review*, 110(3), 526–535. doi: 10.1037/0033-295X.110.3.526
- van Ravenzwaaij, D., & Wagenmakers, E.-J. (2021). *Advantages Masquerading as*

- 'Issues' in Bayesian Hypothesis Testing: A Commentary on Tendeiro and Kiers* (2019) (Preprint). PsyArXiv. doi: 10.31234/osf.io/nf7rp
- Vehtari, A., Gelman, A., & Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5), 1413–1432. doi: 10.1007/s11222-016-9696-4
- Vehtari, A., Simpson, D. P., Yao, Y., & Gelman, A. (2019). Limitations of “Limitations of Bayesian Leave-one-out Cross-Validation for Model Selection”. *Computational Brain & Behavior*, 2(1), 22–27. doi: 10.1007/s42113-018-0020-6
- Wagenmakers, E.-J., Grünwald, P., & Steyvers, M. (2006). Accumulative prediction error and the selection of time series models. *Journal of Mathematical Psychology*, 50(2), 149–166. doi: 10.1016/j.jmp.2006.01.004
- Wong, T. K., Kiers, H., & Tendeiro, J. (2021). *On the potential mismatch between the function of the bayes factor and researchers' expectations*. PsyArXiv. doi: 10.31234/osf.io/86p4k

Figure 1. Four Cauchy priors with varying scale parameter r . The central 50% probability of each prior (i.e., the area between $(-r, r)$ under the prior) is shaded.

Figure 2. Relation between BF_{01} and $\mathcal{P} = P(\delta \in [-0.1, 0.1] | \text{data})$, for the one-sample JZS Bayesian t -test. Each panel pertains to a prior from Figure 1. The relative evidence in favor of $\mathcal{H}_0$ provided by BF_{01} increases with r . Wider priors overstate the evidence in favor of the null in comparison to $\mathcal{P}$, thus rendering the Bayes factor biased in favor of the null. See text for details.