With Bayesian Estimation One Can Get All That Bayes Factors Offer, and More

Jorge N. Tendeiro*

Hiroshima University

Henk A. L. Kiers*

University of Groningen

*Both authors contributed equally to this paper.

Abstract

In classical statistics, there is a close link between Null Hypothesis Significance Testing (NHST) and parameter estimation via confidence intervals. However, for the Bayesian counterpart, a link between Null Hypothesis Bayesian Testing (NHBT) and Bayesian estimation via a posterior distribution is less straightforward, but does exist, and has recently been reiterated by Rouder, Haaf, and Vandekerckhove (2018). It hinges on a combination of a point mass probability and a probability density function as prior (denoted as the spike-and-slab prior). In the present paper it is first carefully explained how the spike-and-slab prior is defined, and how results can be derived for which proofs were not given in Rouder et al. (2018). Next, it is shown that this spike-and-slab prior can be approximated by a pure probability density function with a rectangular peak around the center towering highly above the remainder of the density function. Finally, we will indicate how this ‘hill-and-chimney’ prior may in turn be approximated by fully continuous priors. In this way it is shown that NHBT results can be approximated well by results from estimation using a strongly peaked prior, and it is noted that the estimation itself offers more than merely the posterior odds on which NHBT is based. Thus, it complies with the strong APA requirement of not just mentioning testing results but also offering effect size information. It also offers a transparent perspective on the NHBT approach employing a prior with a strong peak around the chosen point null hypothesis value.

With Bayesian Estimation One Can Get All That Bayes Factors Offer, and More

In classical statistics, there is a close link between Null Hypothesis Significance Testing (NHST) and parameter estimation via confidence intervals. However, for the Bayesian counterpart, a link between Null Hypothesis Bayesian Testing (NHBT) and Bayesian estimation via a posterior distribution is less straightforward, but does exist. This link between these two very important inferential strategies is actually of great relevance. The seemingly dichotomous divide between the adequacy of conducting a hypothesis test versus estimation often leads to polarized discussions and advice on how to choose between the two inferential frameworks (e.g., Kruschke & Liddell, 2018b; Uygun Tunç, Tunç, & Lakens, 2021; Wagenmakers et al., 2018).

Both Bayesian estimation and Bayesian hypothesis testing predicate on the uncontroversial Bayes' theorem (e.g., Puga, Krzywinski, & Altman, 2015). On its own, the Bayes theorem offers a mathematical means of combining information from two distinct sources of information. *Prior distributions* may be used to encapsulate our current understanding of the phenomenon being studied, or to represent the predictions of the theory that we intend to test. The *likelihood function* captures the information provided by the data through the statistical model of choice. The combined information is summarized by the *posterior distributions*, which offer an updated account of our state of knowledge. We refer readers wishing to get started on basic Bayesian inference to a variety of introductory books and papers (e.g., Etz & Vandekerckhove, 2018; Gelman et al., 2013; Kruschke, 2015; van de Schoot et al., 2021).

Inference through Bayesian *estimation* consists of carefully describing the information shown in the posterior distribution, for instance, via *credible intervals* that attach probabilities to specific ranges of values for the parameter at hand. The Bayesian estimation approach is a widely recommended procedure for data analysis (e.g., Gelman et al., 2013; Kruschke, 2011, 2013; Kruschke & Liddell, 2018a, p. 170; Stern, 2016, p. 27; van der Linden & Chryst, 2017; van de Schoot et al., 2014). Bayesian hypothesis *testing*, on the other hand, compares the predictive ability of two competing models or hypotheses. A model that outperforms its competitor in

predicting the observed data leads to a positively revised view of its ability to describe the phenomenon being studied. This update of the relative belief between the two models due to the observed data is processed, precisely, through the Bayes theorem.

The Bayesian testing approaches we refer to in this paper are based on comparing a point null hypothesis to an alternative hypothesis that specifies a parameter to be distributed according to a particular density function. We abbreviate this procedure as NHBT (Null Hypothesis Bayesian Testing). The result is expressed in a so-called Bayes Factor, which captures the evidence from the data for the point null hypothesis versus the specific alternative hypothesis chosen. This approach has been summarized by Kruschke and Liddell (2018a), who also issue five warnings as to its use (see also Tendeiro & Kiers, 2019, 2022; van Ravenzwaaij & Wagenmakers, 2021). A lot of attention nowadays is given to promoting and further developing Bayesian ways of testing null hypotheses. It is sometimes even suggested that testing should routinely precede estimation of effect sizes and their uncertainty in the inferential ladder (e.g., Jeffreys, 1961; van Ravenzwaaij & Wagenmakers, 2021; Wagenmakers et al., 2018).

It is not immediately clear how Bayesian estimation and testing, as described above, relate to each other. In fact, inferences derived from Bayesian estimation can easily lead to results that may appear to conflict with those from NHBT. For instance, when an NHBT is carried out by the programme JASP (JASP Team, 2020), it also offers a posterior distribution and a Highest Density Interval (HDI, also known as credible interval) based on that. An NHBT may indicate fairly strong evidence in favor of 0 while the HDI may not contain the value 0 (Kruschke & Liddell, 2018b; Tendeiro & Kiers, 2019). It is often argued that the researcher's goals are what dictate the preference for either estimation or testing. This, however, has not lessened the discussion among campers on both ends of the spectrum.

One of the goals of this paper is to contribute towards a unification between estimation and testing in the Bayesian paradigm. We aim to do so in a manner such that it is accessible for the general reader who has only little knowledge of probability theory (but knows basic expressions). Such readers may skip the more technical parts,

and notably the appendices; the latter clearly are for the experts. Furthermore, we are not the first ones to have attempted this unification. Of particular importance is the work of Rouder et al. (2018) and Liao, Midya, and Berg (2020). Rouder et al. (2018) drew attention to an earlier established link between Bayesian estimation and NHBT through the so-called spike-and-slab prior ‘density’ (Mitchell & Beauchamp, 1988). See also Kruschke (2018); a less detailed but related account on this is also given by Williams, Bååth, and Philipp (2017). Here we explore the role of the spike-and-slab prior in Bayesian inference, and we reflect on how various *approximations* to the spike-and-slab prior allow a deeper understanding of the link between testing and estimation in Bayesian statistics. Such priors are by no means new. They feature strongly in the literature on selection of variables in regression. See O’Hara and Sillanpää (2009) for an excellent overview on this topic, and also George and McCulloch (1993, 1997), Ishwaran and Rao (2003, 2005), Kuo and Mallick (1998), Malsiner-Walli and Wagner (2011), Ntzoufras, Forster, and Dellaportas (2000), and Wagner and Duller (2012). Morey and Rouder (2011) mentioned such models in the simpler context of testing or estimating single parameters. Kruschke (2018, supplement) offers a more detailed analysis of such variants to the spike-and-slab prior. Thus, for the simple single parameter case, the literature has already surveyed various aspects of the concrete relation between estimation and Bayes factor based testing.

Our contribution

This paper’s offering can be summarized in five main points. First, this paper further contributes towards an integrated view of estimation and testing, under the Bayesian paradigm. Second, we will identify the spike-and-slab prior as an extreme choice of a peaked prior, within a broad class of more or less peaked priors. With this purpose in mind, we will introduce what we dubbed the *hill-and-chimney* prior. Third, we will identify the Bayes factor and notably its related posterior odds, as only one of many interesting concrete measures to evaluate probabilities of events on the basis of a posterior distribution. Fourth, we will argue that the hill-and-chimney’s discontinuous

nature is not ideal and then suggest continuous approximations. We discuss two relatively simple alternatives to this family of approximating continuous distributions. And fifth, we briefly discuss how our integrated inferential approach can be contextualized in the realm of decision theory, following Kruschke and Liddell (2018b). We end by summarizing our main ideas and discuss future avenues for research in the discussion.

All examples and figures can be reproduced by means of the accompanying R script available at the Open Science Framework (<https://osf.io/6rt9m/>). The repository also includes Stan and R scripts that readily allow fitting the estimation models advocated in this paper.

Bayesian estimation

Basics

As in Rouder et al. (2018), consider data $\mathbf{y} = \{y_1, \dots, y_n\}$ of size n , assumed to be randomly sampled from a normal population with unknown mean μ and known variance σ^2 . Although quite unrealistic, the known variance assumption is used only because it facilitates the mathematical formulation and it does not bear consequences conceptually. Let δ denote a standardized effect size measure, defined as μ/σ , that is, the population mean divided by the population standard deviation. Our goal is to infer what can be said about the population effect size δ . For the Bayesian machinery to operate we must choose a prior for δ . Bayesian inference proceeds by combining the information in the prior with that from the data. The posterior distribution for δ will reflect our updated belief in δ in light of the observed data. As a simple example, consider Figure 1. We assume that, for the population of interest, the variance is known and equal to 1. Our goal is to infer about the real value of δ . A prior normal distribution for δ with mean 0 and standard deviation .3 was chosen. Following Bayes' theorem, this prior was multiplied by the normal likelihood based on a sample of size 20 and mean 0.5. The result— the posterior distribution— is the updated compromise for δ combining both the prior and the likelihood.

Of particular importance is the choice of the prior distribution. There are in fact many possible priors that we can choose from, depending on our purposes. What we propose to do in this paper—building a conceptual bridge between Bayesian estimation and testing—is based on using a special type of prior distribution that is not of common use within the social sciences: The so-called *spike-and-slab* prior.

The spike-and-slab prior

We consider the spike-and-slab prior following Rouder et al. (2018). For the idea behind this they refer to Jeffreys (1939), while for the particular name of the model they refer to Mitchell and Beauchamp (1988). The name of this particular prior becomes obvious when looking at its plot: See Figure 2A, which is a copy of Figure 4A in Rouder et al. (2018). The spike-and-slab prior is a weighted combination of a probability density for all non-zero δ values and a probability mass at $\delta = 0$ (in Figure 2A, both parts have equal weight). The probability density is shown as the gentle symmetric curve in the bottom of the picture (which can be seen as a slab of sorts), while the arrow peaks above this curve (clearly a spike) and has half the length of the maximum indicated by the dashed horizontal bar. Note that the slab density and the spike probability are in different units, so the y -axis actually refers to two different scales. To understand what such a combination of a density and a probability mass function means, one could, for instance, assess that $P(\delta < 0) = .25$, and likewise $P(\delta > 0) = .25$. So we see that the three mutually exclusive events, $\delta < 0$, $\delta = 0$, and $\delta > 0$, covering the whole possible range of outcomes, have probabilities .25, .50, and .25, respectively. Even though Figure 2A does not give scale values for the density, for the above example it is clear that the density should be such that the total area under the curve equals .5 (i.e., 1 minus the probability of the spike). For varying weights, the size of the height of the spike varies, and the height of the slab varies with it, ensuring that the total area under the curve always equals one minus the probability mass at 0.

Mathematically, the spike-and-slab prior can be given as follows¹:

$$\begin{cases} p(\delta) &= (1 - \rho_0) \times \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left(-\frac{\delta^2}{2\sigma_0^2}\right) &, \text{ if } \delta \neq 0 \\ P(\delta = 0) &= \rho_0 &, \text{ if } \delta = 0 \end{cases}. \quad (1)$$

Please note that $p(\delta)$ denotes a (non-normalized) density function which is proportional to the $\mathcal{N}(0, \sigma_0^2)$ distribution. The variance σ_0^2 controls the spread of the slab part of the prior. $P(\delta = 0)$, on the other hand, is a probability, hence the density function values are incomparable to the values of the probability.²

Considering the above as the specification of the prior for all values of δ , the next step in estimation is to observe data and establish the posterior distribution for δ given these data. Then we have to work separately for $\delta \neq 0$ and $\delta = 0$. For $\delta \neq 0$, we compute the product of the likelihood and the prior density functions and denote this as the nonnormalized posterior density distribution for $\delta \neq 0$. For $\delta = 0$, we compute the product of ρ_0 and the likelihood at $\delta = 0$ and consider this the nonnormalized posterior probability for $\delta = 0$. Finally, we obtain the full posterior distribution by normalizing the combination of the nonnormalized parts mentioned above. Rouder et al. (2018) say “It is straightforward to update beliefs about δ in the spike-and-slab model using Bayes’ rule”, and then give the posterior in a footnote. For our own interest, but possibly also the reader’s interest, we completely derive this result in Appendix C (and while doing so made a correction).

The hill-and-chimney prior

In general, when estimating a posterior density, the researcher is asked to specify a prior and justify its choice. A common justification is that the prior gives a good representation of the current knowledge and/or the researcher’s belief on the probability density of the parameter of interest. An alternative justification is that a prior should

¹ The slab of the spike-and-slab prior used by Mitchell and Beauchamp (1988) is the uniform distribution on a finite symmetric interval around 0. It is therefore different from Equation 1 but the underlying principle is the same.

² In this paper, we use the notation $P()$ for probabilities and $p()$ for probability densities.

be objective, that is, it should not incorporate a researcher's subjective knowledge or beliefs. The idea is that priors are chosen without committing with too specific information on the parameters of interest (for instance, a wide normal distribution, which ensures that impossible extreme values get low prior densities, while the whole range of realistic values get very similar prior densities). In this way, objective priors are still subjective, or maybe 'intersubjective' (i.e., tentatively agreed upon by a number of people) choices of appropriate almost 'uninformative' priors.

Now as to the spike-and-slab prior, if indeed a researcher has reasons to believe that the value of exactly 0 has a much higher chance of being true than any other values (even values extremely close to 0), the spike-and-slab model offers an appropriate type of prior. It treats the value 0 as qualitatively different from other values for the parameter. We, however, doubt whether the value 0 will often, if ever, have such a special status (see Cohen, 1994; Meehl, 1978; Vardeman, 1987). If choosing such a prior, we think that the researcher must come up with a good reasoning why this could represent his/her current belief and knowledge on the probability distribution of the parameter of interest. This could be difficult without recourse to hard theoretical reasonings, and the latter, unfortunately, do not seem available in behavioral sciences (e.g., Oberauer & Lewandowsky, 2019). The examples often encountered when defending the point null hypothesis are in the realm of extrasensory perception, or precognition (e.g., see Wagenmakers et al., 2018), where on purely theoretical grounds it is considered that exact 0 effects should have a high probability³, but to us such situations seem rare in actual behaviorial science research.

In practice, rather than believing with a sizeable probability that there is exactly no effect at all (even in a population), it seems to make more sense to believe with sizeable probability that the effect is merely *negligible* for all practical purposes (e.g.,

³ Even here one can debate this, if one takes the position that the population parameters should refer to measurements as they actually are being taken (rather than idealized constructs): The procedure of computing the effect size will then not be flawless, and even when a complete full population would be studied, due to operationalization issues, it will be close to impossible that the mean effect size is exactly 0, not even if purported precognition really has no effect whatsoever.

Blume, Greevy, Welty, Smith, & Dupont, 2019). And depending on these practical purposes this could mean, for instance, a mean effect size in the interval $[-.00001, .00001]$ or in the interval $[-.01, .01]$. In fact, 0 is often used as a simplified proxy for such negligible effects. A practically realistic approximation to the spike-and-slab prior could be what we here call the *hill-and-chimney* prior (see Figure 3); the term “chimney” was chosen to replace the word spike, as the spike refers to a line or if you like an extremely narrow rectangle, while the chimney has a noticeable, but still practically negligible width; the term “hill” to replace slab was chosen to come up with a landscape-like metaphor. Rather than assigning probability mass to a spike, the mass is now assigned to an interval, and the graph can now actually express a density across the whole width of this interval. Let the interval of negligibility generally be defined as $[-\varepsilon, \varepsilon]$ where $\varepsilon = .00001$ or $\varepsilon = .01$ in the above two examples. Then the hill-and-chimney prior can be defined as

$$p(\delta) = \begin{cases} k \times \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left(-\frac{\delta^2}{2\sigma_0^2}\right) & , \text{ if } \delta \notin [-\varepsilon, \varepsilon] \\ \frac{\rho_0}{2\varepsilon} & , \text{ if } \delta \in [-\varepsilon, \varepsilon] \end{cases}, \quad (2)$$

where the value k normalizes the function such that the sum of the integrals over $[-\infty, -\varepsilon]$ and $[\varepsilon, \infty]$ equals $(1 - \rho_0)$, and hence, using that the integral over $[-\varepsilon, \varepsilon]$ equals ρ_0 , the integral of the function over $[-\infty, \infty]$ equals 1, which shows that it gives a proper density function (see Appendix D).

To make this more concrete, let’s take a look at the six panels of Figure 3, displaying priors based on $\rho_0 = .5$ and $\sigma_0 = 1$. First of all, panel 3A shows the spike-and-slab prior. In contrast to Figure 2A by Rouder et al. (2018), we distinguish the two sets of units at stake here: Those for the slab and those for the spike, and indicate these alongside the y -axes. We have here chosen to take a length of the spike roughly comparable to that in Figure 2A.

Panels 3B-3F show the hill-and-chimney prior density functions for various values of ε . One might expect that the one for the smallest value of ε (i.e., $\varepsilon = .001$) is the most similar to the spike-and-slab prior. The prior in panel 3B, however, looks quite different compared to the spike-and-slab prior, as the curvy base area seems missing.

This is due to the fact that the chimney has been displayed in its full height, which required shrinking the y -axis considerably. Upon zooming in (panel 3C), now clearly the slightly curvy surface is visible and actually very similar to that in panel 3A. Because the hill-and-chimney approximation employs the same units for the hill and the chimney, we see that the density in the narrow area around 0 is many times higher than the density just outside it. For instance, the density for $\delta = .001$ is roughly 1000 times as high as the density for the closely neighboring value $\delta = .0011$ or the even closer $\delta = .001000001$. The spike-and-slab visualization in Figures 2A and 3A does not reveal this, because the display represents two superimposed incomparable entities in a single graph. By approximating the spike-and-slab prior by the hill-and-chimney prior, we can correctly interpret the plot of superimposed graphs, and we see that for narrow chimneys their heights are immense. In the next section, it will be demonstrated that indeed the hill-and-chimney prior in the limit as ε approaches 0, equals the spike-and-slab prior.

For now, it is interesting to also inspect the other panels of Figure 3. Panels 3D, 3E, and 3F display hill-and-chimney priors for increasing chimney widths, that is, for $\varepsilon = .01$, $\varepsilon = .05$, and $\varepsilon = .1$, respectively. It can be seen that for $\varepsilon = .01$, the chimney still strongly dominates the hill, while for $\varepsilon = .1$ this is no longer strongly the case. However, the interval $[-.1, .1]$ can hardly be considered similar to a spike at 0 (Figure 3F). It could, however, represent a proper way of representing a somewhat strong belief in values of the effect size that are of little or no practical value. One last remark on the hill-and-chimney priors: The ones displayed here have been based on equal probabilities for the chimney and the joint area to the right and left of it. That is, in the plots we took $\rho_0 = .5$. Of course, the degree of dominance of the chimney over the hill would diminish if smaller values for ρ_0 would be taken, and increase if larger values would be taken. If, for instance, one would take $\rho_0 = .001$, the chimney in panel 3C would roughly be at the same height as the curve, and hence no longer dominate it.

Next, it is interesting to see what the full posterior distribution looks like using these hill-and-chimney priors. A derivation of this is given in Appendix D. Two

examples are displayed in Figure 4. Both examples are based on a sample of size 40 and mean of $\mathbf{y}$ equal to .15, but the difference is in the choice of ε . In the left panel $\varepsilon = .1$, and we see a fairly broad chimney for the prior (dashed curve), while in the right panel $\varepsilon = .01$, yielding a much narrower chimney. The posterior for the broad chimney has a somewhat jagged shape, and peaks fairly strongly above $-.1$ and $.1$, but it can also be seen that there is a nonnegligible posterior probability mass to the right of $.1$, which equals $P(\delta > .1|\mathbf{y}) = .14$. Given the posterior distribution, probabilities for any other range of values of δ can be computed: For instance, $P(\delta > 0|\mathbf{y}) = .68$, showing that the probability that δ is positive clearly exceeds that of δ being negative. If for some practical reason only effect sizes above $.3$ are to be taken seriously, it is worthwhile to know that $P(\delta > .3|\mathbf{y}) = .04$, hence it is quite improbable that the effect size is actually relevant in light of this practical relevance threshold. Also, the probability of δ lying in a small interval is interesting. Such an interval could be a “region of practical equivalence” (ROPE) around the value 0 (e.g., see Kruschke, 2018), which could be the interval $[-.3, .3]$. We can directly compute $P(\delta \in [-.3, .3]|\mathbf{y}) = .96$, meaning that, if values within the interval $[-.3, .3]$ refer to negligible (downward or upward) effects, then the probability that the effect is negligible either way is as high as .96. However, possibly the ROPE should be taken much narrower, for instance, as narrow as the chimney. In that case we compute $P(\delta \in [-.1, .1]|\mathbf{y}) = .85$ which shows that the probability of δ being within the area defined by the chimney is still quite high.

Null Hypothesis Bayesian Testing (NHBT)

By Null Hypothesis Bayesian Testing we denote a particular method that allows comparing two hypotheses or models: A point null hypothesis and an alternative hypothesis specifying a range of parameter values, for which a within-hypothesis prior density distribution has been specified. Below, we start by introducing the Bayes factor, and then consider how the Bayes factor can be interpreted directly in terms of the posterior distribution when the spike-and-slab prior has been employed. The spike-and-slab prior will finally allow us to draw a full circle connecting Bayesian

estimation and NHBT.

The Bayes factor

The Bayes factor for $\mathcal{H}_1$ versus $\mathcal{H}_0$, defined by $B_{10} = p(\mathbf{y}|\mathcal{H}_1)/p(\mathbf{y}|\mathcal{H}_0)$, is a measure of the relative predictive ability of both models. The Bayes factor B_{10} is interpreted as a measure of evidence for $\mathcal{H}_1$ versus $\mathcal{H}_0$ provided by the data, and the evidence is in favor of $\mathcal{H}_1$ when $B_{10} > 1$, and for instance, strongly in favor of $\mathcal{H}_1$ if $B_{10} > 10$ (Jeffreys, 1961). The definition of B_{10} implies that $B_{01} = p(\mathbf{y}|\mathcal{H}_0)/p(\mathbf{y}|\mathcal{H}_1) = 1/B_{10}$.

There are now various software packages that allow computing Bayes factors for a wide range of methods (e.g., the R bayesfactor package; Morey & Rouder, 2018; JASP; JASP Team, 2020). JASP in particular offers a very intuitive GUI that facilitates computing Bayes factors for a wide range of informative priors. For illustration, we give the results by JASP for a series of one-sample t -tests for various sample sizes, effect sizes, and priors; see Figure 5⁴. For the standardized effect size δ we used the $\mathcal{N}(0, \sigma_0^2)$ prior available in JASP, for several values of the standard deviation σ_0 . Other common options for the prior available in JASP are the Cauchy and the t -distribution. As can be seen from Figure 5, the larger the effect on the x -axis, the more the evidence in favor of $\mathcal{H}_1$, and this relationship strengthens with sample size. Also, broad priors under $\mathcal{H}_1$ (i.e., larger σ_0 values) are associated with stronger evidence in favor of the null hypothesis. This happens because broad priors dilute the predictive ability of model $\mathcal{H}_1$ over a wide range of parameter values, thus hurting the model's ability to predict the observed data. It is therefore important to choose priors under $\mathcal{H}_1$ judiciously.

The Bayes factor is seen as “strength of evidence from the data about the models, precisely because the strength of evidence should refer to how data lead to revision of beliefs.” (see Rouder et al., 2018, p. 105, referring back to Morey, Romeijn, & Rouder, 2016 and Jeffreys, 1961). Indeed, the Bayes factor is the factor transforming given prior

⁴ In fact, we used R code from Richard Morey that produces exactly the same results as JASP. Code available at <https://osf.io/6rt9m/>.

odds $P(\mathcal{H}_1)/P(\mathcal{H}_0)$ into posterior odds $P(\mathcal{H}_1|\mathbf{y})/P(\mathcal{H}_0|\mathbf{y})$, where here the probabilities refer to probabilities that the hypotheses $\mathcal{H}_1$ and $\mathcal{H}_0$, respectively, are true. This is because $P(\mathcal{H}_1|\mathbf{y})/P(\mathcal{H}_0|\mathbf{y}) = B_{10} \times P(\mathcal{H}_1)/P(\mathcal{H}_0)$, and indeed the role of the Bayes factor is bringing in information from the observed data. So the Bayes factor can be seen as an abstract and general measure of evidence from the data, but it does not lead by itself to concrete conclusions on how probable it is that either $\mathcal{H}_1$ or $\mathcal{H}_0$ is true. For such conclusion drawing and/or decision making one will then need to specify the prior odds⁵ and compute the posterior odds. If, for example, the prior odds are set to 10/90, and $B_{10} = 10$, then the posterior odds are 100/90, so only slightly more in favor of $\mathcal{H}_1$ than of $\mathcal{H}_0$, and clearly not by the 10/1 ratio the Bayes factor might seem to suggest. Only if the prior odds are set to 50/50, the posterior odds equals the Bayes factor and only then the Bayes factor has an appealing interpretation for conclusion drawing. Later in this paper we will also use equal prior odds, but at the same time we want to stress that 50/50 is not any more reasonable than any other prior odds (cf. Hinkley, 1987; Kruschke & Liddell, 2018a).

The NHBT procedure may seem to be a direct alternative for NHST, which compares $\mathcal{H}_0 : \delta = 0$ versus $\mathcal{H}_1 : \delta \neq 0$, but, as Wagenmakers et al. (2018) put it “In Bayesian statistics, this alternative hypothesis needs to be specified exactly.” This means that the hypothesis does not just state that $\delta \neq 0$, but it gives an exact density distribution for δ under $\mathcal{H}_1$. This should be reflected in the process of drawing conclusions, which should not be phrased in terms of evidence for $\delta = 0$ or $\delta \neq 0$, but in terms of “higher (or lower) probability that $\delta = 0$ than that δ has a normal distribution with mean 0 and variance σ_0^2 .” In actual practice, where for instance one is testing the effect of a psychological treatment, formulating the results of a study as “it was found that the probability that the population treatment effect is normally distributed with mean 0 and variance 2 is 12.1 times higher than that the effect of the treatment is 0”,

⁵ As noted by a reviewer, not all Bayes factors allow for prior odds adjustment. One example is based on Bayes factors for inequality constrained hypotheses, which require to use 50/50 prior odds (Klugkist, Laudy, & Hoijtink, 2005).

however, will not appeal to the reader, and one will easily resort to statements like “it was found that the probability that there is a treatment effect is 12.1 times higher than that there is not.” The conclusion is that it seems indeed difficult to deal with the notions of posterior and prior probability of $\mathcal{H}_1$, given that $\mathcal{H}_1$ itself is probabilistic in nature. Luckily, the spike-and-slab prior helps to conceptually bring these concepts together.

The spike-and-slab prior in NHBT

We can conceive of using the spike-and-slab prior given by Equation 1 under the alternative hypothesis $\mathcal{H}_1$. This allows combining the two types of probabilities that we alluded to above. The spike-and-slab prior can be regarded as a weighted sum such that the weights are the prior model probabilities: $\rho_0 = P(\mathcal{H}_0)$ and hence $1 - \rho_0 = P(\mathcal{H}_1)$ (similarly for the spike-and-slab posterior, with $\rho_1 = P(\mathcal{H}_0|\mathbf{y})$ and $1 - \rho_1 = P(\mathcal{H}_1|\mathbf{y})$; see Appendix C). The spike-and-slab prior effectively renders a Bayes model averaged posterior (Hoeting, Madigan, Raftery, & Volinsky, 1999), since the weights of the spike and the slab posterior components correspond to the posterior probabilities of either model. A specification of the probability (density) for each value of δ can be obtained by multiplying the probabilities at stake. That is, for $\delta = 0$ the probability simply is $P(\delta = 0) = \rho_0$, because its probability is specified fully by $P(\mathcal{H}_0)$, while to obtain $p(\delta)$ for other values of δ , we multiply the density specified *within* $\mathcal{H}_1$, that is $p(\delta|\mathcal{H}_1)$, by the probability of $\mathcal{H}_1$ to be true ($P(\mathcal{H}_1)$). Starting from this prior specification, and then estimating the associated posterior probability distribution, leads to a very insightful interpretation of NHBT. Rouder et al. (2018) recently reviewed this earlier noticed intricate relation between NHBT and parameter estimation. In the next section, we will discuss this relation, and prove an interesting approximate relation to priors that are easier to handle.

In line with our running example, let us consider $\mathcal{H}_0 : \delta = 0$ versus $\mathcal{H}_1 : \delta \sim \mathcal{N}(0, \sigma_0^2)$ for known σ . Rouder et al. (2018) give, without proof (but see our Appendix B for a proof), an expression for B_{10} for $\mathcal{H}_0$ and $\mathcal{H}_1$ defined above in terms of

$d = \bar{y}/\sigma$, n , and σ_0 .

Relation between NHBT and the spike-and-slab prior

Rouder et al. (2018, p. 108) point out that “there is an intimate relationship between the spike-and-slab posterior distribution and the Bayes factor.” Indeed, upon observing that

$$\underbrace{\frac{P(\mathcal{H}_1|\mathbf{y})}{P(\mathcal{H}_0|\mathbf{y})}}_{\text{posterior odds}} = B_{10} \underbrace{\frac{P(\mathcal{H}_1)}{P(\mathcal{H}_0)}}_{\text{prior odds}} \quad (3)$$

we have that

$$\frac{1 - \rho_1}{\rho_1} = B_{10} \frac{1 - \rho_0}{\rho_0}, \quad (4)$$

which neatly demonstrates how the Bayes factor relates to the spike-and-slab prior and posterior distributions. However, the spike-and-slab posterior has much more to offer than the posterior odds above: It is possible to compute posterior probabilities for all sorts of ranges of values of δ . So the spike-and-slab based estimation procedure implicitly associated with NHBT actually leads to more information than what NHBT gives on its own. This can be gleaned from Figure 2B, which visualizes the full posterior for a particular data set (in this case, with $\bar{y} = .1$ and $n = 40$, assuming $\sigma = 1$). Here we see that the spike has grown somewhat, while the slab has become less wide and has a smaller surface area under it. From the figure we could conclude, for instance, that $P(\delta > .5|\mathbf{y}) \approx 0$, in other words, the probability that the effect size is larger than .5 is negligible. Also, we can assess that $P(\delta > 0|\mathbf{y}) = .12$, while $P(\delta < 0|\mathbf{y}) = .04$. From NHBT alone, no such conclusion could be drawn. Thus, estimation of the posterior distribution based on the spike-and-slab prior offers the same information as NHBT does, and more.

This relation between NHBT and the spike-and-slab model is bidirectional. We have already mentioned that the spike-and-slab model is a weighted mixture of both hypotheses being tested. Conversely, given the spike-and-slab model, it is simple to derive the hypotheses being tested (Liao et al., 2020, pp. 2-3). However, this is only possible when we know ahead of time the parameter values associated to each

hypothesis. The spike-and-slab mixed model could otherwise be arrived at in many different ways, depending on which of its components were being weighed in together (Kruschke, 2018). Knowing the parameter supports across hypotheses prevents such ambiguity from arising.

We have now seen that posterior density estimation with the spike-and-slab prior offers all information needed for NHBT. That is, using the spike-and-slab prior, from the resulting posterior density one can directly assess B_{10} as

$$\frac{P(\delta \neq 0|\mathbf{y})/P(\delta = 0|\mathbf{y})}{(1 - \rho_0)/\rho_0},$$

which in the special case of prior odds equaling 1, reduces to

$B_{10} = P(\delta \neq 0|\mathbf{y})/P(\delta = 0|\mathbf{y})$. Next we show how the spike-and-slab and hill-and-chimney priors are related and how that is of interest for the Bayes factor.

Relation between spike-and-slab and hill-and-chimney priors and posteriors

In the present section, we demonstrate to what extent using spike-and-slab priors comes down to the same as using hill-and-chimney priors. We do so by offering a formal proof (see Appendix D) to the following intuitive result: *The hill-and-chimney prior converges to the spike-and-slab prior as ε converges to 0*. As a result, a limiting relation between the posterior distributions resulting from the two different priors must also necessarily hold.

It then follows that the posterior odds based on the hill-and-chimney prior, $\frac{P(\delta \notin [-\varepsilon, \varepsilon]|\mathbf{y})}{P(\delta \in [-\varepsilon, \varepsilon]|\mathbf{y})}$, will, in the limit of $\varepsilon \rightarrow 0$, equal the posterior odds based on the spike-and-slab prior (see Equation 4). In other words,

$$\underbrace{B_{10} \frac{1 - \rho_0}{\rho_0} = \frac{P(\delta \neq 0|\mathbf{y})}{P(\delta = 0|\mathbf{y})}}_{\text{spike-and-slab}} \simeq \underbrace{\frac{P(\delta \notin [-\varepsilon, \varepsilon]|\mathbf{y})}{P(\delta \in [-\varepsilon, \varepsilon]|\mathbf{y})}}_{\text{hill-and-chimney}}, \quad (5)$$

where the approximation improves as ε approaches 0.

For practical purposes, it is interesting to get some feeling for how quickly this happens. As an example, we consider data for which $\bar{y} = .05, .15, .25, .35$, and $.45$, and we assume $\sigma = \sigma_0 = 1$ and prior odds equal to 1. The chimney width is gradually

decreased from 0.2 to 0.001, and results were compared with the posterior odds for the spike-and-slab model, which now equal the Bayes factor. To simplify plotting, we focus on the posterior odds $\frac{P(\delta \in [-\varepsilon, \varepsilon] | \mathbf{y})}{P(\delta \notin [-\varepsilon, \varepsilon] | \mathbf{y})}$ and thus on B_{01} , rather than B_{10} . Figure 6 displays the results of these analyses for four different sample sizes, where the posterior odds is displayed on the y -axis against the various chimney widths. It can be seen that indeed, as the chimney width decreases, the posterior odds approaches B_{01} . The speed of the approximation depends on the sample size and on the value of $\bar{y}$, but from all figures it can be seen that substantial differences are seen for the chimney widths from .02 on (albeit only for small effects sizes), while up to .01 differences could be considered negligible.

As mentioned, the hill-and-chimney prior does not approximate the spike-and-slab prior very closely for $\varepsilon = .1$. For example, we can again consider the example previously discussed through Figure 4. The posterior odds $P(\delta \in [-\varepsilon, \varepsilon] | \mathbf{y}) / P(\delta \notin [-\varepsilon, \varepsilon] | \mathbf{y})$ equals 5.59 for $\varepsilon = .1$. This is an approximation to the Bayes factor (associated with the spike-and-slab prior) assuming prior odds equal to 1. A better approximation is obtained for $\varepsilon = .01$, as is displayed in the right-hand panel of Figure 4. Now the posterior odds is 4.23, which is indeed very close to the Bayes factor $B_{01} = 4.13$. The hill-and chimney plot of the prior (dashed line) can be seen to display a far protruding spike (which actually extends till 25), and the posterior is even more spiky, with a top at 42.9. Also here, it can be seen that there still is some mass for values higher than $\varepsilon = .01$ (i.e., to the right of the spike): $P(\delta > .01 | \mathbf{y}) = .16$. Furthermore, one could verify the probability within the ROPE interval $[-.3, .3]$, which now is $P(\delta \in [-.3, .3] | \mathbf{y}) = .97$ or within the interval $[-.1, .1]$, which now is $P(\delta \in [-.1, .1] | \mathbf{y}) = .87$. Interestingly, what we learn now is that, although the two panels in Figure 4 look at first sight quite different, probability statements for even quite small intervals give almost the same values (recall that, based on Figure 4A, $P(\delta \in [-.3, .3] | \mathbf{y}) = .96$ and $P(\delta \in [-.1, .1] | \mathbf{y}) = .85$). Furthermore, we see that it is easy and worthwhile to inspect more information from the posterior distribution than just the posterior odds. When researchers indeed wish to specify a prior in a spike-like

way, using the approximating hill-and-chimney prior could be a good choice which displays all aspects of the prior in a comparable way, that is, in terms of densities. The hill-and-chimney prior, however, still seems unrealistic because of its discontinuity (around the chimney), which also leads to a strangely jagged and discontinuous posterior distribution. In the next section, it will be shown that smooth prior density functions can be found that approximate these hill-and-chimney priors closely.

Approximating hill-and-chimney priors by smooth prior distributions

As mentioned earlier, the hill-and-chimney priors still have a discontinuity at the interval boundaries. It seems desirable to replace it by a continuous density function, so as to avoid that, for instance, prior densities at $\delta = .001$ and $\delta = .0011$ differ enormously⁶. Frühwirth-Schnatter and Wagner (2011) and Wagner and Duller (2012) also discussed a similar search for absolutely continuous distributions in the realm of Bayesian variable selection. Liao et al. (2020) further argued that, in some cases, discontinuous ε -boundary points are associated with poor performance of the Bayes factor, in contrast with the posterior odds. An efficient way of approximating the hill-and-chimney prior distribution (Equation 2) by a smooth continuous density is based on *mollification* (Friedrichs, 1944). The idea consists of considering a so-called mollifier function. The mollifier allows smoothing an irregular or even non-differentiable function by an infinitely differentiable function. The approximation works to a high degree of accuracy.

⁶ As noted by a reviewer, the decision on which prior to use should be ultimately based more on theoretical reasons and information available at the time and less on tradition or convenience.

We first defined the mollifier function as follows⁷:

$$\varphi_\alpha(\delta) = \begin{cases} k \exp \left[\left(\frac{\alpha}{\varepsilon_{\sigma_0}} \right)^2 + \frac{1}{(\delta/\alpha)^2 - (\varepsilon_{\sigma_0}/\alpha)^2} \right], & |\delta| < \varepsilon_{\sigma_0} \\ 0, & |\delta| \geq \varepsilon_{\sigma_0} \end{cases}. \quad (6)$$

The scale parameter α is a positive quantity that tunes the degree of smoothing attained: The larger, the closer the approximation (but too large values may also lead to numerical instability). The value ε_{σ_0} should be some value at least larger than ε (the upper limit of the chimney), in order to allow the approximation around the chimney to work well (in our computations we used $\varepsilon_{\sigma_0} = 4\sigma_0$). Some trial and error to find optimal values for α and ε_{σ_0} is usually required. Constant k adjusts the function so that

$$\int_{\mathbb{R}} \varphi_\alpha(\delta) d\delta = 1.$$

Next, using this mollifier, a continuous, smoothed, hill-and-chimney prior, denoted below as $f_\alpha(\delta)$, is defined as the convolution between the hill-and-chimney prior $p(\delta)$ (Equation 2) and the mollifier in Equation 6:

$$f_\alpha(\delta) = \int_{\mathbb{R}} p(t) \varphi_\alpha(\delta - t) dt. \quad (7)$$

As an example, consider a hill-and-chimney prior defined by $\rho_0 = .5$ and $\varepsilon = .1$ (Figure 7, left panel). As can be seen, the approximation looks very good, and has the advantage of being a smooth continuous function. Now we can again compute a number of specific probabilities under this posterior density function (Figure 7, right panel). For the interval $[-.1, .1]$, we now find a posterior odds of 3.81 compared to 5.59 for the hill-and-chimney prior. We also computed $P(\delta > 0|\mathbf{y}) = .70$, $P(\delta > 0.1|\mathbf{y}) = .18$, and

⁷ Our function was derived from the *standard mollifier* function (Evans, 2010, Section C.5) given by $f(\delta) = k \exp \left(\frac{1}{\delta^2 - 1} \right)$ in the interval $(-1, 1)$ and 0 elsewhere. First we scaled the function to the interval $(-\varepsilon_{\sigma_0}, \varepsilon_{\sigma_0})$: $f_2(\delta) = k \exp \left[\frac{1}{(\delta/\varepsilon_{\sigma_0})^2 - 1} \right]$. Multiplying the numerator and denominator of the argument of the exponential by $\left(\frac{\varepsilon_{\sigma_0}}{\alpha} \right)^2$ leads to the following equivalent function in $(-\varepsilon_{\sigma_0}, \varepsilon_{\sigma_0})$: $f_3(\delta) = k \exp \left[\frac{(\varepsilon_{\sigma_0}/\alpha)^2}{(\delta/\alpha)^2 - (\varepsilon_{\sigma_0}/\alpha)^2} \right]$. Equation 6 results from $f_3(\delta)$ after dropping the numerator of the argument (which is independent from δ) and adding a numerical *adjustment* given by $\left(\frac{\alpha}{\varepsilon_{\sigma_0}} \right)^2$. The adjustment counters the quick convergence of the term $\frac{1}{(\delta/\alpha)^2 - (\varepsilon_{\sigma_0}/\alpha)^2}$ to $-\infty$ as α increases. Observe that the adjustment is independent from δ so it does not affect the mollifier once the normalizing constant k is computed.

$P(\delta > 0.3|\mathbf{y}) = .04$, while for the associated hill-and-chimney, these probabilities were .68, .14, and .04, respectively. Clearly, the odds differs quite a bit, but the probabilities are not that far off.

When, in practice, a researcher wants to specify a prior with a strong peak around $\delta = 0$, it may not be essential that it closely resembles a hill-and-chimney prior, as long as it gets the broad picture. Importantly, one should try to tune the choice of the hill-like and the chimney-like components in such a way that the variance ratio $r = \frac{Var_{\text{spike}}}{Var_{\text{slab}}}$ is much smaller than 1 (e.g., Frühwirth-Schnatter & Wagner, 2011). One alternative way of getting such priors is by using a scaled t -distribution (e.g., Wagner & Duller, 2012) with a low value for the degrees of freedom (e.g., see left panel in Figure 8). For the present situation, after some trial-and-error, we found that the t -distribution with $df = 0.05$ and scale factor 50 (see right panel in Figure 8) gave probabilities fairly close to those found for the (mollified) hill-and-chimney. The posterior odds was 4.06, while $P(\delta > 0|\mathbf{y}) = .63$ and $P(\delta > 0.1|\mathbf{y}) = .18$, which also are fairly close to those associated with the hill-and-chimney results (.68 and .14). In addition, and maybe more importantly, it should be realized that indeed the priors and posteriors for hill-and-chimney and for the smooth priors will lead to similar conclusions for all practical purposes, because by far most of the probability mass is concentrated between $-.1$ and $.1$.

Using the spike-and-slab and hill-and-chimney priors

In the supplementary files available at <https://osf.io/6rt9m/> we have included R scripts and models written in Stan (Stan Development Team, 2021) that allow using various priors, including the hill-and-chimney and approximations based on the normal and t -student distributions. We also include the necessary R code to apply the spike-and-slab model used in this paper.

Decision making

For the bridge between estimation and testing to be fully in place, we must also consider how decisions to be made depend on either inferential framework. First of all,

we think that making discrete decisions is often not a necessity in social sciences. When a researcher is wondering about the merits of a hypothesis such as $\mathcal{H}_0 : \delta = 0$, we think that it is best to express uncertainty about the world through posterior distributions and posterior model probabilities and not necessarily to pick one out of the two hypotheses. The Bayes factor allows comparing the null hypothesis to one, competing, hypothesis, and that will lead to a factor, say, $BF_{01} = 7.3$. But there is little to learn from *one number only*. That the data are 7.3 times more probable under $\mathcal{H}_0$ than under one particular competing hypothesis is not very informative. On the other hand, concluding for instance that the posterior probability of $\mathcal{H}_0$ is .88 for equal prior odds, combined with a range of highly credible values, is much more telling. It is important to recognize that parameter estimation and its uncertainty and magnitude provide qualitatively different information from that provided by Bayes factors. Indeed, as we demonstrate, one can start right away with estimation and still have everything (parameter estimation, uncertainty, *and* testing information like the BF) in one go. Also here, estimation offers much more than the Bayes factor does, even if both can be unified as suggested in this paper.

Having said this, the need to make a binary decision is sometimes unavoidable (also, see Uygun Tunç et al., 2021). NHBT is often associated to decision heuristics aiding at choosing one of either hypotheses being compared, based on the strength of evidence brought about by the data (Jeffreys, 1961; Kass & Raftery, 1995). For example, it is common to use Bayes factor thresholds such as 3 or 10, so that for example if $BF_{01} > 10$ then the researcher is advised to retain the null hypothesis. By using estimation as outlined in this paper, this is of course doable. However, we do not find this procedure sensible in most circumstances. The main reason, as explained above, is that the Bayes factor (i.e., a measure of evidence from the data) differs from measures derived from posterior distributions (i.e., measures of updated belief). The latter reflect the updated probabilistic state of affairs concerning the hypotheses being compared and are arguably what researchers should focus on. The Bayes factor only coincides with the posterior odds if one prespecifies prior odds equal to 1. This

assumption— that both hypotheses being compared are equally likely a priori— is often done in practice so Bayes factor users can somehow ‘get away with it.’ We do not think, however, that equal prior probabilities are always sensible. This is particularly true in the Bayesian framework which lends itself to incorporating prior knowledge in our inferences. As such, options for decision making other than the Bayes factor should also be considered.

Luckily, using Bayesian estimation of parameters opens more possibilities for decision making. As argued by Kruschke (2018), one can resort to the so-called HDI+ROPE criterion. This criterion is based on comparing the HDI for the parameter, which is the range of most credible parameter values, with the ROPE, which defines a range of parameter values that are considered equivalent to each other for all practical purposes. According to the HDI+ROPE criterion, the null hypothesis is accepted if the HDI falls entirely within the ROPE, the alternative hypothesis is accepted if the HDI falls entirely outside the ROPE, and a final decision is suspended otherwise. In this way, a hypothesis such as an interval null may be accepted ‘for practical purposes.’ Kruschke (2018, Supplement) considered decision-theoretic properties of the HDI+ROPE decision rule and showed that this rule is consistent and that it minimizes a certain prototypical loss function, while taking into account information that both the Bayes factor and the posterior odds fail to do.

A simple alternative is to use what Kruschke calls the ROPE alone procedure, suggested for instance by Wellek (2010, sections 2.4, 3.2). This approach consists of computing the posterior probability that the parameter is inside or outside the ROPE. A decision rule can then be set up by comparing these posterior probabilities to a predefined threshold, for instance, .95. Either hypothesis is to be retained if its associated posterior probability is over 95%, otherwise one remains undecided. Kruschke (2018) dismissed this procedure because it does not take into account the varying posterior density associated to the parameter values. However, one could also argue that the total probabilities within or outside the ROPE are all that should count for making a decision. Posterior probabilities are a valid measure of belief and thus we

also think they can serve the purpose of decision making.

Another option for decision making is to focus on the posterior model odds. Unlike the Bayes factor, the posterior model odds are based on the posterior model probabilities, thus incorporate information from both the data and the prior model odds. The posterior model odds are readily available through the hill-and-chimney model (Liao et al., 2020), but some heuristic akin to the Bayes factor thresholds advocated for instance by Jeffreys (1961) would still need to be decided upon. Like the Bayes factor, the posterior odds also fail to consider useful information related to fully expressing model uncertainty that only a full access to the posterior distribution can disclose.

In cases where utilities can be attached to particular effect size values, one can compute the expected utility of, for instance, a treatment, which can then be offset against its costs. A very simple example is where $u(\delta) = 100,000$ if $\delta > 0.2$ and $u(\delta) = 0$ otherwise, while the costs of the treatment are 10,000. For the right panel in Figure 5, $P(\delta > 0.2|\mathbf{y}) = .085$, so the expected utility is 8,500. The costs hence exceed the expected benefit. If, however, $u(\delta) = 100,000$ if $\delta > 0.2$, $u(\delta) = 50,000$ if $\delta \in [.1, .2]$, and $u(\delta) = 0$ otherwise, then, using that $P(\delta \in [.1, .2]|\mathbf{y}) = .099$, the expected benefit is $.099 \times 50,000 + .085 \times 100,000 = 13,450$, which does exceed the costs. Computing the expected utility is beyond reach in NHBT, because NHBT only offers the Bayes factors.

Conclusion

We have seen that, with some effort, one can find a smooth prior that behaves reasonably similarly as a hill-and-chimney prior. In turn, the hill-and-chimney priors can approximate spike-and-slab priors arbitrarily closely as the chimney width tends to 0, but it seems difficult to justify such priors. This is because hill-and-chimney priors specify that the prior belief in δ equaling ε strongly differs from that of δ equaling $\varepsilon + .000001$, while for all practical purposes these two values are equal. Therefore, we think that, in order to justify a prior choice, the prior density should be a smooth continuous function. The mollified and other steep smooth priors discussed before are

continuous and can be scaled such that they have a similar spiky shape as the hill-and-chimney prior, so they seem to be an excellent choice for researchers who actually endorse the spike-and-slab idea that there should be considerable probability mass on values very close to 0. Researchers could, for instance, find and tune spiky priors by adjusting the degrees of freedom and scale factor of the t -distribution at wish, and could find a prior sufficiently close to their ideal of a peaked mass close to 0. Having done so, they can compute a Bayes factor, but also any posterior probabilities they like, as well as HDIs, straightforwardly from the posterior distribution. In this way, one gets full insight in the posterior distribution, which will considerably enhance the understanding of the implications of the evidence provided by the Bayes factor (or any other single value) alone.

Following Rouder et al. (2018), a true unification between estimation and a variant of NHBT has been obtained. Actually, NHBT boils down to obtaining merely one particular probability statement based on aspects of the posterior distribution (i.e., the one that specifies the odds of being close to 0 against its opposite). However, having the posterior distribution, any so desired probability statement on the values of δ can be made, for instance, $P(\delta \in \text{ROPE}|\mathbf{y})$, $P(\delta > 0|\mathbf{y})$, or $P(\delta > \text{practically significant effect}|\mathbf{y})$.

Maybe the most important feature of using smooth priors is that it allows one to move smoothly from strongly peaked priors to gradually flattening priors, and the researcher can choose the prior that is best justifiable according to his/her knowledge. Maybe in some or even many cases, it is unreasonable to hold a strong belief in near zero effect sizes. For instance, when testing treatments, often a long stage of development has preceded this, and it is to be expected that there actually is a substantial effect. In such cases, fairly flat priors, or even shifted priors around a nonzero effect size could be more reasonable. One might counter that this is ‘unfair’ as it leads to confirmation bias. However, the very idea of Bayesian analysis is that, in the data analysis, prior information can be introduced, which in the eyes of others indeed may be a bias with respect to those others’ norms. This prior is subjective, but it

should be justified. In a way, one could say it is meant to ‘bias’/adjust a statement which one should make on the basis of the data alone, because on a particular phenomenon more information is available than just the present data. If there is reason to adjust a result towards small effects, then this can be done by putting a high mass on small values (as in the spike-and-slab prior), but if there is current knowledge about a substantial effect size, nothing is wrong with adjusting one’s results towards this. Typically, in such cases one will rarely find high posterior probabilities in close areas around 0. This, however, is then fully in agreement with current knowledge, and especially when one deals with small ROPEs, it will indeed be very unlikely that the population mean effect will be in there.

Discussion

The present paper elaborated on the unification of NHBT and Bayesian estimation described by Rouder et al. (2018). Following their approach, it did so for only one particular case of NHBT, that is, the one defining the alternative hypothesis in terms of $\delta \sim \mathcal{N}(0, \sigma_0^2)$, while assuming σ , the standard deviation of the population scores y , known. The general reasoning was easy to illustrate for this situation, but it is rather unrealistic in practice, as σ will never be known, and can even be hard to guess. This is a limitation in our model. A more general approach is described by Rouder, Speckman, Sun, and Morey (2009) and uses, in addition to a different prior distribution for δ (viz. a Cauchy distribution), also a prior distribution for σ , which in computations, however is ‘integrated out.’ We do not know whether a particular class of hill-and-chimney priors exists that arbitrarily closely approximates this class of Cauchy priors, but we do believe that hill-and-chimney priors can be constructed that come close to it ‘for all practical purposes.’ Therefore, we believe that our reasoning is of general relevance. Likewise, we expect smooth spiky curves can be found to mimic the Cauchy prior used, while this can be combined with a prior for the unknown σ . Future research may thus aim at generalizing the hill-and-chimney paradigm to more designs. We think that the best way forward is to combine these ideas with MCMC sampling, as

this will help broadening the implementation in practice.

Following Kruschke (2013), even an additional parameter for tuning the thickness of the tails of the prior could be used. We are convinced that such classes of priors can cater for diverse states of knowledge, that is, where there is strong knowledge that the effect size is close to 0 (or another value), as well as cases where only vague knowledge exists. It has been seen now that the spike-and-slab prior focuses only on the former case, which is fine, if justifiable. A more general approach, however, is offered by full posterior estimation, allowing for any type of prior, and compared to NHBT nothing is lost, but a lot gained.

Open Practices Statement

The R scripts including the code required to reproduce all computations and figures available in this paper are available at <https://osf.io/6rt9m/>. No experiments were conducted and therefore preregistered.

Funding

This research was supported by a Japanese JSPS KAKENHI grant awarded to Jorge N. Tendeiro (Grant number 21K20211).

References

- Blume, J. D., Greevy, R. A., Welty, V. F., Smith, J. R., & Dupont, W. D. (2019). An Introduction to Second-Generation p -Values. *The American Statistician*, *73*(sup1), 157–167. doi: 10.1080/00031305.2018.1537893
- Cohen, J. (1994). The earth is round ($p < .05$). *The American Psychologist*, *49*, 997–1003.
- Etz, A., & Vandekerckhove, J. (2018, February). Introduction to Bayesian Inference for Psychology. *Psychonomic Bulletin & Review*, *25*(1), 5–34. doi: 10.3758/s13423-017-1262-3
- Evans, L. C. (2010). *Partial Differential Equations* (Second ed.). Providence, Rhode Island: American Mathematical Society.
- Friedrichs, K. O. (1944). The identity of weak and strong extensions of differential operators. *Transactions of the American Mathematical Society*, *55*, 132–151. doi: 10.1090/S0002-9947-1944-0009701-0
- Frühwirth-Schnatter, S., & Wagner, H. (2011). Bayesian Variable Selection for Random Intercept Modeling of Gaussian and Non-Gaussian Data. In J. M. Bernardo et al. (Eds.), *Bayesian Statistics 9* (pp. 165–200). Oxford University Press. doi: 10.1093/acprof:oso/9780199694587.003.0006
- 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.
- George, E. I., & McCulloch, R. E. (1993). Variable Selection via Gibbs Sampling. *Journal of the American Statistical Association*, *88*(423), 881–889. doi: 10.1080/01621459.1993.10476353
- George, E. I., & McCulloch, R. E. (1997). Approaches for Bayesian variable selection. *Statistica Sinica*, *7*, 339–373.
- Hinkley, D. V. (1987). Testing a Point Null Hypothesis: The Irreconcilability of p Values and Evidence: Comment. *Journal of the American Statistical Association*, *82*, 128–129.
- Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian model

- averaging: A tutorial. *Statistical Science*, 14, 382–401.
- Ishwaran, H., & Rao, J. S. (2003). Detecting Differentially Expressed Genes in Microarrays Using Bayesian Model Selection. *Journal of the American Statistical Association*, 98(462), 438–455. doi: 10.1198/0162145030000224
- Ishwaran, H., & Rao, J. S. (2005). Spike and slab variable selection: Frequentist and Bayesian strategies. *The Annals of Statistics*, 33(2), 730–773. doi: 10.1214/009053604000001147
- 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
- Klugkist, I., Laudy, O., & Hoijsink, H. (2005, December). Inequality Constrained Analysis of Variance: A Bayesian Approach. *Psychological Methods*, 10(4), 477–493. doi: 10.1037/1082-989X.10.4.477
- 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. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, 142(2), 573–603. doi: 10.1037/a0029146
- Kruschke, J. K. (2015). *Doing Bayesian data analysis: A tutorial with R, JAGS, and Stan* (Edition 2 ed.). Boston: Academic Press.
- Kruschke, J. K. (2018). Rejecting or Accepting Parameter Values in Bayesian Estimation. *Advances in Methods and Practices in Psychological Science*, 1(2), 270–280. doi: 10.1177/2515245918771304
- Kruschke, J. K., & Liddell, T. M. (2018a). Bayesian data analysis for newcomers. *Psychonomic Bulletin & Review*, 25, 155–177.
- Kruschke, J. K., & Liddell, T. M. (2018b). The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*, 25, 178–206.

- Kuo, L., & Mallick, B. (1998). Variable selection for regression models. *Sankhyā: The Indian Journal of Statistics, Series B*, 60, 65–81.
- Liao, J. G., Midya, V., & Berg, A. (2020). Connecting and Contrasting the Bayes Factor and a Modified ROPE Procedure for Testing Interval Null Hypotheses. *The American Statistician*, 1–19. doi: 10.1080/00031305.2019.1701550
- Malsiner-Walli, G., & Wagner, H. (2011). Comparing Spike and Slab Priors for Bayesian Variable Selection. *Austrian Journal of Statistics*, 40, 241–264.
- Meehl, P. E. (1978). Theoretical risks and tabular asterisks: Sir Karl, Sir Ronald, and the slow progress of soft psychology. *Journal of Consulting and Clinical Psychology*, 46(4), 806–834. doi: 10.1037/0022-006X.46.4.806
- 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., 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. doi: 10.1016/j.jmp.2015.11.001
- Morey, R. D., & Rouder, J. N. (2011). Bayes Factor approaches for testing interval null hypotheses. *Psychological Methods*, 16, 406–419.
- Morey, R. D., & Rouder, J. N. (2018). *BayesFactor: Computation of Bayes Factors for Common Designs*.
- Ntzoufras, I., Forster, J. J., & Dellaportas, P. (2000). Stochastic search variable selection for log-linear models. *Journal of Statistical Computation and Simulation*, 68(1), 23–37. doi: 10.1080/00949650008812054
- Oberauer, K., & Lewandowsky, S. (2019). Addressing the theory crisis in psychology. *Psychonomic Bulletin & Review*, 26(5), 1596–1618. doi: 10.3758/s13423-019-01645-2
- O’Hara, R. B., & Sillanpää, M. J. (2009). A review of Bayesian variable selection methods: What, how and which. *Bayesian Analysis*, 4, 85–118. doi: 10.1214/09-BA403

- Puga, J. L., Krzywinski, M., & Altman, N. (2015, April). Bayes' theorem. *Nature Methods*, 12(4), 277–278. doi: 10.1038/nmeth.3335
- 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.
- Rouder, J. N., Speckman, P. L., Sun, D., & Morey, R. D. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16, 225–237.
- Stan Development Team. (2021). *Stan Modeling Language Users Guide and Reference Manual*, 2.26. <https://mc-stan.org/>.
- Stern, H. S. (2016). A Test by Any Other Name: P Values, Bayes Factors, and Statistical Inference. *Multivariate Behavioral Research*, 51(1), 23–29. doi: 10.1080/00273171.2015.1099032
- 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
- Uygun Tunç, D., Tunç, M. N., & Lakens, D. (2021). *The Epistemic and Pragmatic Function of Dichotomous Claims Based on Statistical Hypothesis Tests* (Preprint). PsyArXiv. doi: 10.31234/osf.io/af9by
- van de Schoot, R., Depaoli, S., King, R., Kramer, B., Märtens, K., Tadesse, M. G., ... Yau, C. (2021, December). Bayesian statistics and modelling. *Nature Reviews Methods Primers*, 1(1), 1. doi: 10.1038/s43586-020-00001-2
- van de Schoot, R., Kaplan, D., Denissen, J., Asendorpf, J. B., Neyer, F. J., & van Aken, M. A. G. (2014). A Gentle Introduction to Bayesian Analysis: Applications to Developmental Research. *Child Development*, 85(3), 842–860. doi: 10.1111/cdev.12169

- van der Linden, S., & Chryst, B. (2017). No need for Bayes factors: A fully Bayesian evidence synthesis. *Frontiers in Applied Mathematics and Statistics*, 3:12, 1–3.
- van Ravenzwaaij, D., & Wagenmakers, E.-J. (2021, December). Advantages masquerading as “issues” in Bayesian hypothesis testing: A commentary on Tendeiro and Kiers (2019). *Psychological Methods*. doi: 10.1037/met0000415
- van Ravenzwaaij, D., & Wagenmakers, E. (2021). *Advantages masquerading as “issues” in Bayesian hypothesis testing: A commentary on Tendeiro and Kiers (2019)*. (No. Query date: 2021-12-22 15:03:02). psycnet.apa.org.
- Vardeman, S. B. (1987). Testing a Point Null Hypothesis: The Irreconcilability of p Values and Evidence: Comment. *Journal of the American Statistical Association*, 82, 130–131.
- Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage–Dickey method. *Cognitive Psychology*, 60, 158–189.
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., . . . Gronau, Q. F. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review*, 25, 35–57. doi: 10.3758/s13423-017-1343-3
- Wagner, H., & Duller, C. (2012). Bayesian model selection for logistic regression models with random intercept. *Computational Statistics & Data Analysis*, 56(5), 1256–1274. doi: 10.1016/j.csda.2011.06.033
- Wellek, S. (2010). *Testing statistical hypotheses of equivalence and noninferiority* (Second ed.). Boca Raton: Chapman and Hall/CRC Press.
- Williams, M. N., Bååth, R. A., & Philipp, M. C. (2017). Using Bayes Factors to Test Hypotheses in Developmental Research. *Research in Human Development*, 14(4), 321–337. doi: 10.1080/15427609.2017.1370964

Figure 1. Bayesian estimation for parameter δ . Data are assumed to be normally distributed with (unknown) mean δ and standard deviation 1 in the population. The prior distribution is $\mathcal{N}(0, .3^2)$. The sample of size 20 has sample mean equal to 0.5. The posterior distribution is the normalized product of the prior and the data distribution and it offers a compromise between both sources of information. We refer to Appendix A for the mathematical details of this posterior distribution.

Figure 2. Visualisation of the spike-and-slab model by Rouder et al. (2018), copied from their manuscript, with permission by the publisher. A: Prior distribution on effect size $\delta = \mu/\sigma$, with half the mass in the spike and the slab centered around zero. B: Posterior distribution on effect size δ for an observed effect size of $d = .1$ based on a sample size of 40.

Figure 3. The spike-and-slab prior (A), and the chimney-and-hill prior for $\varepsilon = .001$ (B); panel C displays the same prior as panel B, but with the scale of the y -axis set equal to that of panel A, so note that the ‘chimney’ actually continues with more than 60 times the height displayed; the next panels show chimney-and-hill priors for $\varepsilon = .01$ (D), $\varepsilon = .05$ (E), and $\varepsilon = .1$ (F). All priors were based on $\rho_0 = .5$ and $\sigma_0 = 1$. The spike-and-slab prior is based on multiplying the $\mathcal{N}(0, 1)$ density by .5, and the spike is supposed to represent .5 probability on $\delta = 0$. In all cases the priors have been displayed on the ranges $[-1, 1]$. Note that the scales of the density on the y -axes differ markedly.

Figure 4. Two hill-and-chimney priors (dashed curves), with associated posteriors (solid curves) and likelihood (dotted curve), for data from a sample of size 40 and with mean effect size .15. The left panel is based on a chimney for $\varepsilon = .1$ (i.e., width = .2) and the right panel is based on a chimney for $\varepsilon = .01$ (i.e., width = .02).

Figure 5. The one-sample t -test Bayes factor computed as in JASP ($\mathcal{H}_0 : \delta = 0$ versus $\mathcal{H}_1 : \delta \sim \mathcal{N}(0, \sigma_0^2)$), for various sample sizes (20, 40, 80, 160), observed effect sizes ($d = 0, .1, .2, .3$), and priors ($\mathcal{N}(0, \sigma_0^2)$ for $\sigma_0 = 0.05, 0.5, 1, 2$). Data were randomly sampled from a normal distribution such that their sample mean and standard deviation are equal to d and 1, respectively. The y -axis is on the log-scale.

Figure 6. Posterior odds (vertical axis) obtained using hill-and-chimney priors plotted against the various chimney widths (horizontal axis, log-scale), for varying sample sizes (20, 40, 80, 160) and mean values of y (.05, .15, .25, .35, .45). Grey horizontal lines represent the values of the associated spike-and-slab posterior odds.

Figure 7. Approximating the hill-and-chimney prior by mollification. Left panel: The hill-and-chimney prior (dashed line) with $\rho_0 = .5$ and $\varepsilon = .1$, and the mollified prior (solid line), for $\alpha = 600$. Right panel: The hill-and-chimney posterior (dashed line) and the posterior based on the mollified prior (solid line). Based on data with $n = 40$, $\delta = .15$, $\sigma = 1$, and $\sigma_0 = 1$.

Figure 8. Results based on smooth peaked priors, based on data with $n = 40$, $\delta = .15$, $\sigma = 1$, and $\sigma_0 = 1$. The prior in the left panel is the t -distribution with $df = 1$, scaled by factor 10. The prior in the right panel is the t -distribution with $df = .05$, scaled by factor 50.

Appendix A

Useful lemma

The following lemma will be useful to derive the expression of the Bayes factor for the test used in the paper (Appendix B), the spike-and-slab posterior distribution (Appendix C), and the hill-and-chimney posterior distribution (Appendix D). It is a variation of a common result in Bayesian analysis based on conjugating a normal prior with normal data. In the remaining of the paper we will refer to it simply as ‘the Lemma.’

Lemma. *Let $y_i \sim \mathcal{N}(\sigma\delta, \sigma^2)$ for $i = 1, \dots, n$ and σ^2 known. If $\delta \sim \mathcal{N}(0, \sigma_0^2)$ is the prior density for δ , then the (unnormalized) product of the prior for δ by the likelihood is*

$$\text{prior} \times \text{likelihood} = \lambda \exp \left[-\frac{(\delta - \tilde{\mu})^2}{2\tilde{\sigma}^2} \right], \quad (\text{A1})$$

where $\lambda = \frac{1}{\sqrt{2\pi}\sigma_0} \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \exp \left(-\frac{n\bar{y}^2}{2\sigma^2} \right) \exp \left[\frac{(n\bar{y}\sigma_0)^2}{2\sigma^2(1+n\sigma_0^2)} \right]$ does not depend on δ and the exponential is the kernel of a normal density with mean $\tilde{\mu} = \frac{n\bar{y}\sigma_0^2}{\sigma(1+n\sigma_0^2)}$ and variance $\tilde{\sigma}^2 = \frac{\sigma_0^2}{1+n\sigma_0^2}$.

Proof. We have that

$$\text{prior} \times \text{likelihood} = \frac{1}{\sqrt{2\pi}\sigma_0} \exp \left(-\frac{\delta^2}{2\sigma_0^2} \right) \times \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \exp \left[-\frac{\sum_{i=1}^n (y_i - \sigma\delta)^2}{2\sigma^2} \right]. \quad (\text{A2})$$

Next we combine the two exponentials in Equation A2 and rewrite the argument of the resulting exponential as a quadratic function of δ :

$$\begin{aligned} \text{prior} \times \text{likelihood} &= \frac{1}{\sqrt{2\pi}\sigma_0} \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \exp \left(-\frac{\delta^2}{2\sigma_0^2} - \frac{\sum_{i=1}^n y_i^2}{2\sigma^2} - \frac{n\sigma^2\delta^2}{2\sigma^2} + \frac{2n\sigma\delta\bar{y}}{2\sigma^2} \right) \\ &= \frac{1}{\sqrt{2\pi}\sigma_0} \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \exp \left(-\frac{n\bar{y}^2}{2\sigma^2} \right) \\ &\quad \exp \left\{ -\frac{1}{2} \left[\left(\frac{1+n\sigma_0^2}{\sigma_0^2} \right) \delta^2 - \frac{2n\bar{y}}{\sigma} \delta \right] \right\}. \end{aligned} \quad (\text{A3})$$

Finally, we ‘complete the square’ of the quadratic function of δ :

$$\begin{aligned}
\text{prior} \times \text{likelihood} &= \frac{1}{\sqrt{2\pi}\sigma_0} \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \exp \left(-\frac{n\bar{y}^2}{2\sigma^2} \right) \exp \left[\frac{(n\bar{y}\sigma_0)^2}{2\sigma^2(1+n\sigma_0^2)} \right] \\
&\quad \exp \left\{ -\frac{1}{2} \left[\left(\frac{1+n\sigma_0^2}{\sigma_0^2} \right) \delta^2 - \frac{2n\bar{y}}{\sigma} \delta + \left(\frac{n\bar{y}\sigma_0}{\sigma\sqrt{1+n\sigma_0^2}} \right)^2 \right] \right\} \\
&= \lambda \exp \left[-\frac{1}{2} \left(\frac{\sqrt{1+n\sigma_0^2}}{\sigma_0} \delta - \frac{n\bar{y}\sigma_0}{\sigma\sqrt{1+n\sigma_0^2}} \right)^2 \right] \\
&= \lambda \exp \left[-\frac{\left(\delta - \frac{n\bar{y}\sigma_0^2}{\sigma(1+n\sigma_0^2)} \right)^2}{2\frac{\sigma_0^2}{1+n\sigma_0^2}} \right] \\
&= \lambda \exp \left[-\frac{(\delta - \tilde{\mu})^2}{2\tilde{\sigma}^2} \right].
\end{aligned}$$

This completes the proof. □

In particular, the Lemma implies that the posterior over the entire real numbers domain is the normal distribution $\mathcal{N}(\tilde{\mu}, \tilde{\sigma}^2)$. The multiplicative constant λ is of value when Equation A1 is constrained to a subset of the real numbers set, as it is the case in deriving the hill-and-chimney posterior distribution.

Appendix B

Derivation of the Bayes factor

We want to derive the Bayes factor B_{10} to test $\mathcal{H}_0 : \delta = 0$ versus $\mathcal{H}_1 : \delta \sim \mathcal{N}(0, \sigma_0^2)$, for data $\mathbf{y} = \{y_1, \dots, y_n\}$ assumed normally distributed,

$$y_i \sim \mathcal{N}(\mu, \sigma^2) = \mathcal{N}(\sigma\delta, \sigma^2), i = 1, \dots, n. \quad (\text{B1})$$

Observe that δ is the standardized effect size $\delta = \mu/\sigma$ and that σ^2 is assumed known.

Here we will derive B_{10} directly by its definition⁸. We have that

$$\begin{aligned} B_{10} &= \frac{p(\mathbf{y}|\mathcal{H}_1)}{p(\mathbf{y}|\mathcal{H}_0)} \\ &= \frac{\int_{-\infty}^{\infty} p(\mathbf{y}, \delta|\mathcal{H}_1) p(\delta|\mathcal{H}_1) d\delta}{\prod_{i=1}^n \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\frac{(y_i-0)^2}{2\sigma^2}\right]}. \end{aligned}$$

Applying the Lemma to the integrand in the numerator we have that

$$\begin{aligned} B_{10} &= \frac{\frac{1}{\sqrt{2\pi}\sigma_0} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^n \exp\left(-\frac{n\bar{y}^2}{2\sigma^2}\right) \exp\left[\frac{(n\bar{y}\sigma_0)^2}{2\sigma^2(1+n\sigma_0^2)}\right] \int_{-\infty}^{\infty} \exp\left[-\frac{(\delta - \tilde{\mu})^2}{2\tilde{\sigma}^2}\right] d\delta}{\left(\frac{1}{\sqrt{2\pi}\sigma}\right)^n \exp\left(-\frac{n\bar{y}^2}{2\sigma^2}\right)} \\ &= \frac{1}{\sqrt{2\pi}\sigma_0} \exp\left[\frac{(n\bar{y}\sigma_0)^2}{2\sigma^2(1+n\sigma_0^2)}\right] \sqrt{2\pi}\tilde{\sigma} \\ &= \frac{1}{\sqrt{1+n\sigma_0^2}} \exp\left[\frac{(n\bar{y}\sigma_0)^2}{2\sigma^2(1+n\sigma_0^2)}\right]. \end{aligned} \quad (\text{B2})$$

This equals the expression of B_{10} provided by Rouder et al. (2018, Footnote 2).

⁸ Alternatively, B_{10} can also be derived by means of the Savage-Dickey density ratio for nested models (Wagenmakers, Lodewyckx, Kuriyal, & Grasman, 2010).

Appendix C

Derivation of the spike-and-slab posterior distribution

Here we derive the posterior distribution for δ , given the prior defined in Equation 1 and the same normal data model as given by Equation B1. The proof consists of firstly computing the posterior probability of $\mathcal{H}_0$, and then assessing the posterior density for $\delta \neq 0$. The posterior is:

$$\begin{cases} p(\delta|\mathbf{y}) &= (1 - \rho_1) \times \frac{1}{\sqrt{2\pi}\sigma_1} \exp \left[-\frac{(\delta - \delta_1)^2}{2\sigma_1^2} \right] & , \text{ if } \delta \neq 0 \\ P(\delta = 0|\mathbf{y}) &= \rho_1 & , \text{ if } \delta = 0 \end{cases}, \quad (\text{C1})$$

where

$$\begin{aligned} \rho_1 &= \frac{\rho_0}{\rho_0 + (1 - \rho_0)B_{10}} \\ \delta_1 &= \frac{n\bar{y}\sigma_1^2}{\sigma}, \text{ and} \\ \sigma_1^2 &= \left(n + \frac{1}{\sigma_0^2} \right)^{-1}. \end{aligned}$$

Note that we did not use Rouder et al.'s shortcut $d = \bar{y}/\sigma$. Hence this posterior for δ depends on $\bar{y}$, n , σ_0 , and σ .

We start by considering the basic equation that relates the Bayes factor with the prior and posterior odds (Equation 3):

$$\frac{P(\mathcal{H}_1|\mathbf{y})}{P(\mathcal{H}_0|\mathbf{y})} = B_{10} \times \frac{P(\mathcal{H}_1)}{P(\mathcal{H}_0)}.$$

We know that $P(\mathcal{H}_0) = \rho_0$ and $P(\mathcal{H}_1) = 1 - \rho_0$ (see Equation 1), and that $P(\mathcal{H}_1|\mathbf{y}) = 1 - P(\mathcal{H}_0|\mathbf{y})$. We therefore have the following:

$$\begin{aligned} \frac{1 - P(\mathcal{H}_0|\mathbf{y})}{P(\mathcal{H}_0|\mathbf{y})} &= B_{10} \times \frac{1 - \rho_0}{\rho_0} \\ \frac{1}{P(\mathcal{H}_0|\mathbf{y})} &= 1 + B_{10} \times \frac{1 - \rho_0}{\rho_0} \\ P(\mathcal{H}_0|\mathbf{y}) &= \frac{\rho_0}{\rho_0 + (1 - \rho_0)B_{10}} \\ P(\mathcal{H}_0|\mathbf{y}) &= \rho_1. \end{aligned}$$

Thus we established that the posterior probability of $\mathcal{H}_0 : \delta = 0$ is equal to ρ_1 .

In order to find the posterior density $p(\delta|\mathbf{y})$ for $\delta \neq 0$, we observe that this posterior is proportional to the product given by Equation A1 from the Lemma, that is,

$$p(\delta|\mathbf{y}) \propto \exp \left[-\frac{(\delta - \tilde{\mu})^2}{2\tilde{\sigma}^2} \right].$$

The total mass for $\delta \neq 0$ should be $(1 - \rho_1)$, so we conclude that the sought density must be

$$p(\delta|\mathbf{y}) = (1 - \rho_1) \frac{1}{\sqrt{2\pi}\tilde{\sigma}} \exp \left[-\frac{(\delta - \tilde{\mu})^2}{2\tilde{\sigma}^2} \right].$$

This is equal to the expression in Equation C1 after noting that $\sigma_1 = \tilde{\sigma}$ and $\delta_1 = \tilde{\mu}$.

Appendix D

The hill-and-chimney prior (Equation 2)

In this section we establish two results based on the hill-and-chimney prior (Equation 2). Firstly, we derive the expression of constant k in Equation 2. Next, we derive the expression of the posterior distribution for δ , given the hill-and-chimney prior and the normal data model from Equation B1.

Constant k

In order to find k , it is useful to consider the following relation:

$$\frac{1 - \rho_0}{2} = \int_{-\infty}^{-\varepsilon} p(\delta) d\delta = k\Phi\left(-\frac{\varepsilon}{\sigma_0}\right). \quad (\text{D1})$$

This first equality holds because the hill-and-chimney prior is symmetric around 0 and the chimney area is by definition equal to ρ_0 (since $\int_{-\varepsilon}^{\varepsilon} \frac{\rho_0}{2\varepsilon} d\delta = \rho_0$). Therefore, the area under the slab ($= 1 - \rho_0$) is evenly divided across the intervals $(-\infty, -\varepsilon)$ and (ε, ∞) . The second equality holds because the slab is based on the (rescaled) normal density $\mathcal{N}(0, \sigma_0^2)$. Solving Equation D1 with respect to k implies that

$$k = \frac{1 - \rho_0}{2\Phi\left(-\frac{\varepsilon}{\sigma_0}\right)}. \quad (\text{D2})$$

Derivation of the hill-and-chimney posterior distribution

We first find the unnormalized posterior (i.e., the product of the prior by the likelihood), and then we find the normalization factor, which is the marginal likelihood. Jointly these will lead to the posterior distribution.

The unnormalized posterior. We must consider two situations, namely, $\delta \in [-\varepsilon, \varepsilon]$ and $\delta \notin [-\varepsilon, \varepsilon]$.

When $\delta \in [-\varepsilon, \varepsilon]$ we have that

$$\begin{aligned} \text{prior} \times \text{likelihood} &= p(\delta)L(\delta|\mathbf{y}) \\ &= \frac{\rho_0}{2\varepsilon} \left(\frac{1}{\sqrt{2\pi}\sigma}\right)^n \exp\left[-\frac{\sum_{i=1}^n (y_i - \sigma\delta)^2}{2\sigma^2}\right]. \end{aligned} \quad (\text{D3})$$

The exponential in Equation D3 can be decomposed as follows:

$$\begin{aligned} \exp \left[-\frac{\sum_{i=1}^n (y_i - \sigma\delta)^2}{2\sigma^2} \right] &= \exp \left[-\frac{1}{2\sigma^2} \left(n\sigma^2\delta^2 - 2\sum_{i=1}^n y_i\sigma\delta + \sum_{i=1}^n y_i^2 \right) \right] \\ &= \exp \left[-\frac{n}{2\sigma^2} \underbrace{(\sigma^2\delta^2 - 2\bar{y}\sigma\delta + \bar{y}^2)}_{\text{quadratic in } \delta} \right]. \end{aligned}$$

Next we complete the square of the quadratic function in δ :

$$\begin{aligned} \exp \left[-\frac{n}{2\sigma^2} (\sigma^2\delta^2 - 2\bar{y}\sigma\delta + \bar{y}^2) \right] &= \exp \left[-\frac{n}{2\sigma^2} (\sigma^2\delta^2 - 2\bar{y}\sigma\delta + \bar{y}^2 - \bar{y}^2 + \bar{y}^2) \right] \\ &= \beta_1 \exp \left[-\frac{n}{2\sigma^2} (\sigma\delta - \bar{y})^2 \right] \\ &= \beta_1 \exp \left[-\frac{n}{2} \left(\delta - \frac{\bar{y}}{\sigma} \right)^2 \right], \end{aligned} \quad (\text{D4})$$

where $\beta_1 = \exp \left[-\frac{n}{2\sigma^2} (\bar{y}^2 - \bar{y}^2) \right]$. Replacing the exponential in Equation D3 by Equation D4 finally leads to

$$\text{prior} \times \text{likelihood} = \frac{\rho_0}{2\varepsilon} \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \beta_1 \exp \left[-\frac{n}{2} \left(\delta - \frac{\bar{y}}{\sigma} \right)^2 \right]. \quad (\text{D5})$$

The rightmost exponential factor in Equation D5 is the kernel of a normal density with mean $\mu_{\text{chimney}} = \frac{\bar{y}}{\sigma}$ and variance $\sigma_{\text{chimney}}^2 = \frac{1}{n}$.

When $\delta \notin [-\varepsilon, \varepsilon]$ we can use the Lemma:

$$\begin{aligned} \text{prior} \times \text{likelihood} &= p(\delta) L(\delta | \mathbf{y}) \\ &= k\lambda \exp \left[-\frac{(\delta - \tilde{\mu})^2}{2\tilde{\sigma}^2} \right]. \end{aligned} \quad (\text{D6})$$

We conclude that the unnormalized posterior is given by Equation D5 when $\delta = 0$ and by Equation D6 when $\delta \neq 0$.

The marginal likelihood. The marginal likelihood is given by $\int_{-\infty}^{\infty} \text{prior} \times \text{likelihood} d\delta$, by definition. This integral can be computed as a sum of two terms, namely,

$$\int_{\delta \in [-\varepsilon, \varepsilon]} \text{prior} \times \text{likelihood} d\delta + \int_{\delta \notin [-\varepsilon, \varepsilon]} \text{prior} \times \text{likelihood} d\delta. \quad (\text{D7})$$

The first term of Equation D7 is equal to (recall Equation D5)

$$\begin{aligned} \int_{\delta \in [-\varepsilon, \varepsilon]} \text{prior} \times \text{likelihood} d\delta &= \frac{\rho_0}{2\varepsilon} \left(\frac{1}{\sqrt{2\pi}\sigma} \right)^n \beta_1 \sqrt{2\pi} \sigma_{\text{chimney}} \\ &\quad \left[\Phi \left(\frac{\varepsilon - \mu_{\text{chimney}}}{\sigma_{\text{chimney}}} \right) - \Phi \left(\frac{-\varepsilon - \mu_{\text{chimney}}}{\sigma_{\text{chimney}}} \right) \right]. \end{aligned} \quad (\text{D8})$$

The second term of Equation D7 is equal to (recall Equation D6)

$$\int_{\delta \notin [-\varepsilon, \varepsilon]} \text{prior} \times \text{likelihood} d\delta = k\lambda\sqrt{2\pi}\tilde{\sigma} \left[1 - \Phi\left(\frac{\varepsilon - \tilde{\mu}}{\tilde{\sigma}}\right) + \Phi\left(\frac{-\varepsilon - \tilde{\mu}}{\tilde{\sigma}}\right) \right]. \quad (\text{D9})$$

The hill-and-chimney posterior distribution. Finally, the hill-and-chimney posterior distribution is equal to

$$\frac{\text{prior} \times \text{likelihood}}{\int_{-\infty}^{\infty} \text{prior} \times \text{likelihood} d\delta}. \quad (\text{D10})$$

The numerator of Equation D10 is equal to either Equation D5 (for $\delta \in [-\varepsilon, \varepsilon]$) or Equation D6 (for $\delta \notin [-\varepsilon, \varepsilon]$). The denominator of Equation D10 is given by Equation D7.

The hill-and-chimney prior converges to the spike-and-slab

We first show that the hill part of the hill-and-chimney prior density converges to the slab part of the spike-and-slab prior density as ε converges to 0. Observe that both densities are rescaled versions of the normal density $\mathcal{N}(0, \sigma_0^2)$. What distinguishes both expressions is the rescaling factor, which equals k (Equation D2) for the hill-and-chimney prior and $(1 - \rho_0)$ for the spike-and-slab prior. Thus, all we need is to show is that $\lim_{\varepsilon \rightarrow 0} k = 1 - \rho_0$. And indeed,

$$\lim_{\varepsilon \rightarrow 0} k = \lim_{\varepsilon \rightarrow 0} \frac{1 - \rho_0}{2\Phi\left(-\frac{\varepsilon}{\sigma_0}\right)} = \frac{1 - \rho_0}{2\Phi(0)} = 1 - \rho_0.$$

Moreover, since the probability for the interval $[-\varepsilon, \varepsilon]$ (i.e., $\int_{-\varepsilon}^{\varepsilon} \frac{\rho_0}{2\varepsilon} = \rho_0$) is constant for any value of ε , we conclude that the probability for the interval $[-\varepsilon, \varepsilon]$ converges to the probability mass ρ_0 at $\delta = 0$ as $\varepsilon \rightarrow 0$. That is, in the limit, $P(\delta = 0) = \rho_0$, which is the expression for the spike (Equation 1). This concludes the argument.