A New Bayesian Two-Sample t Test and Solution to the Behrens–Fisher Problem Based on Gaussian Mixture Modelling with Known Allocations

Riko Kelter¹

Received: 27 March 2020 / Revised: 27 September 2021 / Accepted: 6 October 2021 /
Published online: 10 December 2021
© The Author(s) 2021

Abstract

Testing differences between a treatment and control group is common practice in biomedical research like randomized controlled trials (RCT). The standard two-sample t test relies on null hypothesis significance testing (NHST) via p values, which has several drawbacks. Bayesian alternatives were recently introduced using the Bayes factor, which has its own limitations. This paper introduces an alternative to current Bayesian two-sample t tests by interpreting the underlying model as a two-component Gaussian mixture in which the effect size is the quantity of interest, which is most relevant in clinical research. Unlike p values or the Bayes factor, the proposed method focusses on estimation under uncertainty instead of explicit hypothesis testing. Therefore, via a Gibbs sampler, the posterior of the effect size is produced, which is used subsequently for either estimation under uncertainty or explicit hypothesis testing based on the region of practical equivalence (ROPE). An illustrative example, theoretical results and a simulation study show the usefulness of the proposed method, and the test is made available in the R package `bayest`. In sum, the new Bayesian two-sample t test provides a solution to the Behrens–Fisher problem based on Gaussian mixture modelling.

Keywords Bayesian t test · Reproducibility in medical research · Region of practical equivalence (ROPE) · Behrens–Fisher-problem

1 Introduction

In medical research, the t test is one of the most popular statistical procedures conducted. In randomized controlled trials (RCT), the goal often is to test the efficacy of new treatments or drugs and find out the size of an effect. Usually, a

✉ Riko Kelter
riko.kelter@uni-siegen.de

¹ University of Siegen, Walter-Flex-Strasse 3, 57072 Siegen, Germany

treatment and control group are used, and differences in a response variable like blood pressure or cholesterol level between both groups are observed. The gold standard for deciding if the new treatment or drug is more effective than the status quo treatment or drug is the p value, which is the probability, under the null hypothesis H_0 , of obtaining a difference equal to or more extreme than what was actually observed. The dominance of p values when comparing two groups in medical (and other) research is overwhelming [19, 20, 28].

The original two-sample t test belongs to the class of frequentist solutions. These are based on sampling statistics, which allow to reject the null hypothesis via the use of p values. The misuse and drawbacks of p values in medical research have been detailed in a variety of papers, including an official ASA statement in 2016 [38]. On the other side, Bayesian versions of the two-sample t test have become more popular recently. Examples include the proposals in [12, 32, 37, 40] and [14]. All of these focus on the Bayes factor (BF) for testing a null hypothesis $H_0 : \delta = 0$ of no effect against a one- or two-sided alternative $H_1 : \delta > 0$, $H_1 : \delta < 0$ or $H_1 : \delta \neq 0$. Bayes factors themselves are also not without problems: (1) Bayes factors are sensible to prior modelling [18]. (2) Bayes factors require the researcher to calculate marginal likelihoods, the calculations of which can be complex except when conjugate distributions exist. (3) In the setting of the two-sample t test, Bayes factors weigh the evidence for $H_0 : \delta = 0$ against the evidence for $H_1 : \delta \neq 0$ (or $H_1 : \delta < 0$ or $H_1 : \delta > 0$) given the data x . In the case, when $BF_{10} = 20$, H_1 is 20 times more likely after observing the data than H_0 . The natural question following in such cases is: How large is δ ? A Bayes factor cannot answer this question and was not designed to answer such questions, but often this is of most relevance in applied biomedical research. Last, in most applied research, estimation of the effect size δ is more desirable than a mere rejection or acceptance of a point or composite hypothesis [19, 20, 24].

To be fair enough, Bayes factors can be computed alongside posterior estimates, so testing and estimation do not mutually exclude each other. However, as the Bayes factor is often proposed as a replacement for the p value, it is even more questionable if practitioners will really combine testing with estimation of effect sizes, especially when scales translating the size of a Bayes factor into evidence (similar to $p < .05$, $p < .01$) are regularly provided by now, see [35] and [2]. p values and the Bayes factor are useful tools if explicit hypothesis testing is necessary. However, exactly this necessity may be questioned in a wide range of applied biomedical research.

Therefore, this paper proposes an alternative Bayesian two-sample t test by formulating the statistical model as a two-component Gaussian mixture with known allocations and using the region of practical equivalence (ROPE). Instead of focussing on rejection or confirmation of hypotheses, the proposed method's focus lies on estimation of the effect size under uncertainty. Still, if desired the method also provides a Bayesian solution to the Behrens–Fisher problem by enabling to test the equality $\mu_1 = \mu_2$ of two groups of normally distributed data $\mathcal{N}(\mu_1, \sigma_1^2)$ and $\mathcal{N}(\mu_2, \sigma_2^2)$ without assuming $\sigma_1^2 = \sigma_2^2$.

2 Method

2.1 Modelling the Bayesian t Test as a Mixture Model with Known Allocations

In this section, the two-sample t test is modelled as a two-component Gaussian mixture with known allocations.

“Consider a population made up of K subgroups, mixed at random in proportion to the relative group sizes $\eta_1, \dots, \eta_K$. Assume interest lies in some random feature Y which is heterogeneous across and homogeneous within the subgroups. Due to heterogeneity, Y has a different probability distribution in each group, usually assumed to arise from the same parametric family $p(y|\theta)$ however, with the parameter θ differing across the groups. The groups may be labeled through a discrete indicator variable S taking values in the set $\{1, \dots, K\}$.

When sampling randomly from such a population, we may record not only Y , but also the group indicator S . The probability of sampling from the group labeled S is equal to η_S , whereas conditional on knowing S , Y is a random variable following the distribution $p(y|\theta_S)$ with θ_S being the parameter in group S . (...) The marginal density $p(y)$ is obviously given by the following mixture density

$$p(y) = \sum_{S=1}^K p(y, S) = \eta_1 p(y|\theta_1) + \dots + \eta_K p(y|\theta_K)$$

” [9, p. 1]

Clearly, this resembles the situation of the two-sample t test, in which the allocations S are known. While traditionally mixtures are treated with missing allocations, in the setting of the two-sample t test, these are known, leading to a “degenerate” mixture¹. While this assumption does not only remove computational difficulties like label switching, it also makes sense from a semantic perspective: the inherent assumption of a researcher is that the population is indeed made up of $K = 2$ subgroups, which differ in a random feature Y which is heterogeneous across groups and homogeneous within each group. The group indicator S of course is recorded. When conducting a randomized controlled trial (RCT), the clinician will choose the patients according to a sampling plan, which could be set to achieve equally sized groups, that is, $\eta_1 = \eta_2$. Therefore, when sampling the population with the goal of equally sized groups, the researcher randomizes the patients with equal probability from the population into treatment and control group. After the RCT is conducted, the resulting histogram of observed Y values will take the form of the mixture density $p(y)$ above

¹ The mixture is called degenerate here, because when allocations are known, the likelihood is not mixed in the classical sense.

and express bimodality due to the mixture model of the data-generating process.² After fixing the mixture weights, the family of distributions for the single groups needs to be chosen. The above considerations lead to consider finite mixtures of normal distributions, as these “*occur frequently in many areas of applied statistics such as [...] medicine*” [9, p. 169]. The components $p(y|\theta_k)$ become $f_N(y; \mu_k, \sigma_k^2)$ for $k = 1, 2$ in this case, where $f_N(y; \mu_k, \sigma_k^2)$ is the density of the univariate normal distribution. Parameter estimation in finite mixtures of normal distributions consists of estimation of the component parameters (μ_k, σ_k^2) , the allocations $S_i, i = 1, \dots, n$ and the weight distribution (η_1, η_2) based on the available data $y_i, i = 1, \dots, n$. In the case of the two-sample Bayesian t test, the allocations S_i (where $S_i = 1$ if y_i belongs to the first component and $S_i = 2$ else) are known for all observations $y_i, i = 1, \dots, n$. Also, the weights η_1, η_2 are known. Therefore, inference is concerned with the component parameters μ_k, σ_k^2 for $k = 1, 2$ given the complete data (S, y) .

Definition 1 (*Bayesian two-sample t test model*) Let S, Y be random variables with S taking values in the set $\{1, 2\}$ and Y in $\mathbb{R}$. If $Y|S = k \sim \mathcal{N}(\mu_k, \sigma_k^2)$ for $k = 1, 2$, so conditional on S the component densities of Y are Gaussian with unknown parameters μ_k and σ_k^2 , and if the marginal density is a two-component Gaussian mixture with known allocations:

$$p(y) = \eta_1 f_N(y; \mu_1, \sigma_1^2) + \eta_2 f_N(y; \mu_2, \sigma_2^2),$$

where $\eta_2 := \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{S_i=2}(y_i, S_i)$ and $\eta_1 = 1 - \eta_2$, the complete data (S, Y) follow the Bayesian two-sample t test model.

2.2 Inference via Gibbs Sampling

From the above line of thought, it is clear that due to the representation via a mixture model with known allocations, no prior is placed directly on the effect size $\delta := \frac{\mu_1 - \mu_2}{s}$ itself, where

$$s := \sqrt{\frac{(n_1 - 1)s_1^2 + (n_2 - 1)s_2^2}{n_1 + n_2 - 2}}$$

and s_1^2 and s_2^2 are the empirical variances of the two groups, see also Cohen [5]. This is the common approach in existing Bayesian t tests [14]. Instead, in the proposed mixture model, priors are assigned to the parameters of the Gaussian mixture components μ_1, μ_2 and σ_1^2, σ_2^2 . This has several benefits: Incorporation of available prior knowledge is easier achieved with the mixture component parameters than for the effect size, which is an aggregate of these component parameters. Consider a drug

² If unbalanced groups are the goal, the weights could be adjusted accordingly. As in most cases, equally sized groups are considered, $\eta_1 = \eta_2 = 0.5$ is a justified assumption regarding the sampling process in the study or experiment conducted, when balanced groups are used.

where from biochemical properties, it can safely be assumed that the mean in the treatment group will become larger, but the variance will increase, too. Incorporating such knowledge on μ_k and σ_k is much easier than incorporating it in the prior of the effect size δ . This situation holds in particular, when group sizes n_1, n_2 are not balanced.

These practical gains of translating prior knowledge into prior parameters come at a cost: In contrast to existing solutions [14], the model implies that no closed-form expression for the posterior of δ is available. Therefore, sampling methods are used here, to first construct the joint posterior $p(\mu_1, \mu_2, \sigma_1, \sigma_2 | S, y)$ and subsequently use a sample

$$(\mu_1^{(1)}, \mu_2^{(1)}, \sigma_1^{(1)}, \sigma_2^{(1)}, \dots, \mu_1^{(m)}, \mu_2^{(m)}, \sigma_1^{(m)}, \sigma_2^{(m)})$$

of size m , to produce a sample $(\delta^{(1)}, \delta^{(2)}, \dots, \delta^{(m)})$ of δ , where $\delta^{(i)} := \frac{\mu_1^{(i)} - \mu_2^{(i)}}{s^{(i)}}$ and

$$s^{(i)} = \sqrt{\frac{(n_1 - 1)(s_1^{(i)})^2 + (n_2 - 1)(s_2^{(i)})^2}{n_1 + n_2 - 2}}.$$

In summary, via Gibbs sampling, the posterior of δ can be approximated reasonably well. In order to apply Gibbs sampling, the conditional distributions need to be derived.

2.3 Derivation of the Full Conditionals Using the Independence Prior

To derive the full conditionals, it first has to be decided which priors should be used on the mixture component parameters. There are multiple priors available, the most prominent among them the conditionally conjugate prior and the independence prior [7, 9]. While the conditionally conjugate prior has the advantage of leading to a closed-form posterior $p(\mu, \sigma^2 | S, y)$, the main difficulty in the setting of the Bayesian two-sample t test is that while a priori the component parameters $\theta_k = (\mu_k, \sigma_k^2)$ are pairwise independent across both groups, inside each group the mean μ_k and variance σ_k^2 are dependent. This is in contrast to the assumption in the setting of the Bayesian two-sample t test, and therefore, the independence prior is chosen, which is used in [7] and [30]. The independence prior assumes the mean μ_k and the variance σ_k^2 are a priori independent, which is $p(\mu, \sigma^2) = \prod_{k=1}^2 p(\mu_k) \prod_{k=1}^2 p(\sigma_k^2)$, with $\mu_k \sim \mathcal{N}(b_0, B_0)$ and $\sigma_k^2 \sim \mathcal{G}^{-1}(c_0, C_0)$, where $\mathcal{G}^{-1}(\cdot)$ is the inverse Gamma distribution. The normal prior on the means μ_k seems reasonable as the parameters b_0 and B_0 can be chosen to keep the influence of the prior only weakly informative.³ The inverse Gamma prior is chosen because for a two-component Gaussian mixture to show any signs of bimodality – in which case, one would assume differences between two subgroups in the whole sample – the variance should not be huge, because otherwise, the modes (or the bell-shape) of the two normal components of

³ Another option would be a t_n prior, but this would also imply another free hyperparameter to be estimated simultaneously, the degrees of freedom n .

the mixture will flatten out more and more, until unimodality is reached. Thus, the inverse Gamma prior connects this model aspect by giving more probability mass to smaller values of σ_k^2 , while extremely large values get much less prior probability mass.⁴ The hyperparameters c_0 and C_0 then offer control over this kind of shrinkage on σ_k^2 towards zero. In the simulation study below, the prior sensitivity will also be studied briefly.

The independence prior is, therefore, used and leads to the following full conditionals:

Theorem 1 *For the Bayesian two-sample t-test model, the full conditional distributions under the independence prior*

$$p(\mu, \sigma^2) = \prod_{k=1}^2 p(\mu_k) \prod_{k=1}^2 p(\sigma_k^2)$$

with $\mu_k \sim \mathcal{N}(b_0, B_0)$ and $\sigma_k^2 \sim \mathcal{G}^{-1}(c_0, C_0)$ (where $\mathcal{G}^{-1}(\cdot)$ is the inverse Gamma distribution) are given as:

$$\begin{aligned} p(\mu_1 | \mu_2, \sigma_1^2, \sigma_2^2, S, y) &= p(\mu_1 | \sigma_1^2, S, y) \sim \mathcal{N}(b_1(S), B_1(S)) \\ p(\mu_2 | \mu_1, \sigma_1^2, \sigma_2^2, S, y) &= p(\mu_2 | \sigma_2^2, S, y) \sim \mathcal{N}(b_2(S), B_2(S)) \\ p(\sigma_1^2 | \mu_1, \mu_2, \sigma_2^2, S, y) &= p(\sigma_1^2 | \mu_1, S, y) \sim \mathcal{G}^{-1}(c_1(S), C_1(S)) \\ p(\sigma_2^2 | \mu_1, \mu_2, \sigma_1^2, S, y) &= p(\sigma_2^2 | \mu_2, S, y) \sim \mathcal{G}^{-1}(c_2(S), C_2(S)) \end{aligned}$$

with $B_1(S), b_1(S), B_2(S), b_2(S)$ as defined in Equations (20) and (21), and $c_1(S), c_2(S), C_1(S)$ and $C_2(S)$ as defined in Equations (22) and (23) in Appendix A.2.

A proof is given in Appendix A.2, which builds on the derivations in Appendix A.1. Note that when $\eta_1 \neq \eta_2$, $N_1(S)$ and $N_2(S)$ in the Appendices just need to be changed accordingly. For example, if the first group consists of 30 observations, and the second group of 70, setting $N_1(S) = 30$ and $N_2(S) = 70$ implies $\eta_1 = 0.3$ and $\eta_2 = 0.7$, handling the case of unequal group sizes easily.

2.4 Derivation of the Single-Block Gibbs Sampler

Based on the full conditionals derived in the last section, this section now derives a single-block Gibbs sampler to obtain the joint posterior distribution

$$p(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2 | S, y)$$

given the complete data (S, y) . The resulting Gibbs sampler is given as follows:

⁴ Another option would be an exponential prior with parameter λ , but as the exponential distribution is just a special case of the gamma distribution, the more general gamma prior is selected here.

Table 1 Two conceptual distinctions in the practice of data analysis, replicated from Kruschke and Liddell [24]

	Frequentist	Bayesian
Hypothesis test	p value (null hypothesis significance test)	Bayes factor
Estimation with uncertainty	MLE with CI (The “New Statistics”)	Posterior Distribution with the highest posterior density interval (HPD)

Corollary 1 (Single-block Gibbs sampler for the Bayesian two-sample t test) *The joint posterior distribution*

$$p(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2 | S, y)$$

in the Bayesian two-sample t test model can be simulated under the independence prior as follows:

Conditional on the classification $S = (S_1, \dots, S_N)$:

1. Sample σ_k^2 in each group k , $k = 1, 2$ from an inverse Gamma distribution $\mathcal{G}^{-1}(c_k(S), C_k(S))$
2. Sample μ_k in each group k , $k = 1, 2$, from a normal distribution $\mathcal{N}(b_k(S), B_k(S))$

where $B_k(S)$, $b_k(S)$ and $c_k(S)$, $C_k(S)$ are given by Equations (20), (21), (22) and (23) in the Appendix.

A proof is given in Appendix A.3.

2.5 The Shift from Hypothesis Testing to Estimation Under Uncertainty and the ROPE

As already mentioned, instead of explicit hypothesis testing via p values and Bayes factors, we follow the proposed shift from hypothesis testing to estimation under uncertainty. Cumming [6] proposed such a shift originally from frequentist hypothesis testing to estimation, called the “*New Statistics*”, a process observable in a broad range of scientific fields, see Wasserstein and Lazar [39]. Note that one of the six principles for properly interpreting p values in the 2016 ASA statement stressed that a p value “*does not measure the size of an effect or the importance of a result*”. [38, p. 132]. Cumming [6], therefore, included in his proposal a focus on “*estimation based on effect sizes*” [6, p. 7]. To promote this shift, Kruschke and Liddell [24] offered two conceptual distinctions, which are replicated in Table 1.

While Cumming [6] originally proposed a shift from frequentist hypothesis testing to frequentist estimation under uncertainty, in this paper, it is proposed that this shift can be achieved easier by Bayesian methods. The main reasons are that confidence intervals as quantities for estimation are still “*highly sensitive to the stopping*

and testing intentions”. [24, p. 184], while the Bayesian posterior distributions are not, see also Berger and Wolpert [1, Chapter 4].

2.6 The Proposal of a Region of Practical Equivalence

To facilitate the shift to an estimation-oriented perspective, Kruschke and Liddell [24] advertised the *region of practical equivalence (ROPE)*. As they note: “*ROPE*’s go by different names in the literature, including “interval of clinical equivalence”, “range of equivalence”, “equivalence interval”, “indifference zone”, “smallest effect size of interest,” and “good-enough belt” ...’ [24, p. 185], where these terms come from a wide spectrum of scientific domains, see Carlin and Louis [3], Freedman, Lowe and Macaskill [8], Hobbs and Carlin [16], Lakens [25] and Schuirmann [34]. The uniting idea is to establish a region of practical equivalence around the null value of the hypothesis, which expresses “the range of parameter values that are equivalent to the null value for current practical purposes”. [24, p. 185]. With a caution not to slip back into dichotomic thinking, the following decision rule was proposed by Kruschke and Liddell [24]: Reject the null value, if the 95% highest posterior density interval (HPD) falls entirely outside the ROPE. Accept the null value, if the 95% HPD falls entirely inside the ROPE. In the first case, with more than 95% probability, the parameter value is not inside the ROPE, and therefore not practically equivalent to the null value. A rejection of the null value then seems legitimate. In the second case, the parameter value is inside the ROPE with at least 95% posterior probability, and therefore, practically equivalent to the null value. It seems legitimate to accept the null value. Of course, it would also be possible to accept the null value iff the whole posterior is located inside the ROPE, leading to an even stricter decision rule.

While the idea of a ROPE is intriguing, a limitation should be noted which is that it can only apply in situations where scientific standards of practically equivalent parameter values exist and are widely accepted by researchers. Luckily, this is the case for effect sizes, which have a long tradition of being categorized in biomedical and psychological research, see Cohen [5].

Definition 2 (*ROPE*) The region of practical equivalence (ROPE) R for (or around) a hypothesis $H \subset \Theta$ is a subset of the parameter space Θ with $H \subset R$.

A statistical hypothesis H is now described via a region of practical equivalence R , e.g. $H : \delta = \delta_0$ can be described as $R := [\delta_0 - \varepsilon, \delta_0 + \varepsilon]$ for $\varepsilon > 0$. By definition, any set $R \subset \Theta$ with $H \subset R$ is allowed to describe H and should be selected depending on how precise the measuring process of the experiment or study is assumed to be. Next, we define two options for the ROPE:

Definition 3 (*Correctness*) Let $R \subset \Theta$ a ROPE around a hypothesis $H \subset \Theta$, that is $H \subset R$, where H makes a statement about the unknown model parameter θ . If the true parameter value θ_0 lies in R that is $\theta_0 \in R$, then R is called correct, otherwise incorrect.

A correct ROPE, therefore, contains the true parameter value θ_0 , while an incorrect one does not.

Note that the ROPE is equivalent to what is also known as an interval hypothesis $H_0 : \theta \in [\theta_0 - \Delta_0, \theta_0 + \Delta_0]$ for some $\Delta_0 \in \mathbb{R}$ [21]. The test of an interval hypothesis against its alternative, which structures the parameter space into values which are materially different from θ_0 in the scientific context at hand and values which are not ranges back to Hodges and Lehmann [17], see also Kelter [21, 22]. The ROPE is, thus, primarily a synonym for an interval hypothesis.

2.7 The Proposal for a Shift Towards Estimation Under Uncertainty

The two major drawbacks of the proposal of Kruschke [23] and Kruschke and Liddell [24] are that the ROPE still facilitates hypothesis testing, enforcing a binary decision of rejection or acceptance, while it is also unclear what to do when the 95%-HPD lies partly inside and partly outside the ROPE. Therefore, a different proposal is made in this paper, which is estimation of the *mean probable effect size* (MPE) instead of hypothesis testing. This procedure will be used in the proposed t test afterwards. First, the acceptance or rejection of a hypothesis H can be formalized as follows:

Definition 4 (*α -accepted / α -rejected*) Let θ the unknown parameter (or vector of unknown parameters) in an experiment $E := \{X, \theta, \{f_\theta\}\}$, where the random variable X taking values in $\mathbb{R}$ and having density f_θ for some $\theta \in \Theta$ is observed. Let $f(\theta|x)$ the posterior distribution of θ (under any prior $\pi(\theta)$), and let C_α the corresponding $\alpha\%$ highest posterior density interval of $f(\theta|x)$. Let $R \subset \Theta$ a ROPE around the hypothesis H of interest, which makes a statement about θ . Then, if $C_\alpha \subset R$, the hypothesis H is called α -accepted, else α -rejected. If $\alpha = 1$, then H is simply called accepted, else rejected.

Thus, if C_α lies completely inside the ROPE R and if $\alpha = 1$, the entire posterior probability mass indicates that θ is practically equivalent to the values described by the ROPE R . Thus, H can be accepted. If $\alpha < 1$, the strength of this statement becomes less with decreasing α of course. For example, if H is 0.75 accepted for a given ROPE R , 25% of the posterior indicates that θ may take values different than the ones included in the ROPE R . It is clear that little is gained if the value of α is small or close to zero when speaking of α -acceptance. Therefore, instead of forcing an acceptance or rejection (which only makes sense for substantial values of α), a perspective focussing on continuous estimation is preferred:

Definition 5 (*Posterior mass percentage*) Let $f(\theta|x)$ a posterior for θ and $\delta_{\text{MPE}} := \mathbb{E}[\theta|x]$ the mean posterior effect size (where the expectation $\mathbb{E}$ is taken with respect to the posterior). Let $R_1, \dots, R_m$ be a partition of the support of the posterior $f(\theta|x)$ into different ROPEs corresponding to different hypotheses $H_1, \dots, H_m$, which make statements about the unknown parameter (vector) θ . Without loss of

generality, let R_j the ROPE for which $\delta_{\text{MPE}} \subset R_j$, $j \in \{1, \dots, m\}$. The posterior mass percentage $\text{PMP}_{R_j}(\delta_{\text{MPE}})$ of δ_{MPE} is given as follows:

$$\text{PMP}_{R_j}(\delta_{\text{MPE}}) := \int_{R_j} f(\theta|x) d\theta$$

That is, the percentage of the posterior distribution's probability mass inside the ROPE R_j around δ_{MPE} .

For simplicity of notation, the subscript R_j is omitted whenever it is clear which ROPEs R_j are used for partitioning the support of $f(\theta|x)$. Now, in contrast to strict α -acceptance or α -rejection rules based on the ROPE R_j , we propose to use δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$ together to estimate the effect size δ under uncertainty, and to quantify this uncertainty via $\text{PMP}(\delta_{\text{MPE}})$. If δ_{MPE} is non-zero, the t test found a difference between both groups. The size of this difference is quantified by δ_{MPE} itself. The uncertainty in this statement is quantified by $\text{PMP}(\delta_{\text{MPE}})$. For the developed two-sample t test, we propose the following procedure:

1. For a fixed credible level α , the *effect size range (ESR)* should be reported. That is, which effect sizes δ are assigned positive probability mass by the $\alpha\%$ HPD interval, $0 \leq \alpha \leq 1$. The ESR is a first estimate of credible effect sizes a posteriori.
2. The support of the posterior distribution $f(\delta|S, Y)$ in the Bayesian t test model is partitioned into the standardized ROPEs of the effect size δ of [5], leading to a partition $\mathcal{P}$ of the support as given in the definition of $\text{PMP}(\delta_{\text{MPE}})$.
3. The *mean posterior effect size* δ_{MPE} is calculated as an estimate of the true effect size δ . The surrounding ROPE R_j with $\delta_{\text{MPE}} \subset R_j$ of the partition $\mathcal{P}$ is selected, and the exact percentage inside R_j is reported as the *posterior mass percentage* $\text{PMP}(\delta_{\text{MPE}})$.

The above procedure leads to an estimation of the effect size δ under uncertainty instead of a hypothesis testing perspective. Additionally to the posterior mean δ_{MPE} , the posterior mass percentage $\text{PMP}(\delta_{\text{MPE}})$ gives a *continuous* measure of the trustworthiness of the estimate ranging from 0% to 100% (actually from zero to one, but for better interpretability how much of the posteriors mass is allocated in the ROPE R_j the values zero to 100 per cent will be used in what follows). δ_{MPE} estimates with $\text{PMP}(\delta_{\text{MPE}}) > 0.5$ (or 50%) could be interpreted as decisive but do not need to. $\text{PMP}(\delta_{\text{MPE}})$ can be treated as a continuous measure of support for the effect size estimated by δ_{MPE} . There are multiple advantages of utilizing a ROPE and combining it with δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$, the most important of which may be

Theorem 2 *Let $R_j \subset \Theta$ a ROPE around δ_{MPE} that is $\delta_{\text{MPE}} \subset R_j$. If R_j is correct, then $\text{PMP}(\delta_{\text{MPE}}) \rightarrow 1$ for $n \rightarrow \infty$ almost surely, and if R_j is incorrect, then $\text{PMP}(\delta_{\text{MPE}}) \rightarrow 0$ for $n \rightarrow \infty$ almost surely, except possibly on a set of π -measure zero for any prior π on θ .*

A proof is given in Appendix A.4. Using δ_{MPE} together with $\text{PMP}(\delta_{\text{MPE}})$, therefore, will eventually lead to the correct estimation of δ in the sense that when a correct ROPE is chosen, the posterior mass percentage will converge to one, and if an incorrect ROPE is chosen, the posterior mass percentage will converge to zero. Thus, the procedure indicates whether δ is practically equivalent to the values given by the ROPE or not. If necessary, explicit hypothesis testing can be performed via α -rejection. The advantages compared to p values and Bayes factors which favour an explicit hypothesis testing perspective are

1. As Greenland et al. [13, p. 338] stress with regard to the dichotomy induced by hypothesis testing, “*estimation of the size of effects and the uncertainty surrounding our estimates will be far more important for scientific inference and sound judgement than any such classification*”.
2. In contrast to the Bayes factor (BF), the ROPE and δ_{MPE} have important advantages: they do not encourage the same automatic calculation routines as Bayes factors. For example, Gigerenzer and Marewski [11] warned explicitly against Bayes factors becoming the new p values due to the same automatic calculation routines, and the approach via the ROPE fosters estimation and judging the evidence based on the *continuous* support for δ_{MPE} provided by $\text{PMP}(\delta_{\text{MPE}})$, instead of using thresholds.
3. The ROPE is supported by the following argument, which questions the use of testing point hypothesis like $H_0 : \delta = 0$ (no matter if rejecting or confirming is the goal): In practice, measuring is always done with finite precision (like blood pressure, or the heart rate), and therefore, the goal rarely is to show (or reject) that the effect size δ is exactly *equal* to zero but much more that δ is *negligibly small* to deny the existence of any existing, (clinically) relevant effect. Therefore, invariances like $\delta = 0$ can be interpreted as not existing, at least not exactly, and the search for approximate invariances, as described by a ROPE $R = (-.2, .2)$ around $\delta = 0$ is intellectually (more) compelling. A clinician will be satisfied by the statement that the true effect size is not exactly zero, but with 95% probability negligibly small. This general problem of precise hypothesis testing was first noted by Hodges & Lehmann [17] as mentioned earlier, compare also Rao & Lovric [29] and Kelter [21].

2.8 Illustrative Example

To clarify the above line of thought, the following example combines the developed Gibbs sampler for the Bayesian t test with the ROPE, δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$. It uses data from Wagenmakers et al. [36], who conducted a randomized controlled trial in which participants had to fill out a personality questionnaire while rolling a kitchen roll clockwise or counter clockwise. The mean score of both groups was compared afterwards.

Fig. 1 Posterior distribution of the effect size δ and analysis for the kitchen roll RCT of [36] via δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$

2.8.1 Frequentist Analysis via Welch's Two-Sample t Test

A traditional two-sided two-sample Welch's t test indicates that there is no significant difference between both groups, yielding a p value of 0.4542.⁵ What is missing is the effect size, which is of much more interest. Note that computing the effect size from the raw study data does not quantify the uncertainty in the data, which is undesirable. From the p value, a clinician can only judge that the results are unlikely to be observed under the null hypothesis. However, if the effect is clinically relevant or negligible remains unknown (or at best only point and interval estimates are reported). In this case, as the p value is quite large, neither can the null hypothesis of no effect be rejected, nor can be said with certainty that there is indeed no effect in the sense of confirming the null hypothesis, leaving the clinician without any trace to proceed with but to collect more data.

⁵ The data and code including a full replication script for all simulations and figures presented in the paper are available at <https://osf.io/cvwr5/>.

2.8.2 Bayesian Analysis via the Two-Sample t Test Based on Gaussian Mixtures

Figure 1 in contrast shows an analysis of the posterior of δ produced by the Gibbs sampler for the Bayesian t test. In it, the ROPE, δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$ are used. The posterior distribution of δ is given in the upper plot and shows that the posterior mean is 0.149 and the posterior mode 0.156, that is, the mean posterior effect size given the data is 0.149, no effect discernible from zero. The 95% highest density interval ranges from 0.114 to 0.182, showing that with 95% probability, there is no effect discernible from $\delta = 0$, given the data. Even when taking the 100% highest posterior density interval (HPD), this situation does not change as indicated by the upper plot. The coloured horizontal lines and vertical dotted lines represent the boundaries of the different ROPEs according to Cohen [5]. The lower plot shows the results of partitioning the posterior mass of δ into the ROPEs for different effect sizes, which are standardized as small, if $\delta \in [0.2, 0.5)$, medium, if $\delta \in [0.5, 0.8)$ and large, if $\delta \in [0.8, \infty)$ [5]. 100% of this posterior probability mass lies inside the ROPE $(-0.2, 0.2)$ of no effect. So $\delta_{\text{MPE}} = 0.149$ indicates that no effect discernible from zero is apparent, and the posterior mass percentage $\text{PMP}(\delta_{\text{MPE}}) = 1$ (or 100%) shows that the estimate δ_{MPE} is trustworthy, as the entire posterior probability mass is located inside the ROPE $(-0.2, 0.2)$ of no effect (also indicated by the upper plot). Based on this analysis, one can conclude that given the data, it is highly probable that there exists no effect. The method provides more insight than the information a p value is giving: In the example, the p value cannot reject the null hypothesis $H_0 : \delta = 0$ and neither can the null hypothesis be confirmed. Even if the p value would have been significant, this means only that the result is unlikely to be observed under the null hypothesis. The non-significant p value of 0.4542 in this case does not enable researchers to accept the null hypothesis of no effect. The proposed procedure in contrast does.

2.8.3 Bayes Factor Analysis

A Bayes factor would have to be combined with estimation to yield the same information, and a Bayes factor alone of course would not have provided this information. In the example, the Bayes factor BF_{01} of $H_0 : \delta = 0$ against $H_1 : \delta \neq 0$ is $\text{BF}_{01} = 5.015$ when using the recommended wide Cauchy $C(0, 1)$ prior of Rouder et al. [32], which indicates only moderate evidence for the null hypothesis $H_0 : \delta = 0$ according to Van Doorn et al. [35]. This is in sharp contrast to the $\text{PMP}(\delta_{\text{MPE}})$ value of 100%, which strongly suggests that the null hypothesis $H_0 : \delta = 0$ is confirmed. The posterior in Fig. 1 is obtained by the Gibbs sampler given in Corollary 1.

3 Simulation Study

Primary interest now lies in the ability to correctly estimate different sizes of effects via the combination of the derived t test, δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$. The effect size ROPEs are oriented at the standard effect sizes of Cohen [5], where an effect is categorized as *small*, if $\delta \in [0.2, 0.5)$ or $\delta \in (-0.5, -0.2]$, *medium*, if $\delta \in [0.5, 0.8)$ or

$\delta \in (-0.8, -0.5]$ and *large*, if $\delta \geq 0.8$ or $\delta \leq -0.8$. Secondary interest lies in analysing if the Gibbs sampler achieves better performance regarding the type I and II error compared with Welch's t test, the standard NHST solution.⁶ The plan of the study is as follows: If there is indeed an effect, the Gibbs sampler should lead to a posterior distribution of δ which lies outside the ROPE $(-0.2, 0.2)$, which is equivalent to the rejection of the null hypothesis $H_0 : \delta = 0$. The precise estimation of the size of an effect is a second task, one more demanding than the sole rejection of $H_0 : \delta = 0$. If the sampler correctly rejects the null hypothesis, because the posteriors concentrate in the set $(-\infty, -0.2] \cup [0.2, \infty)$, this indicates that it makes no type II error and subsequently achieves a power of nearly 100%. Of course, this will depend on the sample sizes in both groups. If additionally, the 95%-credible intervals of the posteriors concentrate in the set $\{(-0.5, -0.2] \cup [0.2, 0.5)\}$, then the Gibbs sampler is also consistent for small effect sizes, again depending on the sample size. The same rationale applies for medium effect sizes and the ROPE $\{(-0.8, 0.5] \cup [0.5, 0.8)\}$ and for large effect sizes and the ROPE $\{(-\infty, 0.8] \cup [0.8, \infty)\}$. Therefore, three two-component Gaussian mixtures have been fixed in advance, each representing one of the three effect sizes. For the small effect, the first component is $\mathcal{N}(2.89, 1.84)$ and the second component $\mathcal{N}(3.5, 1.56)$, resulting in a effect size of $\delta = (2.89 - 3.5)/\sqrt{((1.56^2 + 1.84^2)/2)} = -0.35$. For a medium effect, the first and second group are simulated as $\mathcal{N}(254.08, 2.36)$ and $\mathcal{N}(255.84, 3.04)$, yielding a true effect size of

$$\delta = \frac{(255.84 - 254.08)}{\sqrt{((3.04^2 + 2.36^2)/2)}} = 0.6467 \approx 0.65$$

For the large effect, the first and second groups are simulated as $\mathcal{N}(15.01, 3.4)$ and $\mathcal{N}(19.91, 5.8)$, yielding a true effect size of

$$\delta = \frac{(19.91 - 15.01)}{\sqrt{((5.8^2 + 3.4^2)/2)}} = 1.03$$

Note that in each setting $\sigma_1^2 \neq \sigma_2^2$ which is the premise in the Behrens–Fisher problem [33]. In each of the three effect size scenarios, 100 datasets of the corresponding two-component mixture were simulated for different sample sizes, and the Gibbs sampler was run for each of the 100 datasets for 10000 iterations, using a burn-in of 5000. δ_{MPE} , the ESR and the ROPE criterion together with α -acceptance are applied, that is, the hypothesis H stating a small, medium or large effect size is α -accepted if the 95%-HPD lies completely inside the corresponding ROPE $\{(-0.5, -0.2] \cup [0.2, 0.5)\}$, $\{(-0.8, -0.5] \cup [0.5, 0.8)\}$ or $\{(-\infty, -0.8] \cup [0.8, \infty)\}$. The recommended wide prior was used for all simulations, which is detailed later in

⁶ We do not compare the type I error rate with Bayes factors, as these error types are formally not defined for the Bayes factor.

Fig. 2 Posterior Means δ_{MPE} and 95% credible intervals for δ_{MPE} for 100 datasets consisting of sample sizes $2n$, with n observations in each group; dotted lines represent the ROPE boundaries; *upper row*: Small effect size; *middle row*: Medium effect size; *lower row*: Large effect size

the prior sensitivity analysis. In total, the Gibbs sampler should stabilize around the true effect size δ .⁷

3.1 Results

The upper row of Fig. 2 shows the results for small effect sizes. The two left plots show the results for $n = 100$ and $n = 200$ observations in each group. It is clear that the 95%-HPDs in both cases fluctuate strongly, indicating that anything from no effect to a medium effect is possible. The two right plots of the upper row show the results when increasing to $n = 300$ and $n = 700$ observations per group. The 95%-HPDs get narrower and stabilize inside the ROPE. While for $n = 300$, there are still some outliers, for $n = 700$, all HPDs have concentrated inside the ROPE of a small effect – that is $\text{PMP}(\delta_{\text{MPE}}) = 1$ (100%) for all iterations – and the estimates δ_{MPE} (blue points) have already converged closely to the true effect size indicated by the solid black line. The necessary sample size for this precision is not small, but a small effect requires a large sample size to be detected.

⁷ Note, that here balanced groups are treated, but unbalanced groups could also easily be treated by setting $N_1(S)$ and $N_2(S)$ accordingly, as described above.

Fig. 3 Posterior means δ_{MPE} and 95% credible intervals for δ for 100 datasets consisting of different sample sizes $2n$, with n observations from a $\mathcal{N}(148.3, 1.34)$ distribution and n observations from a $\mathcal{N}(148.3, 2.04)$ distribution; dotted lines represent the ROPE $(-0.2, 0.2)$ of no effect size around $\delta = 0$; posterior distributions are based on 10000 iterations of the Gibbs sampler with a burn-in of 5000 iterations

The middle row of Fig. 2 shows the results for medium effect sizes. The two left plots in it show the result for $n = 100$ and $n = 200$ observations in each group for a medium effect size. Increasing the sample size to $n = 400$ and $n = 600$ leads to the results shown in the two right plots. These figures show that even for sample sizes of $n = 100$ in both groups, no 95% HPD lies completely inside the ROPE $(-0.2, 0.2)$ around $\delta_0 = 0$ of no effect, indicating that while the size of the effect may still not be estimated accurately, a null hypothesis of no effect $\delta_0 = 0$ could always be rejected when using sample sizes of at least $n = 100$ in each group and the underlying effect has medium size. When it comes to precisely estimating the size of the effect, larger sample sizes similar to those needed to detect small effect sizes are necessary, as shown by the right plots of the second row.

The lower row of Fig. 2 shows the results for large effect sizes. About $n = 50$ observations in each group suffice to produce δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$ which estimate small to large effects, and thereby reject a null hypothesis of no effect, while about $n = 150$ to $n = 200$ seem reasonable to precisely estimate a large effect size. When using δ_{MPE} (blue points) as an estimator for δ , sample sizes of $n = 200$ produce an estimate close to the true effect size, which in this case was $\delta_0 = 1.030723$.

3.2 Controlling the Type I Error Rate

In frequentist NHST, the Neyman–Pearson theory aims at controlling the type I error rate α , which is the probability to reject the null hypothesis H_0 when it is true. In the setting of the two-sample Bayesian t test, this equals the rejection of $H_0 : \delta = 0$ although the true effect size is $\delta_0 = 0$. Following Cohen [5], an effect is considered small if the effect size is at least $|\delta| \geq 0.2$, so effect sizes in the interval $(-0.2, 0.2)$ can be considered as noise, or *practically equivalent to zero*. Therefore, a ROPE of $(-0.2, 0.2)$ is set around the null value $\delta_0 = 0$ to compare the type I error rate of the proposed method against the standard frequentist NHST solution, Welch’s t

test. Again 100 datasets of different sample sizes are simulated where the true effect size δ_0 is set to zero. The Gibbs sampler should produce a posterior distribution of δ which concentrates inside the ROPE, so that the null hypothesis $H : \delta_0 = 0$ is α -accepted for $\alpha = 0.95$, see Definition 4. If the 95%-HPD interval lies (entirely) outside the ROPE, this equals α -rejection for $\alpha = 0.95$, or in frequentist terms the rejection of the null hypothesis $H_0 : \delta_0 = 0$ of no effect, and therefore, the commitment of a type I error. The following two definitions formalize the type I and II error building on the concept of α -rejection:

Definition 6 (*α type I error*) An α type I error happens if the true parameter value $\delta_0 \in H$, with $H \subset R$ for a ROPE $R \subset \Theta$, but H is α -rejected for α .

Definition 7 (*α type II error*) An α type II error happens if the true parameter value $\delta_0 \notin H$ and $\delta_0 \notin R$, with $H \subset R$ for a ROPE $R \subset \Theta$, but H is α -accepted for α .

The left plot in Fig. 3 shows the results of 100 datasets of size $n = 50$ in each group. The first group was simulated as $\mathcal{N}(148.3, 1.34)$, and the second group as $\mathcal{N}(148.3, 2.03)$. The true effect size is

$$\delta_0 = \frac{\mu_2 - \mu_1}{\sqrt{(\sigma_1^2 + \sigma_2^2)/2}} = 0$$

The blue points represent δ_{MPE} and the blue-dotted lines, the 95%-HPDs of the posterior of δ . While the estimates fluctuate strongly for $n = 50$, increasing sample size in each group successively to $n = 200$ as shown by the progression of the plots from left to right shows that false-positive results – α type I errors with $\alpha = .95$ – get completely eliminated for sufficiently large sample size. The right plot with sample size $n = 300$ shows that no α type I error with $\alpha = .95$ occurs anymore. Also, δ_{MPE} stabilizes around the true value $\delta_0 = 0$ of no effect, indicating its convergence to the true effect size δ .

The simulations show that the α type I error rate converges to zero when the sample size is increased. The number of credible intervals which lie partly inside and partly outside the ROPE decreases to zero. In contrast, p values are uniformly distributed under the null hypothesis, so that no matter what size the samples in both groups are, in the long-run one will still obtain $\alpha\%$ (most often 5%) type I errors. Conducting Welch's t tests will, thus, inevitably lead to a type I error rate of 5%, if the test level is set to $\alpha = .05$. If the sample size is at least $n = 200$ in each group, the proposed Bayesian t test together with δ_{MPE} and $\text{PMP}(\delta_{\text{MPE}})$ performs better with respect to control the α type I error rate.

From a theoretical perspective, it is of course of interest for which values of α this fact does hold, and indeed, using the two generalized types of type I and II errors, it can also be shown that the number of type I (type II) errors converges to zero for any $\alpha \neq 0$, when a correct (incorrect) ROPE is chosen:

Fig. 4 Prior sensitivity analysis for the $\mathcal{N}(b_0, B_0)$ prior on the means μ_k and inverse Gamma prior $\mathcal{G}^{-1}(c_0, C_0)$ on the variances σ_k^2 , $k = 1, 2$ for 100 datasets with first group simulated as $\mathcal{N}(0, 1)$ and the second as $\mathcal{N}(1, 1)$

Theorem 3 *For the Bayesian two-sample t test model, the probability of making an α type I error for any $\alpha \neq 0$ converges to zero for any correct ROPE R around the hypothesis H which makes a statement about the unknown parameter δ . Also, the probability of making a α type II error for any $\alpha \neq 0$ converges to zero for any incorrect ROPE R .*

A proof is given in Appendix A.5. The implications of Theorem 3 are that if a correct ROPE R is chosen, then eventually the probability of making a α type I error will become zero. Thus, when the ROPE R includes the true parameter δ_0 , eventually the hypothesis H will be α -accepted for $\alpha = 1$, that is, accepted. If on the other hand, an incorrect ROPE is selected, which does not include the true parameter δ_0 , then eventually the probability of making an α type II error – that is, accepting H although $\delta_0 \notin H$ – will become zero for any $\alpha \neq 0$ if sample size is large enough.

3.3 Prior Sensitivity Analysis

Section 2.3 detailed the independence prior used in the model and of specific interest is of course the influence of this prior on the results produced by the procedure. Therefore, three different hyperparameter settings were selected to resemble a wide, medium and narrow prior, where the shrinkage effect on the σ_k^2 , $k = 1, 2$ caused

by the inverse Gamma prior $\mathcal{G}^{-1}(c_0, C_0)$ on σ_k^2 increases with the prior getting narrower (that is, σ_k^2 is shrunk towards zero). The same applies for the normal prior $\mathcal{N}(b_0, B_0)$ on the means μ_k , $k = 1, 2$. The following hyperparameters were chosen for the three different settings: For the wide prior, $b_0 := \bar{x}$ and $B_0 := 10 \cdot s^2(x)$ where $\bar{x}$ and $s^2(x)$ are the complete sample mean and variance. c_0 and C_0 were selected both as 0.01 for the wide prior, implying fatter tails of the inverse Gamma prior than in the medium or narrow prior. For the medium prior, B_0 was decreased to $5 \cdot s^2(x)$, and c_0 and C_0 increased to 0.1 both. For the narrow prior finally, $B_0 := s^2(x)$ and $c_0 = C_0 = 1$, which is the most informative of all three priors.

Subsequently, 100 datasets with $n = 100$ observations in each group were simulated, where the first group was generated as $\mathcal{N}(0, 1)$ and the second as $\mathcal{N}(1, 1)$. The Gibbs sampler was run for 10000 iterations with a burn-in of 5000 once for each prior on each dataset. Fig. 4 shows the results of the simulations. Here, the resulting posterior densities of δ_{MPE} are overlayed in Fig. 4, and it becomes clear that the wide and medium prior do result in barely differing posteriors. When using the narrow prior the shrinkage moves the posterior slightly towards smaller values of δ .

The lower plot in Fig. 4 additionally shows the posterior distributions of differences between means obtained from the three priors: The left-hand plot shows the posterior distribution of differences between δ_{MPE} obtained via a wide and a medium prior. The middle plot shows the posterior distribution of differences between δ_{MPE} obtained via a wide and a narrow prior, and the right-hand plot the posterior distribution of differences between δ_{MPE} obtained via a medium and a narrow prior. The results show that in all cases the differences are of tiny magnitude, indicating that the proposed t test is quite robust to the prior hyperparameters selected. Of course, it can happen that δ_{MPE} will be drawn towards smaller values when switching from the wide to the narrow prior, but the posterior mass percentage supporting a large effect will not vary much as shown by the nearly identical resulting posterior densities in Fig. 4, which is a strength of the continuous quantification through $\text{PMP}(\delta_{\text{MPE}})$. Based on the sensitivity analysis, all three hyperparameter settings differ only slightly, and therefore, the wide prior seems suitable for most applications, as it is the least informative one and can be interpreted as most objective because it introduces the smallest amount of subjectivity into the analysis.

4 Discussion

In this paper, a new Bayesian two-sample t test based on Gaussian mixture modelling with known allocations was introduced which provides a Bayesian solution to the Behrens–Fisher problem, while simultaneously focussing on effect size estimation. Following the proposal of a shift from hypothesis testing to estimation under uncertainty, a Bayesian two-sample t test was derived for inference on the effect size δ , which is the quantity of interest in most biomedical research. Also, the dichotomy of the ROPE decision rule of Kruschke and Liddell [24] was resolved by introducing the mean probable effect size δ_{MPE} as an estimator of δ , combined with the posterior

mass percentage $\text{PMP}(\delta_{\text{MPE}})$, a continuous measure which quantifies the support for the evidence suggested by δ_{MPE} .

Theoretical results showed that the use of the proposed method leads to a consistent estimation procedure which shifts from hypothesis testing to estimation under uncertainty. Also, theoretical results showed that under any correct ROPE R , the number of introduced α type I errors converges to zero a.s. under the law determined by the prior π on δ , while under any incorrect ROPE (which means the ROPE picked by the researcher does not cover the true parameter value δ_0), the number of introduced α type II errors converges to zero a.s. under the law determined by the prior π on δ . Together, these properties and the introduced concept of α -rejection make the proposed method an attractive alternative to existing solutions via p values or the Bayes factor.

One important limitation of the approach is that the results depend on the chosen priors for μ_k and σ_k^2 . Despite being quite robust, especially the choice of the inverse Gamma prior for the variances may be questioned. While it is beyond the scope of this paper to perform a sensitivity analysis using different priors for the variances, one remedy which allows changing the priors would be to switch to different sampling techniques, for example Hamiltonian Monte Carlo in Stan [4]. Also, the speed of convergence to entire elimination of type I errors may be slow, although the simulation results are promising.

The proposed method, therefore, may be helpful in improving the reproducibility in biomedical research, especially by reducing the number of false-positive results, which is one of the biggest problems of the medical sciences, see McElreath and Smaldino [27]. Finally, it should be noted that both Bayes factors and p values are reasonable tools when hypothesis testing is the goal [19]. The proposed new Bayesian two-sample t test is not intended to be a replacement of these tools, but as a complementary tool for cases in which the test of a *precise* point null hypothesis $H_0 : \theta = \theta_0$ for some $\theta_0 \in \Theta$ is not desired but effect size estimation under uncertainty is more important. Such situations are, in particular, important in the medical sciences [17, 21, 26]. Therefore, the main advantage may be seen in the shift towards effect size estimation. Future work could extend the model to more than two groups, leading to an equivalent of the ANOVA, use more robust component distributions like t -distributions, or derive the posterior under different priors on the mixture components parameters.

Appendix

A.1 Derivation of the Single-Block Gibbs Sampler

This supplementary material provides the derivation of the joint posterior and full conditionals for the single-block Gibbs sampler.

Bayesian Parameter Estimation for Known Allocations

In this section, for the setting when the Allocations S are known the posterior distribution of μ_k, σ_k^2 given the complete data (S, y) are derived, see also [9]. As already mentioned above, the weights (η_1, η_2) are known in this case because for every observation $y_i \in (y_1, \dots, y_N)$ it is known to which group y_i belongs, that is, the quantities $S_i = k, k \in \{1, 2\}$ are available for all $i \in \{1, \dots, N\}$.

In the setting of a t test between two groups with equal sample sizes $n_1 = n_2$ and $n_1 + n_2 = N$, the belonging of an observation y_i to its group normally is known for all observations $i \in \{1, \dots, N\}$. The underlying data-generating process, therefore, can be assumed to consist of a mixture of $K = 2$ components with weights $\eta_1 = \eta_2 = 0.5$. This makes inference in the mixture model much easier compared to the case when both the weights (η_1, η_2) as well as the component parameters (μ_1, μ_2) and (σ_1^2, σ_2^2) are unknown.

To conduct inference about the unknown parameters, the necessary group-specific quantities are the number $N_k(S)$ of observations in group k for $k = 1, 2$, the within-group variance $s_{y,k}^2(S)$ and the group mean $\bar{y}_k(S)$:

$$\begin{aligned} N_k(S) &= |\{i : S_i = k\}| \\ \bar{y}_k &= \frac{1}{N_k(S)} \sum_{i:S_i=k} y_i \\ s_{y,k}^2(S) &= \frac{1}{N_k(S)} \sum_{i:S_i=k} (y_i - \bar{y}_k(S))^2 \end{aligned} \quad ,$$

where $|\cdot|$ denotes the cardinality of a set. These quantities depend on S , so the classification of the observation y_i to the component $S_i = k$ needs to be available. When $S_i = k$ for an observation y_i holds, then the observational model for observation y_i is $\mathcal{N}(\mu_k, \sigma_k^2)$ and y_i contributes to the complete-data likelihood $p(y|\mu, \sigma^2, S)$ by a factor of

$$\frac{1}{\sqrt{2\pi\sigma_k^2}} \exp\left(-\frac{1}{2\sigma_k^2}(y_i - \mu_k)^2\right).$$

Taking into account all observations $y_1, \dots, y_N$, the complete-data likelihood function can be written as follows:

$$p(y|\mu, \sigma^2, S) = \prod_{k=1}^2 \left(\frac{1}{2\pi\sigma_k^2} \right)^{N_k(S)/2} \exp\left(-\frac{1}{2} \sum_{i:S_i=k} \frac{(y_i - \mu_k)^2}{\sigma_k^2}\right).$$

The complete-data likelihood is a product of $K = 2$ components, of which each summarizes the information about the i th group, $i \in \{1, 2\}$. These $K = 2$ factors are then combined in a Bayesian analysis with a prior. While the ultimate interest lies in the posterior of both μ_k, σ_k^2 , we consider first two different cases, which will eventually lead to the solution of the joint posterior for μ_k and σ_k^2 .

In the first case, when the variance σ_k^2 is fixed, the complete-data likelihood function as a function of μ is the kernel of a univariate normal distribution. Choosing a $\mathcal{N}(b_0, B_0)$ -distribution as a conjugate prior, the posterior density of μ_k given σ_k^2 and the $N_k(S)$ observations in group k for $k = 1, 2$ can be derived as follows:

$$p(\mu_k | \sigma_k^2, S, y) \propto p(y | \mu_k, \sigma_k^2, S) \cdot p(\mu_k, \sigma_k^2, S) \quad (1)$$

$$\stackrel{(1)}{=} p(y | \mu_k, \sigma_k^2, S) \cdot p(\mu_k) \quad (2)$$

$$= \left(\frac{1}{2\pi\sigma_k^2} \right)^{N_k(S)/2} \exp \left(-\frac{1}{2} \sum_{i:S_i=k} \frac{(y_i - \mu_k)^2}{\sigma_k^2} \right) \quad (3)$$

$$\cdot \frac{1}{2\pi B_0} \exp \left(-\frac{1}{2} \frac{(\mu_k - b_0)^2}{B_0} \right) \quad (4)$$

where in (1) the fact that σ_k^2 is assumed to be given and the allocations S are known constants, too, was used.

In general, for a sample of size n from a $\mathcal{N}(\mu, \sigma)$ distribution with known variance σ^2 , a standard Bayesian analysis, see e.g. [15, p. 181], yields that the likelihood

$$L(\mu) \propto \exp \left(-\frac{n}{2\sigma^2} (\mu - \bar{x})^2 \right)$$

when combined with a prior $\mu \sim \mathcal{N}(\nu, \tau^2)$ leads to the posterior

$$\mu | x \sim \mathcal{N} \left(\left(\frac{n}{\sigma^2} + \frac{1}{\tau^2} \right)^{-1} \cdot \left(\frac{n\bar{x}}{\sigma^2} + \frac{\nu}{\tau^2} \right), \left(\frac{n}{\sigma^2} + \frac{1}{\tau^2} \right)^{-1} \right) \quad (5)$$

Substituting $\nu = b_0$ and $\tau^2 = B_0$ for the prior of μ as well as μ_k for μ and σ_k^2 for σ^2 in the likelihood, the posterior $p(\mu_k | \sigma_k^2, S, y)$ in Eq. (1) becomes

$$p(\mu_k | \sigma_k^2, S, y) \sim \mathcal{N} \left[\left(\frac{N_k(S)}{\sigma_k^2} + \frac{1}{B_0} \right)^{-1} \cdot \left(\frac{N_k(S)\bar{y}_k(S)}{\sigma_k^2} + \frac{b_0}{B_0} \right), \left(\frac{N_k(S)}{\sigma_k^2} + \frac{1}{B_0} \right)^{-1} \right] \quad (6)$$

By Eq. (6), the posterior can be written as follows:

$$\mu_k | \sigma_k^2, S, y \sim \mathcal{N}(b_k(S), B_k(S))$$

with

$$B_k(S)^{-1} = B_0^{-1} + \sigma_k^{-2} N_k(S) \quad , \quad (7)$$

$$b_k(S) = B_k(S)(\sigma_k^{-2} N_k(S)\bar{y}_k(S) + B_0^{-1} b_0) \quad , \quad (8)$$

where for an empty group k for $k = 1, 2$ the term $N_k(S)\bar{y}_k(S)$ is defined as zero.

On the other hand, if the mean μ_k is regarded as fixed, the complete-data likelihood as a function of σ_k^2 is the kernel of an inverse Gamma density. Choosing the conjugate inverse Gamma prior $\sigma_k^2 \sim \mathcal{G}^{-1}(c_0, C_0)$, a standard Bayesian analysis – for details, see for example Held and Sabanés-Bové [15, p. 181] – yields the posterior of $\sigma_k^2 | \mu_k, S, y$ as

$$p(\sigma_k^2 | \mu_k, S, y) \sim \mathcal{G}^{-1}(c_k(S), C_k(S)) \quad (9)$$

with

$$c_k(S) = c_0 + \frac{1}{2}N_k(S) \quad (10)$$

$$C_k(S) = C_0 + \frac{1}{2} \sum_{i:S_i=k} (y_i - \mu_k)^2 \quad (11)$$

The case of interest here is when both μ_k and σ_k^2 are unknown, and in this case, a closed-form solution for the joint posterior $p(\mu_k, \sigma_k^2 | S, y)$ does exist only under specific conditions. That is, the prior variance of the mean of group k , μ_k , must depend on σ_k^2 through the relation $B_{0,k} = \frac{\sigma_k^2}{N_0}$, where N_0 is a newly introduced hyperparameter in the prior of μ_k , that is, the prior $\mu_k \sim \mathcal{N}(b_0, B_0)$ then becomes $\mu_k \sim \mathcal{N}(b_0, \sigma_k^2/N_0)$. The joint posterior now can be rewritten as follows:

$$\begin{aligned} p(\mu, \sigma^2 | S, y) &= p(\mu_1, \dots, \mu_K, \sigma_1^2, \dots, \sigma_K^2 | S, y) \\ &\stackrel{(1)}{=} \prod_{k=1}^2 p(\mu_k, \sigma_k^2 | S, y) \\ &\stackrel{(2)}{=} \prod_{k=1}^2 \underbrace{p(\mu_k | \sigma_k^2, S, y)}_{=:(A)} \cdot \underbrace{p(\sigma_k^2 | S, y)}_{=:(B)} \end{aligned} \quad (12)$$

where (1) follows from the fact that the group parameters μ_k, σ_k^2 are assumed to be independent across groups and (2) follows from factorizing the joint posterior as follows:

$$\begin{aligned} p(\mu_k, \sigma_k^2 | S, y) &= \frac{p(\mu_k, \sigma_k^2 | S, y) \cdot p(\sigma_k^2 | S, y)}{p(\sigma_k^2 | S, y)} \\ &= p(\mu_k | \sigma_k^2, S, y) \cdot p(\sigma_k^2 | S, y) \end{aligned}$$

As the factor (A) in Eq. (12) was already derived in Eq. (6) with corresponding parameters in Eqs. (7) and (8) for $k = 1, 2$, the factor (A) of the posterior in Eq. (12) is normal distributed $\mathcal{N}(b_k(S), B_k(S))$ with parameters

$$B_k(S) \stackrel{(1)}{=} \frac{1}{B_0^{-1} + \sigma_k^{-2} N_k(S)} \stackrel{(2)}{=} \frac{1}{\sigma_k^{-2} N_0 + \sigma_k^{-2} N_k(S)} \quad (13)$$

$$\frac{1}{N_0 + N_k(S)} \sigma_k^2 \quad (14)$$

and

$$b_k(S) \stackrel{(3)}{=} B_k(S) (\sigma_k^{-2} N_k(S) \bar{y}_k(S) + B_0^{-1} b_0) \quad , \quad (15)$$

$$\stackrel{(4)}{=} B_k(S) (\sigma_k^{-2} N_k(S) \bar{y}_k(S) + \frac{N_0}{\sigma_k^2} b_0) \quad , \quad (16)$$

$$\stackrel{(5)}{=} \frac{1}{N_0 + N_k(S)} \sigma_k^2 (\sigma_k^{-2} N_k(S) \bar{y}_k(S) + \frac{N_0}{\sigma_k^2} b_0) \quad , \quad (17)$$

$$= \frac{N_k(S) \bar{y}_k(S) + N_0 b_0}{N_0 + N_k(S)} \quad , \quad (18)$$

$$= \frac{N_0}{N_k(S) + N_0} b_0 + \frac{N_k(S)}{N_k(S) + N_0} \bar{y}_k(S) \quad , \quad (19)$$

where in (1) $B_k(S)^{-1} = B_0^{-1} + \sigma_k^{-2} N_k(S)$ from Eq. (7) was used and in (2) the relation $B_{0,k} = \frac{\sigma_k^2}{N_0}$, where N_0 is the newly introduced hyperparameter. In (3), Equation (8) was used, in (4) again the relation $B_{0,k} = \frac{\sigma_k^2}{N_0}$, in (5) the right-hand side of Eq. (13) was substituted for $B_k(S)$ in Eq. (16).

The remaining term (B) of Eq. (12) is the marginal posterior of σ_k^2 that is

$$\int p(\sigma_k^2 | \mu_k, S, y) d\mu_k$$

and by integrating out μ_k , a standard Bayesian analysis shows that the marginal posterior of σ_k^2 is distributed as inverse Gamma $\mathcal{G}^{-1}(c_k(S), C_k(S))$, where $c_k(S)$ is already given in Eq. (10), and the parameter $C_k(S)$ in Eq. (11) changes to

$$C_k(S) = C_0 + \frac{1}{2} \left(N_k(S) s_{y,k}^2(S) + \frac{N_k(S) N_0}{N_k(S) + N_0} (\bar{y}_k(S) - b_0)^2 \right) \quad .$$

This is because combining an inverse-gamma prior with the normal likelihood with known mean yields an inverse-gamma posterior as shown above and marginalizing this posterior for the variance yields exactly another inverse-gamma distribution with different parameters. Details can be found in [15].

Application to the Two-Sample t Test: Derivation of the Marginal and Joint Posterior Distributions

Summarizing, for two groups, the two-component Gaussian mixture can be interpreted as a data-generating process which consists of $K = 2$ components. The weights η_1 and η_2 are both equal to $1/2$ for equally sized groups, that is, $N = n_1 + n_2$ with n_1 being the sample size of group one and n_2 the sample size of group two and $n_1 = n_2$.

Taking into account all observations $y_1, \dots, y_N$, the complete-data likelihood function can be written as follows:

$$\begin{aligned} p(y|\mu, \sigma^2, S) &= \prod_{k=1}^2 \left(\frac{1}{2\pi\sigma_k^2} \right)^{N_k(S)/2} \cdot \exp \left(-\frac{1}{2} \sum_{i:S_i=k} \frac{(y_i - \mu_k)^2}{\sigma_k^2} \right) \\ &= \left(\frac{1}{2\pi\sigma_1^2} \right)^{N_1(S)/2} \cdot \exp \left(-\frac{1}{2} \sum_{i:S_i=1} \frac{(y_i - \mu_1)^2}{\sigma_1^2} \right) \\ &\quad \cdot \left(\frac{1}{2\pi\sigma_2^2} \right)^{N_2(S)/2} \cdot \exp \left(-\frac{1}{2} \sum_{i:S_i=2} \frac{(y_i - \mu_2)^2}{\sigma_2^2} \right) \end{aligned}$$

where $N_1(S) = N_2(S) = N/2$. The posteriors $p(\mu_k | \sigma_k^2, S, y)$ for $k = 1, 2$ in Eq. (1) then are $\mathcal{N}(b_k(S), B_k(S))$ -distributed with

$$B_k(S) = \frac{1}{N_0 + N_k(S)} \sigma_k^2 \quad (20)$$

$$b_k(S) = \frac{N_0}{N_k(S) + N_0} b_0 + \frac{N_k(S)}{N_k(S) + N_0} \bar{y}_k(S) \quad (21)$$

where also $N_1(S) = N_2(S) = N/2$ are half of total sample size and $\bar{y}_k(S)$ is the mean of group $k = 1, 2$. After choosing the conjugate prior $\mu_k \sim \mathcal{N}(b_0, B_0)$, these posteriors can be computed. The posteriors

$$\sigma_k^2 | \mu_k, S, y \sim \mathcal{G}^{-1}(c_k(S), C_k(S))$$

with

$$c_k(S) = c_0 + \frac{1}{2} N_k(S) \quad , \quad (22)$$

$$C_k(S) = C_0 + \frac{1}{2} \sum_{i:S_i=k} (y_i - \mu_k)^2 \quad , \quad (23)$$

for $k = 1, 2$ become

$$\sigma_1^2 | \mu_1, S, y \sim \mathcal{G}^{-1} \left(c_0 + \frac{1}{2} N_1(S), C_0 + \frac{1}{2} \sum_{i: S_i=1} (y_i - \mu_1)^2 \right)$$

$$\sigma_2^2 | \mu_2, S, y \sim \mathcal{G}^{-1} \left(c_0 + \frac{1}{2} N_2(S), C_0 + \frac{1}{2} \sum_{i: S_i=2} (y_i - \mu_2)^2 \right)$$

and again, after selecting a conjugate inverse-gamma prior $\sigma_k^2 \sim \mathcal{G}^{-1}(c_0, C_0)$ for $k = 1, 2$, these posteriors are also completely determined.

The necessary marginal posteriors for σ_k^2 for $k = 1, 2$ are then obtained, following the derivations in the above section, as

$$p(\sigma_1^2 | S, y) \sim \mathcal{G}^{-1} \left(c_0 + \frac{1}{2} N_1(S), C_0 + \frac{1}{2} \left(N_1(S) s_{y,1}^2(S) + \frac{N_1(S) N_0}{N_1(S) + N_0} (\bar{y}_1(S) - b_0)^2 \right) \right)$$

and

$$p(\sigma_2^2 | S, y) \sim \mathcal{G}^{-1} \left(c_0 + \frac{1}{2} N_2(S), C_0 + \frac{1}{2} \left(N_2(S) s_{y,2}^2(S) + \frac{N_2(S) N_0}{N_2(S) + N_0} (\bar{y}_2(S) - b_0)^2 \right) \right).$$

These marginal posteriors are completely determined, once c_0, C_0 is given by the selected prior $\mathcal{G}^{-1}(c_0, C_0)$ and the group variances $s_{y,1}^2(S)$ and $s_{y,2}^2(S)$ are calculated. Again here, $N_1(S) = N_2(S) = N/2$ due to equal sizes of both groups, and the $\bar{y}_1(S)$ and $\bar{y}_2(S)$ are the means of the two groups. The joint posterior, which is the ultimate quantity of interest, then can be rewritten as follows:

$$\begin{aligned} p(\mu, \sigma^2 | S, y) &= p(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2 | S, y) = \prod_{k=1}^2 p(\mu_k, \sigma_k^2 | S, y) \\ &= \prod_{k=1}^2 \underbrace{p(\mu_k | \sigma_k^2, S, y)}_{=: (A)} \cdot \underbrace{p(\sigma_k^2 | S, y)}_{=: (B)} \\ &= p(\mu_1 | \sigma_1^2, S, y) p(\sigma_1^2 | S, y) \\ &\quad \cdot p(\mu_2 | \sigma_2^2, S, y) p(\sigma_2^2 | S, y) \\ &= \mathcal{N}(b_1(S), B_1(S)) \cdot \mathcal{G}^{-1}(c_1(S), C_1(S)) \\ &\quad \cdot \mathcal{N}(b_2(S), B_2(S)) \cdot \mathcal{G}^{-1}(c_2(S), C_2(S)) \end{aligned}$$

A.2 Proof of Theorem 1: Derivation of the Full Conditionals for the Single-Block Gibbs Sampler

Proof To make Gibbs sampling possible, the full conditionals of

$$p(\mu, \sigma^2 | S, y) = p(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2 | S, y)$$

need to be derived. We start with $p(\mu_1 | \mu_2, \sigma_1^2, \sigma_2^2, S, y)$:

$$\begin{aligned} p(\mu_1 | \mu_2, \sigma_1^2, \sigma_2^2, S, y) &= \frac{p(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2, S, y)}{p(\mu_2, \sigma_1^2, \sigma_2^2, S, y)} \\ &= \frac{p(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2 | S, y) p(S, \overline{y})}{p(\mu_2, \sigma_1^2, \sigma_2^2 | S, y) p(S, \overline{y})} \\ &\stackrel{(1)}{=} \frac{\prod_{k=1}^2 p(\mu_k | \sigma_k^2, S, y) p(\sigma_k^2 | S, y)}{p(\mu_2, \sigma_2^2 | S, y) p(\sigma_1^2 | S, y)} \\ &\stackrel{(2)}{=} \frac{p(\mu_1 | \sigma_1^2, S, y) p(\sigma_1^2 | S, \overline{y}) p(\mu_2 | \sigma_2^2, S, y) p(\sigma_2^2 | S, \overline{y})}{p(\mu_2 | \sigma_2^2, S, y) p(\sigma_2^2 | S, \overline{y}) p(\sigma_1^2 | S, \overline{y})} \\ &= p(\mu_1 | \sigma_1^2, S, y) \end{aligned}$$

where in (1), the independence of σ_1^2 and μ_2, σ_2^2 was used, so that $p(\mu_2, \sigma_1^2, \sigma_2^2 | S, y) = p(\mu_2, \sigma_2^2 | S, y) \cdot p(\sigma_1^2 | S, y)$ holds, and in (2), the factorization

$$p(\mu_2, \sigma_2^2 | S, y) = p(\mu_2 | \sigma_2^2, S, y) \cdot p(\sigma_2^2 | S, y)$$

was used.

The full conditional of μ_2 is given by

$$p(\mu_2 | \mu_1, \sigma_1^2, \sigma_2^2, S, y),$$

which is derived as

$$\begin{aligned} p(\mu_2 | \mu_1, \sigma_1^2, \sigma_2^2, S, y) &= \frac{p(\mu_2, \mu_1, \sigma_1^2, \sigma_2^2, S, y)}{p(\mu_1, \sigma_1^2, \sigma_2^2, S, y)} \\ &= \frac{p(\mu_2, \mu_1, \sigma_1^2, \sigma_2^2 | S, y) p(S, \overline{y})}{p(\mu_1, \sigma_1^2, \sigma_2^2 | S, y) p(S, \overline{y})} \\ &= \frac{\prod_{k=1}^2 p(\mu_k | \sigma_k^2, S, y) p(\sigma_k^2 | S, y)}{p(\mu_1, \sigma_1^2 | S, y) p(\sigma_2^2 | S, y)} \\ &= \frac{p(\mu_1 | \sigma_1^2, S, \overline{y}) p(\sigma_1^2 | S, \overline{y}) p(\mu_2 | \sigma_2^2, S, y) p(\sigma_2^2 | S, \overline{y})}{p(\mu_1 | \sigma_1^2, S, \overline{y}) p(\sigma_2^2 | S, \overline{y}) p(\sigma_1^2 | S, \overline{y})} \\ &= p(\mu_2 | \sigma_2^2, S, y) \end{aligned}$$

where the reasoning is the same as in the derivation of the full conditional for μ_1 .

The full conditional of σ_1^2 is $p(\sigma_1^2 | \mu_1, \mu_2, \sigma_2^2, S, y)$, which is derived as follows:

$$\begin{aligned}
 p(\sigma_1^2 | \mu_1, \mu_2, \sigma_2^2, S, y) &= \frac{p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2, S, y)}{p(\mu_1, \mu_2, \sigma_2^2, S, y)} \\
 &= \frac{p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2 | S, y) \cancel{p(S, y)}}{p(\mu_1, \mu_2, \sigma_2^2 | S, y) \cancel{p(S, y)}} \\
 &= \frac{p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2 | S, y)}{p(\mu_1, \mu_2, \sigma_2^2 | S, y)} \\
 &\stackrel{(1)}{=} \frac{\prod_{k=1}^2 p(\sigma_k^2 | \mu_k, S, y) \cdot p(\mu_k | S, y)}{p(\mu_1 | S, y) \cdot p(\mu_2, \sigma_2^2 | S, y)} \\
 &\stackrel{(2)}{=} \frac{p(\sigma_1^2 | \mu_1, S, y) \cdot p(\mu_1 | S, y) \cdot \cancel{p(\sigma_2^2 | \mu_2, S, y)} \cdot \cancel{p(\mu_2 | S, y)}}{p(\mu_1 | S, y) \cdot \cancel{p(\sigma_2^2 | \mu_2, S, y)} \cdot \cancel{p(\mu_2 | S, y)}} \\
 &= \frac{\cancel{p(\mu_1 | S, y)} \cdot p(\sigma_1^2 | \mu_1, S, y)}{\cancel{p(\mu_1 | S, y)}} \\
 &= p(\sigma_1^2 | \mu_1, S, y)
 \end{aligned}$$

where in (1) first the independence of parameters between both groups was used, that is, the independence of μ_1, σ_1^2 and μ_2, σ_2^2 and second (as a special case of this fact) the independence of μ_1 and μ_2, σ_2^2 was used. Therefore, $p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2 | S, y) = \prod_{k=1}^2 p(\sigma_k^2 | \mu_k, S, y) \cdot p(\mu_k | S, y)$ holds and also $p(\mu_1, \mu_2, \sigma_2^2 | S, y) = p(\mu_1 | S, y) \cdot p(\mu_2, \sigma_2^2 | S, y)$. In (2), the factorization $p(\mu_2, \sigma_2^2 | S, y) = p(\sigma_2^2 | \mu_2, S, y) \cdot p(\mu_2 | S, y)$ was used.

The full conditional $p(\sigma_2^2 | \mu_1, \mu_2, \sigma_1^2, S, y)$ of σ_2^2 is similarly given by

$$\begin{aligned}
 p(\sigma_2^2 | \mu_1, \mu_2, \sigma_1^2, S, y) &= \frac{p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2, S, y)}{p(\mu_1, \mu_2, \sigma_1^2, S, y)} \\
 &= \frac{p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2 | S, y) \cancel{p(S, y)}}{p(\mu_1, \mu_2, \sigma_1^2 | S, y) \cancel{p(S, y)}} \\
 &= \frac{p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2 | S, y)}{p(\mu_1, \mu_2, \sigma_1^2 | S, y)} \\
 &\stackrel{(1)}{=} \frac{\prod_{k=1}^2 p(\sigma_k^2 | \mu_k, S, y) \cdot p(\mu_k | S, y)}{p(\mu_2 | S, y) \cdot p(\mu_1, \sigma_1^2 | S, y)} \\
 &\stackrel{(2)}{=} \frac{\cancel{p(\sigma_1^2 | \mu_1, S, y)} \cdot \cancel{p(\mu_1 | S, y)} \cdot p(\sigma_2^2 | \mu_2, S, y) \cdot p(\mu_2 | S, y)}{p(\mu_2 | S, y) \cdot \cancel{p(\sigma_1^2 | \mu_1, S, y)} \cdot \cancel{p(\mu_1 | S, y)}} \\
 &= \frac{\cancel{p(\mu_2 | S, y)} \cdot p(\sigma_2^2 | \mu_2, S, y)}{\cancel{p(\mu_2 | S, y)}} \\
 &= p(\sigma_2^2 | \mu_2, S, y)
 \end{aligned}$$

where in (1) first the independence of parameters between both groups, that is, the independence of μ_1, σ_1^2 and μ_2, σ_2^2 and second (as a special case of this) the independence of μ_2 and μ_1, σ_1^2 was used. Therefore, $p(\sigma_1^2, \mu_1, \mu_2, \sigma_2^2 | S, y) = \prod_{k=1}^2 p(\sigma_k^2 | \mu_k, S, y) \cdot p(\mu_k | S, y)$ holds and one also has $p(\mu_1, \mu_2, \sigma_2^2 | S, y) = p(\mu_1 | S, y) \cdot p(\mu_2, \sigma_2^2 | S, y)$. In (2), the factorization

$p(\mu_1, \sigma_1^2 | S, y) = p(\sigma_1^2 | \mu_1, S, y) \cdot p(\mu_1 | S, y)$ was used. In total, the full conditionals, therefore, are derived as follows:

$$p(\mu_1 | \mu_2, \sigma_1^2, \sigma_2^2, S, y) = p(\mu_1 | \sigma_1^2, S, y)$$

$$p(\mu_2 | \mu_1, \sigma_1^2, \sigma_2^2, S, y) = p(\mu_2 | \sigma_2^2, S, y)$$

$$p(\sigma_1^2 | \mu_1, \mu_2, \sigma_2^2, S, y) = p(\sigma_1^2 | \mu_1, S, y)$$

$$p(\sigma_2^2 | \mu_1, \mu_2, \sigma_1^2, S, y) = p(\sigma_2^2 | \mu_2, S, y)$$

When using the independence prior, the full conditionals are, therefore, given by

$$p(\mu_1 | \mu_2, \sigma_1^2, \sigma_2^2, S, y) \sim \mathcal{N}(b_1(S), B_1(S)) \quad (24)$$

$$p(\mu_2 | \mu_1, \sigma_1^2, \sigma_2^2, S, y) \sim \mathcal{N}(b_2(S), B_2(S)) \quad (25)$$

$$p(\sigma_1^2 | \mu_1, \mu_2, \sigma_2^2, S, y) \sim \mathcal{G}^{-1}(c_1(S), C_1(S)) \quad (26)$$

$$p(\sigma_2^2 | \mu_1, \mu_2, \sigma_1^2, S, y) \sim \mathcal{G}^{-1}(c_2(S), C_2(S)) \quad (27)$$

with $b_1(S), B_1(S), b_2(S), B_2(S), c_1(S), c_2(S), C_1(S)$ and $C_2(S)$ as specified in the previous section in Eqs. (20), (21), (22) and (23), which completes the proof. $\square$

A.3 Proof of Corollary 1

Proof From standard MCMC theory, see e.g. Robert and Casella [31], it is clear that the full conditionals derived in Theorem 1 can be used to construct a Gibbs sampler, by iteratively updating each parameter via simulating step by step from the full conditionals (24) to (27). Using (24) to (27), this leads to the following Gibbs sampling algorithm:

1. Sample σ_k^2 in each group $k, k = 1, 2$ from an inverse Gamma distribution $\mathcal{G}^{-1}(c_k(S), C_k(S))$ (which depends on μ_k)
2. Sample μ_k in each group $k, k = 1, 2$, from a normal distribution $\mathcal{N}(b_k(S), B_k(S))$ (which depends on σ_k^2)

where $B_k(S), b_k(S)$ and $c_k(S), C_k(S)$ are given by Equations (20), (21), (22) and (23). The convergence to the joint posterior $p(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2 | S, y)$ then follows then from standard MCMC theory, see Robert and Casella [31]. $\square$

A.4 Proof of Theorem 2

Proof In the case, the ROPE R_j is correct, R_j includes the true effect size δ_0 , so that $\delta_0 \subset R_j$. Estimation of δ via δ_{MPE} is then consistent: A posterior distribution p_n is said to be consistent for the parameter δ_0 , if for every neighbourhood U of δ_0 , $p_n(U) \xrightarrow[n \rightarrow \infty]{a.s.} 1$ almost surely under the law determined by δ_0 . By Theorem 1 in Ghosal [10], the consistency of the posterior follows for any prior π when choosing any ROPE U including the true parameter δ_0 , except possibly on a set of π -measure zero. As a direct consequence one, therefore, obtains: If the ROPE $R_j \neq \emptyset$ is correct and contains the true parameter δ_0 , any prior π leads to a consistent posterior for which $p_n(R_j) \xrightarrow[n \rightarrow \infty]{a.s.} 1$ almost surely under the law determined by δ_0 except possibly on a set of π -measure zero. This means that

$$\text{PMP}(\delta_{\text{MPE}}) = \int_{R_j} f(\theta|x) d\theta \xrightarrow[n \rightarrow \infty]{a.s.} 1.$$

If on the other hand R_j is incorrect, then $\delta_0 \notin R_j$. Then there exists a neighbourhood $N := (\delta_0 - \varepsilon, \delta_0 + \varepsilon)$ for $\varepsilon > 0$ around δ_0 , so that $N \cap R_j = \emptyset$. Then, as $p_n(N) \xrightarrow[n \rightarrow \infty]{a.s.} 1$ almost surely under the law determined by δ_0 except possibly on a set of π -measure zero, it follows that on the complement N^c , $p_n(N^c) \xrightarrow[n \rightarrow \infty]{a.s.} 0$ and because of $R_j \subset N^c$ also that $p_n(R_j) \xrightarrow[n \rightarrow \infty]{a.s.} 0$ almost surely under the law determined by δ_0 except possibly on a set of π -measure zero. Thereby it follows that

$$\text{PMP}(\delta_{\text{MPE}}) = \int_{R_j} f(\theta|x) d\theta \xrightarrow[n \rightarrow \infty]{a.s.} 0$$

□

Proof of Theorem 3

Proof An α type I error happens if the true parameter value $\delta_0 \in H$, with $H \subset R$ for a ROPE $R \subset \Theta$, but H is α -rejected for α . If any correct ROPE R is selected around the hypothesis $H \subset \Theta$ which makes a statement about the unknown parameter δ , then the true value δ_0 of δ is inside R , that is, $\delta_0 \subset R$. Then, under any prior π on δ , the posterior $p_n(R) \xrightarrow[n \rightarrow \infty]{a.s.} 1$ except on a set of π -measure zero, compare [10]. Therefore, the corresponding $\alpha\%$ HPD interval C_α of $f(\delta|x)$ lies inside R for $n \rightarrow \infty$, too, that is $C_\alpha \subset R$, and by definition, H is then α -accepted. As α was arbitrary, the above holds in particular without loss of generality for $\alpha = 1$, and therefore, H is accepted always for $n \rightarrow \infty$ almost surely under the law determined by π for any correct ROPE R around the hypothesis H . This in turn implies that H can only be α -rejected for $\alpha = 0$ under the above conditions.⁸

⁸ Note that α rejection of H for $\alpha = 0$ is only stating that zero per cent of the HPD interval are located outside the ROPE R around H , which means that the $\alpha\%$ HPD lies fully inside the ROPE. This in fact confirms H , being in agreement with the intuition of no type I error occurring.

An α type II error happens if the true parameter value $\delta_0 \notin H$ and $\delta \notin R$, with $H \subset R$ for a ROPE $R \subset \Theta$, but H is α -accepted for α . If any incorrect ROPE R is selected around the hypothesis $H \subset \Theta$ which makes a statement about the unknown parameter δ , then the true value δ_0 of δ is not inside R , that is, $\delta_0 \notin R$. Then, under any prior π on δ , the posterior $p_n(R) \xrightarrow[n \rightarrow \infty]{a.s.} 0$ except on a set of π -measure zero, compare [10]. Therefore, the corresponding $\alpha\%$ HPD interval C_α of $f(\delta|x)$ lies not inside R for $n \rightarrow \infty$, which means $C_\alpha \notin R$, and by definition, H is then α -rejected. As α was arbitrary, the above holds in particular for $\alpha = 1$, and therefore, H is rejected always for $n \rightarrow \infty$ almost surely under the law determined by π for any incorrect ROPE R around the hypothesis H . This also implies that H can only be α -accepted for $\alpha = 0$ under the above conditions.⁹ $\square$

Data Availability All data used in the simulations and illustrative example are available at <https://osf.io/cvwr5/>.

Code Availability A complete replication script to run the simulations and produce all results is available at the Open Science Framework under <https://osf.io/cvwr5/>. The proposed Bayesian t test is made available in the R package `bayest` at CRAN: <https://cran.r-project.org/web/packages/bayest/index.html>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Berger J, Wolpert RL (1988) The likelihood principle. Institute of Mathematical Statistics, Hayward
2. van den Bergh D, Doorn JV, Marsman M, Gupta KN, Sarafoglou A, Jan G, Stefan A, Ly A, Hinne M (2019) A tutorial on conducting and interpreting a Bayesian ANOVA in JASP. psyarxiv preprint. <https://psyarxiv.com/spreb>
3. Carlin B, Louis T (2009) Bayesian methods for data analysis. Chapman & Hall, CRC Press, Boca Raton
4. Carpenter B, Guo J, Hoffman MD, Brubaker M, Gelman A, Lee D, Goodrich B, Li P, Riddell A, Betancourt M (2017) Stan: a probabilistic programming language. J Stat Softw 76(1):1–32. <https://doi.org/10.18637/jss.v076.i01>
5. Cohen J (1988) Statistical power analysis for the behavioral sciences, 2nd edn. Routledge, Hillsdale

⁹ Here too, α -acceptance for $\alpha = 0$ can be interpreted as zero per cent of the HPD being located inside the ROPE R around H , being in agreement with the intuition of no type II error occurring.

6. Cumming G (2014) The new statistics: why and how. *Psychol Sci* 25(1):7–29. <https://doi.org/10.1177/0956797613504966>
7. Escobar MD, West M (1995) Bayesian density estimation and inference using mixtures. *J Am Stat Assoc* 90(430):577–588. <https://doi.org/10.1080/01621459.1995.10476550>
8. Freedman LS, Lowe D, Macaskill P (1983) Stopping rules for clinical trials. *Stat Med* 2(2):167–174. <https://doi.org/10.1002/sim.4780020210>
9. Frühwirth-Schnatter S (2006) Finite mixture and Markov switching models. Springer, New York
10. Ghosal S (1996) A review of consistency and convergence of posterior distribution. In: Proceedings of Varanashi symposium in Bayesian inference. Banaras Hindu University
11. Gigerenzer G, Marewski JN (2015) Surrogate science: the idol of a universal method for scientific inference. *J Manag* 41(2):421–440. <https://doi.org/10.1177/0149206314547522>
12. Gönen M, Johnson WO, Lu Y, Westfall PH (2005) The Bayesian two-sample t test. *Am Stat* 59(3):252–257. <https://doi.org/10.1198/000313005X55233>
13. Greenland S, Senn SJ, Rothman KJ, Carlin JB, Poole C, Goodman SN, Altman DG (2016) Statistical tests, P values, confidence intervals, and power: a guide to misinterpretations. *Eur J Epidemiol* 31(4):337–350. <https://doi.org/10.1007/s10654-016-0149-3>
14. Gronau QF, Ly A, Wagenmakers EJ (2020) Informed Bayesian t-tests. *Am Stat* 74(2):137–143. <https://doi.org/10.1080/00031305.2018.1562983>
15. Held L, Sabanés Bové D (2014) Applied statistical inference. Springer, Berlin. <https://doi.org/10.1007/978-3-642-37887-4>
16. Hobbs BP, Carlin BP (2007) Practical Bayesian design and analysis for drug and device clinical trials. *J Biopharm Stat* 18(1):54–80
17. Hodges JL, Lehmann EL (1954) Testing the approximate validity of statistical hypotheses. *J R Stat Soc: Ser B (Methodol)* 16(2):261–268. <https://doi.org/10.1111/j.2517-6161.1954.tb00169.x>
18. Kamary K, Mengersen K, Robert CP, Rousseau J (2014) Testing hypotheses via a mixture estimation model, pp 1–37. arXiv preprint. <https://arxiv.org/abs/1412.2044>. <https://doi.org/10.16373/j.cnki.ahr.150049>
19. Kelter R (2020) Analysis of Bayesian posterior significance and effect size indices for the two-sample t-test to support reproducible medical research. *BMC Med Res Methodol*. <https://doi.org/10.1186/s12874-020-00968-2>
20. Kelter R (2020) Bayesian survival analysis in STAN for improved measuring of uncertainty in parameter estimates. *Meas: Interdiscip Res Perspect* 18(2):101–119. <https://doi.org/10.1080/15366367.2019.1689761>
21. Kelter R (2021) Bayesian Hodges-Lehmann tests for statistical equivalence in the two-sample setting: power analysis, type I error rates and equivalence boundary selection in biomedical research. *BMC Med Res Methodol*. <https://doi.org/10.1186/s12874-021-01341-7>
22. Kelter R (2021) On the measure-theoretic premises of Bayes factor and full Bayesian significance tests: a critical reevaluation. *Comput Brain Behav*. <https://doi.org/10.1007/s42113-021-00110-5>
23. Kruschke JK (2018) Rejecting or accepting parameter values in Bayesian estimation. *Adv Methods Pract Psychol Sci* 1(2):270–280. <https://doi.org/10.1177/2515245918771304>
24. Kruschke JK, Liddell T (2018) The Bayesian new statistics : hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychon Bull Rev* 25:178–206. <https://doi.org/10.3758/s13423-016-1221-4>
25. Lakens D (2014) Performing high-powered studies efficiently with sequential analyses. *Eur J Soc Psychol* 44(7):701–710. <https://doi.org/10.1002/ejsp.2023>
26. Lakens D, Scheel AM, Isager PM (2018) Equivalence testing for psychological research: a tutorial. *Adv Methods Pract Psychol Sci* 1(2):259–269. <https://doi.org/10.1177/2515245918770963>
27. McElreath R, Smaldino PE (2015) Replication, communication, and the population dynamics of scientific discovery. *PLoS ONE* 10(8):1–16. <https://doi.org/10.1371/journal.pone.0136088>
28. Nuijten MB, Hartgerink CH, van Assen MA, Epskamp S, Wicherts JM (2016) The prevalence of statistical reporting errors in psychology (1985–2013). *Behav Res Methods* 48(4):1205–1226. <https://doi.org/10.3758/s13428-015-0664-2>
29. Rao CR, Lovric MM (2016) Testing point null hypothesis of a normal mean and the truth: 21st century perspective. *J Mod Appl Stat Methods* 15(2):2–21. <https://doi.org/10.22237/jmasm/1478001660>
30. Richardson S, Green P (1997) On Bayesian analysis of mixtures with an unknown number of components. *J R Stat Soc Ser B (Methodol)* 59(4):731–792
31. Robert C, Casella G (2004) Monte Carlo statistical methods. Springer, New York

32. Rouder JN, Speckman PL, Sun D, Morey RD, Iverson G (2009) Bayesian t tests for accepting and rejecting the null hypothesis. *Psychon Bull Rev* 16(2):225–237. <https://doi.org/10.3758/PBR.16.2.225>
33. Rüschemdorf L (2014) *Mathematische Statistik*. Springer, Berlin
34. Schuirmann D (1987) A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. *J Pharmacokinet Biopharm* 15(6):657–80
35. van Doorn J, van den Bergh D, Bohm U, Dablander F, Derks K, Draws T, Evans NJ, Gronau QF, Hinne M, Kucharský S, Ly A, Marsman M, Matzke D, Raj A, Sarafoglou A, Stefan A, Voelkel JG, Wagenmakers EJ (2019) The JASP guidelines for conducting and reporting a Bayesian analysis. *psyarxiv preprint*. <https://psyarxiv.com/yqxfr>. <https://doi.org/10.31234/osf.io/yqxfr>
36. Wagenmakers EJ, Beek T, Rotteveel M, Gierholz A, Matzke D, Steingroever H, Ly A, Verhagen J, Selker R, Sasiadek A, Gronau QF, Love J, Pinto Y (2015) Turning the hands of time again: a purely confirmatory replication study and a Bayesian analysis. *Front Psychol* 6:494
37. Wang M, Liu G (2016) A simple two-sample Bayesian t-test for hypothesis testing. *Am Stat* 70(2):195–201. <https://doi.org/10.1080/00031305.2015.1093027>
38. Wasserstein RL, Lazar NA (2016) The ASA’s statement on p-values: context, process, and purpose. *Am Stat* 70(2):129–133. <https://doi.org/10.1080/00031305.2016.1154108>
39. Wasserstein RL, Schirm AL, Lazar NA (2019) Moving to a world beyond “ $p < 0.05$ ”. *Am Stat* 73(sup1):1–19. <https://doi.org/10.1080/00031305.2019.1583913>
40. Wetzels R, Matzke D, Lee MD, Rouder JN, Iverson GJ, Wagenmakers EJ (2011) Statistical evidence in experimental psychology: an empirical comparison using 855 t tests. *Perspect Psychol Sci* 6(3):291–298. <https://doi.org/10.1177/1745691611406923>