The Pearson Bayes factor: An analytic formula for computing evidential value from minimal summary statistics

Thomas J. Faulkenberry

Department of Psychological Sciences, Tarleton State University, Stephenville, Texas,
76402, USA, e-mail: faulkenberry@tarleton.edu

SUMMARY

In Bayesian hypothesis testing, evidence for a statistical model is quantified by the Bayes factor, which represents the relative likelihood of observed data under that model compared to another competing model. In general, computing Bayes factors is difficult, as computing the marginal likelihood of data under a given model requires integrating over a prior distribution of model parameters. In this paper, I capitalize on a particular choice of prior distribution that allows the Bayes factor to be expressed without integral representation, and I develop a simple formula – the Pearson Bayes factor – that requires only minimal summary statistics as commonly reported in scientific papers, such as the t or F score and the degrees of freedom. In addition to presenting this new result, I provide several examples of its use and report a simulation study validating its performance. Importantly, the Pearson Bayes factor gives applied researchers the ability to compute exact Bayes factors from minimal summary data, and thus easily assess the evidential value of any data for which these summary statistics are provided, even when the original data is not available.

Key words: Bayesian statistics, Bayes factor, Pearson Type VI distribution, summary statistics, t -test; analysis of variance

1. Introduction

Across many fields of study, the hypothesis test remains one of the primary tools used by researchers to distinguish signal from noise in empirical observations. Two of the most commonly used tools are the t -test and the analysis of variance. Though classically considered within a frequentist paradigm, Bayesian versions of these tests are becoming increasingly popular. Bayesian hypothesis testing makes use of the *Bayes factor* (Kass and Raftery, 1995),

which represents the relative likelihood of observed data under one model – say, a model hypothesizing a nonzero group difference – compared to a competing model (e.g., a model hypothesizing a null difference). Bayes factors have many desirable properties for inference, but their difficulty of computation has likely hindered their widespread adoption across the sciences.

One method that has been proposed for computing Bayes factors is the BIC method (Kass and Raftery, 1995; Wagenmakers, 2007), which can be used even when only summary statistics are available (Faulkenberry, 2018, 2019a). This popular and well-cited method only requires the use of a pocket calculator, but it suffers from one primary drawback – it is an *approximate* method whose performance suffers in small samples. In this paper, I will present a new formula for computing Bayes factors from minimal summary statistics – the *Pearson Bayes Factor* (PBF). This formula is based on the work of Min Wang and colleagues (Wang and Sun, 2014; Wang and Liu, 2016) and retains the easy-to-calculate nature of the BIC method while providing the user with an exact Bayes factor computation, regardless of sample size. The formula is summarized in the following proposition:

Proposition 1 (Pearson Bayes factor). *Given an analysis of variance summary reported in the form $F(x, y)$, where x equals the between-treatments degrees of freedom and y equals the residual (within-treatments) degrees of freedom, the Bayes factor can be expressed exactly as*

$$\text{PBF}_{10} = \frac{\Gamma\left(\frac{x}{2} + \alpha + 1\right) \cdot \Gamma\left(\frac{y}{2}\right)}{\Gamma\left(\frac{x+y}{2}\right) \cdot \Gamma(\alpha + 1)} \left(\frac{y}{y + xF}\right)^{\alpha - \frac{y}{2} + 1}.$$

The paper is roughly organized as follows. After first providing some general statistical background, I will review some of the recent methods used to compute Bayes factors for t -tests and analysis of variance designs and illustrate the various problems associated with their use. Then, I will outline the development of the Pearson Bayes factor and provide examples of its computation. Finally, I will present a simulation study which shows that the Pearson Bayes factor compares favorably against other well-known Bayes factors, and in some cases, it has properties which make it a better choice to use.

2. Background

2.1. t -test

Originally devised by William Sealy Gosset, who published his work under the pseudonym Student (1908), the t -test uses the mean and variance of two observed samples to detect differences in the underlying population means. The t -test works by first assuming a null hypothesis $\mathcal{H}_0$, and then calculating a t -score, which indexes the likelihood of obtaining some sample of observed data under the null hypothesis. If this probability is small, we reject the null hypothesis $\mathcal{H}_0$ in favor of some alternative hypothesis $\mathcal{H}_1$.

Specifically, let us consider the following general setup. Let x_{ij} denote measurements for the i^{th} participant ($i = 1, \dots, N$) in the j^{th} experimental condition ($j = 1, 2$). Further, we assume that the x_{ij} are drawn from independent and normally distributed populations with mean μ_j and variance σ^2 . Then we can test the hypotheses

$$\mathcal{H}_0 : \mu_1 = \mu_2 \text{ versus } \mathcal{H}_1 : \mu_1 \neq \mu_2$$

by computing a test statistic

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\hat{\sigma}_p / \sqrt{N_\delta}}$$

whose components are computed as follows. $\bar{x}_j$ represents the sample mean of the measurements in group j . $\hat{\sigma}_p$ is the “pooled” estimate of σ , defined by the relationship

$$\hat{\sigma}_p^2 = \frac{\hat{\sigma}_1^2(N_1 - 1) + \hat{\sigma}_2^2(N_2 - 1)}{N_1 + N_2 - 2},$$

where each $\hat{\sigma}_j$ is the sample standard deviation of the measurements in group j . Finally, $N_\delta = (1/N_1 + 1/N_2)^{-1}$; this is often called the effective sample size for the experiment.

Under the null hypothesis $\mathcal{H}_0$, the distribution of these t scores is well known as *Student’s t distribution*, with density function

$$f_\nu(x) = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}},$$

where ν represents the degrees of freedom of the test and $x \in (-\infty, \infty)$. The cumulative distribution function $F_\nu(x) = \int_{-\infty}^x f_\nu(u)du$ can then be used to

index the probability of observing data at least as extreme as that which we observed under the null hypothesis $\mathcal{H}_0$. Specifically, for an observed t -statistic t_{obs} , we compute $P(|x| > t_{\text{obs}}) = 2(1 - F_\nu(t_{\text{obs}}))$, a quantity commonly known as a p -value. If this probability is small (say, less than 5%), then one may decide to reject $\mathcal{H}_0$ in favor of $\mathcal{H}_1$ and conclude that $\mu_1 \neq \mu_2$, thus implying that the two populations from which we sampled are indeed different.

2.2. Analysis of variance

In 1925, Sir Ronald Fisher introduced the analysis of variance (Fisher, 1925). As a generalization of the t -test procedure that works for comparing multiple population means, the analysis of variance (or ANOVA) has become one of the most enduring techniques of hypothesis testing in the experimental sciences. In its simplest form, the analysis of variance is designed to test for differences among multiple group means, making it useful for a wide variety of scientific investigations. Since its introduction, the analysis of variance has become widespread in its popularity and a core topic in most introductory statistics textbooks. In fact, its use is so ubiquitous in the psychological and behavioral sciences that Rouder et al. (Rouder et al., 2016) referred to the analysis of variance as the “workhorse” of experimental psychology.

Roughly, the analysis of variance works by partitioning the total variance in a set of observed data $\mathcal{D}$ into two sources: the variance between experimental treatment groups, and the residual, or left over, variance. Then, one calculates an F statistic, which is defined as the ratio of the between-groups variance to the residual variance. Like the t -test above, inference about differences between the treatment groups is then performed by quantifying the likelihood of the observed data $\mathcal{D}$ under a null hypothesis $\mathcal{H}_0$. Specifically, this is done by computing the probability of obtaining the observed F statistic (or greater) under $\mathcal{H}_0$. If this p -value is small, this indicates that the data $\mathcal{D}$ are *rare* under $\mathcal{H}_0$, so the researcher may reject $\mathcal{H}_0$ in favor of an alternative hypothesis $\mathcal{H}_1$ which posits that there are significant differences among the treatment groups.

2.3. Alternatives to p -values?

Despite the popularity of the previously described approaches to testing group differences, there have been many recent criticisms against their use, and more generally, against null hypothesis significance testing (see Wagen-

makers, 2007)). In fact, the American Statistical Association has recently recommended against the use of p -values and significance testing for scientific inference (Wasserstein and Lazar, 2016). One alternative that has been recommended in its place is the *Bayes factor* (Jeffreys, 1961; Raftery, 1995), which indexes the extent to which the observed data $\mathcal{D}$ are more likely under one hypothesis (e.g., $\mathcal{H}_1$) than another (e.g., $\mathcal{H}_0$). The Bayes factor has several advantages over the p -value as a tool for inference, including the ability to index support for either the null hypothesis $\mathcal{H}_0$ or the alternative hypothesis $\mathcal{H}_1$. However, the widespread adoption of Bayes factors has likely been prevented due to the general difficulty of their computation.

In the next section, I will give a conceptual definition of Bayes factors and outline several approaches to computing Bayes factors for t -tests and ANOVAs.

3. Bayes factors

3.1. Conceptual definition

To introduce the Bayes factor, let us first recall Bayes' theorem. Here, we will relax our notation and let p represent both prior probabilities and likelihoods. In later sections, we will denote priors and likelihoods by the more conventional notations π and f , respectively. Bayes' theorem states

$$\underbrace{P(\mathcal{H} \mid \mathcal{D})}_{\text{Posterior beliefs about model}} = \underbrace{P(\mathcal{H})}_{\text{Prior beliefs about model}} \times \underbrace{\frac{P(\mathcal{D} \mid \mathcal{H})}{P(\mathcal{D})}}_{\text{predictive updating factor}}. \quad (1)$$

which tells us that our posterior belief about model $\mathcal{H}$ can be found by taking our *prior* belief about $\mathcal{H}$ and multiplying it by the fraction $P(\mathcal{D} \mid \mathcal{H})/P(\mathcal{D})$. This fraction is a *predictive updating factor* in the sense that it indexes how well our data is predicted by $\mathcal{H}$ compared to how well it is predicted by all other possible models. In hypothesis testing, our goal is to directly compare the predictive adequacy of two models $\mathcal{H}_0$ and $\mathcal{H}_1$. To this end, we can use Equation 1 to derive the following:

$$\frac{P(\mathcal{H}_0 \mid \mathcal{D})}{P(\mathcal{H}_1 \mid \mathcal{D})} = \frac{P(\mathcal{H}_0) \cdot \frac{P(\mathcal{D} \mid \mathcal{H}_0)}{P(\mathcal{D})}}{P(\mathcal{H}_1) \cdot \frac{P(\mathcal{D} \mid \mathcal{H}_1)}{P(\mathcal{D})}}$$

which we can simplify to

$$\underbrace{\frac{P(\mathcal{H}_0 | \mathcal{D})}{P(\mathcal{H}_1 | \mathcal{D})}}_{\text{posterior beliefs about models}} = \underbrace{\frac{P(\mathcal{H}_0)}{P(\mathcal{H}_1)}}_{\text{prior beliefs about models}} \times \underbrace{\frac{P(\mathcal{D} | \mathcal{H}_0)}{P(\mathcal{D} | \mathcal{H}_1)}}_{\text{predictive updating factor}}. \quad (2)$$

The predictive updating factor

$$\text{BF}_{01} = \frac{P(\mathcal{D} | \mathcal{H}_0)}{P(\mathcal{D} | \mathcal{H}_1)}$$

is called the *Bayes factor*, and as we can see in Equation (2), it represents the factor by which our relative belief in $\mathcal{H}_0$ over $\mathcal{H}_1$ can be updated after observing data. Note that we can easily express our relative belief in $\mathcal{H}_1$ over $\mathcal{H}_0$ by taking reciprocals; that is, $\text{BF}_{10} = 1/\text{BF}_{01}$.

One interesting consequence of Equation 2 is that we can use the Bayes factor to compute the posterior probability of $\mathcal{H}_0$ as a function of the prior model probabilities. To see this, consider the following. If we solve Equation 2 for the posterior probability $P(\mathcal{H}_0 | \mathcal{D})$ and then use Bayes' theorem, we see

$$\begin{aligned} P(\mathcal{H}_0 | \mathcal{D}) &= \text{BF}_{01} \cdot \frac{P(\mathcal{H}_0)}{P(\mathcal{H}_1)} \cdot P(\mathcal{H}_1 | \mathcal{D}) \\ &= \frac{\text{BF}_{01} \cdot P(\mathcal{H}_0) \cdot P(\mathcal{D} | \mathcal{H}_1) \cdot P(\mathcal{H}_1)}{P(\mathcal{H}_1) \cdot P(\mathcal{D})} \\ &= \frac{\text{BF}_{01} \cdot P(\mathcal{H}_0) \cdot P(\mathcal{D} | \mathcal{H}_1)}{P(\mathcal{D} | \mathcal{H}_0) \cdot P(\mathcal{H}_0) + P(\mathcal{D} | \mathcal{H}_1) \cdot P(\mathcal{H}_1)}. \end{aligned}$$

Dividing both numerator and denominator by the marginal likelihood $P(\mathcal{D} | \mathcal{H}_1)$ gives us

$$P(\mathcal{H}_0 | \mathcal{D}) = \frac{\text{BF}_{01} \cdot P(\mathcal{H}_0)}{\text{BF}_{01} \cdot P(\mathcal{H}_0) + P(\mathcal{H}_1)}.$$

Since $\text{BF}_{10} = 1/\text{BF}_{01}$, it can then be shown similarly that

$$P(\mathcal{H}_1 | \mathcal{D}) = \frac{\text{BF}_{10} \cdot P(\mathcal{H}_1)}{\text{BF}_{10} \cdot P(\mathcal{H}_1) + P(\mathcal{H}_0)}.$$

In practice, researchers often assume both models are *a priori* equally likely, and thus set both $P(\mathcal{H}_0) = P(\mathcal{H}_1) = 0.5$. In this case, we obtain the

simplified forms

$$P(\mathcal{H}_0 \mid \mathcal{D}) = \frac{\text{BF}_{01}}{\text{BF}_{01} + 1}, \quad P(\mathcal{H}_1 \mid \mathcal{D}) = \frac{\text{BF}_{10}}{\text{BF}_{10} + 1}. \quad (3)$$

To motivate the work in this paper, let us consider the general problem of computing Bayes factors. From the conceptual definition given in Equation 2, we see that the Bayes factor BF_{01} is the ratio of likelihoods for $\mathcal{D}$ under $\mathcal{H}_0$ and $\mathcal{H}_1$, respectively. Technically, they are *marginal* likelihoods, where the likelihood function $f_i(\mathcal{D} \mid \theta_i) = f(\mathcal{D} \mid \theta_i, \mathcal{H}_i)$ ($i = 0, 1$) is integrated over the prior distribution $\pi_i(\theta_i) = \pi(\theta_i, \mathcal{H}_i)$ for the vector of model parameters θ_i associated with model $\mathcal{H}_i$. That is,

$$\text{BF}_{01} = \frac{m_0(\mathcal{D})}{m_1(\mathcal{D})}$$

where

$$m_0(\mathcal{D}) = \int f_0(\mathcal{D} \mid \theta_0) \pi_0(\theta_0) d\theta_0$$

and

$$m_1(\mathcal{D}) = \int f_1(\mathcal{D} \mid \theta_1) \pi_1(\theta_1) d\theta_1.$$

In general, computing these marginal likelihoods can be difficult. One reason is that as sample size increases, the likelihood becomes highly peaked at its maximum. Without prior knowledge of the location of this maximum, numerical integration methods can fail to converge around this highly peaked area. Related to this is the problem of dimensionality: the parameter vector θ_i is often of high dimension. Combined with the peaked nature of the integrand, numerical methods often come up against difficulties – a problem that Thiele, Rouder, and Haaf refer to as “finding a needle in a haystack” (Thiele et al., 2017). Thus, when possible, it is desirable to make simplifying assumptions which lead to analytic solutions (i.e., without integral representation) for these marginal likelihoods.

3.2. BIC approximation

The most popular approach to this problem is the BIC approximation, described in detail by various papers of Kass and Raftery (e.g., Raftery, 1995; Kass and Raftery, 1995) and introduced to the applied psychological

science community by Wagenmakers (2007) and Masson (2011). The basic idea is to construct a second-order Taylor approximation of $\ln m_i(\mathcal{D}) = \ln P(\mathcal{D} \mid \theta_i, \mathcal{H}_i)$ about the posterior mode for θ_i . Of course, by removing all terms beyond the second derivative, the result is an approximation of the log marginal likelihood. However, Raftery (1995) pointed out that with a certain choice of noninformative prior on θ_i (the *unit information prior*), the order of the approximation is $\mathcal{O}(1/\sqrt{n})$. This results in the approximation

$$\ln m_i(\mathcal{D}) = \ln P(\mathcal{D} \mid \hat{\theta}_i, \mathcal{H}_i) - \frac{k_i}{2} \ln n + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right) \quad (4)$$

where $\hat{\theta}_i$ is the maximum likelihood estimate for θ_i under $\mathcal{H}_i$, k_i is the number of parameters in $\mathcal{H}_i$, and n is the total number of observations in $\mathcal{D}$. Equation 4 has a fortuitous relation to the Bayesian information criterion (BIC) of Schwarz (1978):

$$\text{BIC}(\mathcal{H}_i) = -2 \ln L_i + k_i \ln n, \quad (5)$$

where L is the maximum likelihood estimate for model $\mathcal{H}_i$. Combining Equations 4 and 5, we have

$$\ln m_i(\mathcal{D}) = -\frac{1}{2} \text{BIC}(\mathcal{H}_i) + \mathcal{O}\left(\frac{1}{\sqrt{n}}\right),$$

or equivalently

$$m_i(\mathcal{D}) \approx \exp\left(-\frac{1}{2} \text{BIC}(\mathcal{H}_i)\right).$$

It is then easy to derive the following expression for the BIC approximation to the Bayes factor:

$$\begin{aligned} \text{BF}_{01} &= \frac{m_0(\mathcal{D})}{m_1(\mathcal{D})} \approx \frac{\exp\left(-\frac{1}{2} \text{BIC}(\mathcal{H}_0)\right)}{\exp\left(-\frac{1}{2} \text{BIC}(\mathcal{H}_1)\right)} \\ &= \exp\left(\frac{\text{BIC}(\mathcal{H}_1) - \text{BIC}(\mathcal{H}_0)}{2}\right). \end{aligned} \quad (6)$$

Compared to computing marginal likelihoods by integrating the prior-weighted likelihoods, Equation 6 requires only that we know the BIC values

for models $\mathcal{H}_0$ and $\mathcal{H}_1$. For ANOVA models, the calculation is straightforward. In the ANOVA context, the BIC can be calculated (Raftery, 1995) as

$$BIC = n \ln(1 - R^2) + k \ln n = n \ln \left(\frac{SSR}{SST} \right) + k \ln n$$

where SST represents the total sum of squares (i.e., $\sum_i \sum_j (Y_{ij} - \bar{Y}_{..})^2$) for the model, SSR represents the residual sum of squares left over after accounting for the model's postulated treatment effects (i.e., $SST - SSA$, where $SSA = r \sum_i (\bar{Y}_{i.} - \bar{Y}_{..})^2$), and k equals the number of distinct treatment conditions in the model.

The BIC approach from Equation 6 requires that we have enough “raw” data available in order to compute SSR and SST . One improvement to Equation 6 is to remove this dependency and make the Bayes factor computation accessible with only summary statistics. Faulkenberry (2018) did exactly this, further refining the BIC approximation to produce a formula that can be used by the applied researcher who has access to the commonly reported summary statistics from the ANOVA, usually reported in the standard form $F(x, y)$, where F is the observed F -statistic, x is the degrees of freedom between treatments, and y is the residual (within treatments) degrees of freedom. This results in the simple closed form expression

$$BF_{01} \approx \sqrt{n^x \left(1 + \frac{Fx}{y} \right)^{-n}}. \quad (7)$$

The advantage of Faulkenberry's formula in Equation 7 is that researchers may compute approximations to Bayes factors from summary statistics only, whereas the approach of Wagenmakers (2007) and Masson (2011) that uses the BIC approximation in Equation 6 requires knowing the underlying sum of squares terms in the ANOVA calculation. Thus, Equation 7 can be useful to researchers who are trying to assess the evidential value of observed data from published summary statistics, especially where raw data (or even ANOVA summary tables) are not available (Faulkenberry, 2019b).

The Bayes factor approximation in Equation 7 can be easily adapted to use with the t -test by simply noting that $t^2 = F$ and $x = 1$. Thus, for a

t -test with degrees of freedom equal to ν , we have the approximation

$$\text{BF}_{01} \approx \sqrt{n \left(1 + \frac{t^2}{\nu}\right)^{-n}}.$$

Before proceeding further, it will be instructive to consider a simple example of computing the BIC Bayes factor for an ANOVA model.

Example 1. Consider the following data representing 6 measurements in each of 3 treatment groups:

Treatment 1	Treatment 2	Treatment 3
2	5	8
3	9	6
8	10	12
6	13	11
5	8	11
6	9	12

From here, it is a routine exercise to complete the following ANOVA summary table:

Source	SS	df	MS	F
Treatments	84	2	42	7.16
Residual	88	15	5.87	
Total	172	17		

To use Equation 6, we need to compute the BIC of $\mathcal{H}_0$ and $\mathcal{H}_1$. First, we compute

$$\text{BIC}(\mathcal{H}_1) = n \ln \left(\frac{SSR}{SST} \right) + k \ln n = 18 \ln \left(\frac{88}{172} \right) + 3 \ln(18) = -3.392.$$

Under $\mathcal{H}_0$, we note that all data are predicted by the grand mean μ only. Thus, all variation is to be considered as residual (i.e., $SSR = 172$), and there is $k = 1$ treatment condition. This gives us

$$\text{BIC}(\mathcal{H}_0) = 18 \ln \left(\frac{172}{172} \right) + 1 \ln(18) = 2.890.$$

Thus,

$$\text{BF}_{01} \approx \exp\left(\frac{\text{BIC}(\mathcal{H}_1) - \text{BIC}(\mathcal{H}_0)}{2}\right) = \exp\left(\frac{-3.392 - 2.890}{2}\right) = 0.0432.$$

Taking the reciprocal, we see

$$\text{BF}_{10} = \frac{1}{\text{BF}_{01}} \approx \frac{1}{0.0432} = 23.15$$

indicating that the observed data $\mathcal{D}$ are 23.15 times more likely under $\mathcal{H}_1$ than $\mathcal{H}_0$. Note that this also agrees with the formula using only the summary statistics from the ANOVA table: $F(2, 15) = 7.16$. Taking from this expression $x = 2$ and $y = 15$, we evaluate Equation 7 as

$$\text{BF}_{01} \approx \sqrt[n^x \left(1 + \frac{Fx}{y}\right)^{-n}]{} = \sqrt[18^2 \left(1 + \frac{7.16 \cdot 2}{15}\right)^{-18}]{} = 0.0432,$$

thus demonstrating the equivalence of the approaches from Equations 6 and 7.

Though Equations 6 and 7 provide simple ways to compute Bayes factors from ANOVA summaries, it is important to remember that the BIC formulas represent an *approximation* to the Bayes factor. This fact is underscored by the following computation. Sellke et al. (2001) showed that under a reasonable class of prior distributions for p -values, an upper bound for the Bayes factor can be computed directly from the p -value as

$$\text{BF}_{10} \leqslant -\frac{1}{e \cdot p \ln(p)}.$$

Applying this bound to our previous example, we first note that for $F(2, 15) = 7.16$, the associated p -value is $p = 0.0066$. This gives us the upper bound

$$\text{BF}_{10} \leqslant -\frac{1}{e \cdot 0.0066 \cdot \ln(0.0066)} = 11.10.$$

Thus, it is clear that our BIC Bayes factor of 23.15 is quite an overestimate of the actual Bayes factor, so our desire for *exact* Bayes factors is quite warranted.

4. Analytic methods for computing Bayes factors

Since its introduction by Kass and Raftery (1995), the BIC Bayes factor has been a popular method for estimating evidential value from empirical data. One of the attractive features of the BIC Bayes factor is that it can be computed with a simple calculator. This makes it easy to implement not only for the researcher, but also in the context of a beginning statistics course. However, it is a large-sample approximation, and the discussion above calls into question its computational stability for smaller examples, especially those which might be used in a traditional statistics course.

Thus, it is desirable to find *analytic methods* – that is, solutions without the need for approximating integrals – to replace the BIC approximation. In other words, we want a method which removes the integral representation, but unlike the BIC, gives an *exact* value for the Bayes factor. In the following sections, I will discuss some historical approaches to developing these methods for the t -test and the analysis of variance, followed by a discussion of their limitations.

4.1. Analytic Bayes factors for the t -test

One approach to developing a Bayesian t -test with an analytic solution was provided by Gönen et al. (2005). Gönen et al. reparameterized $\mathcal{H}_0$ and $\mathcal{H}_1$ in terms of “effect size” $\delta = \mu_1 - \mu_2$. Their goal was to test

$$\mathcal{H}_0 : \delta = 0 \text{ versus } \mathcal{H}_1 : \delta \neq 0,$$

using a parameterization $(\mu = (\mu_1 + \mu_2)/2, \delta, \sigma^2)$. Following Jeffreys (1961), they placed a noninformative prior on the mean and variance:

$$\pi(\mu, \sigma^2) \propto 1/\sigma^2$$

and assumed, under $\mathcal{H}_1$, that δ was drawn from a normal distribution centered at 0 with a variance parameter σ_a^2 . This variance term serves as a hyperparameter that needed to be specified in advance by the analyst. Given this setup, Gönen et al. (2005) showed

$$\text{BF}_{10} = \left(\frac{1 + \frac{t^2}{\nu}}{1 + \frac{t^2}{\nu(1+n_\delta\sigma_a^2)}} \right)^{\frac{\nu+1}{2}} (1 + n_\delta\sigma_a^2)^{-\frac{1}{2}}. \quad (8)$$

Certainly, Equation 8 meets our requirements of (1) being an exact Bayes factor, and (2) being solvable by means of a calculator, but it does have

some drawbacks. While the ability for the researcher to specify the prior variance of the effect size under $\mathcal{H}_1$ may be initially appealing, it is fraught with some unfortunate consequences.

Issue 1: It is a common practice to specify a large value for σ_a^2 , as such a choice would minimize the information provided by the prior and thus potentially better serve as an “objective” choice of prior. However, this causes a problem. Consider a fixed data summary; suppose we observed $t = 3.00$ with $N_1 = N_2 = 20$. Applying Equation 8 with increasing values of σ_a^2 , we see a troubling pattern, shown in the following table:

σ_a^2	1	10	100	1000	10000
BF ₁₀	8.99	6.01	2.08	0.67	0.21

Indeed, this table indicates that the Gönen Bayes factor exhibits *Bartlett’s paradox*; that is, for fixed information from the data, the Bayes factor tends to 0 (and thus always favors $\mathcal{H}_0$) with increasing prior variance (see also Wang and Liu, 2016).

Issue 2: Fix a set of observed data (samples) and suppose these observed data are always generated under $\mathcal{H}_1$. Then as t increases with fixed σ_a^2 , the Bayes factor in Equation 8 tends to $(1 + n_\delta \sigma_a^2)^{\eta/2}$ (see Figure 1). However, since the samples are drawn from $\mathcal{H}_1$, we would expect that as t increases without bound, the Bayes factor should also increase without bound. The unexpected asymptote in Figure 1 is referred to as the *information paradox*.

4.2. Analytic Bayes factors for ANOVA

Similar to the Gönen et al. (2005) t -test, it is also possible to calculate analytic solutions for Bayes factors with ANOVA summaries. One approach is to use Zellner’s g -prior approach, which was originally developed for the context of linear regression models (Zellner, 1986). In the g -prior approach, the standardized regression coefficients β are given a multivariate normal prior centered at 0 with a scaled covariance matrix that is proportional to the inverse Fisher information matrix for β . This scaling factor, denoted g , must be chosen by the analyst. Given the construction of the prior, g serves as a scale factor for the analyst to weigh the relative “importance” of the prior and the sample in the resulting inference. If $g = 1$, then the prior and the sample have the same importance in the inference. If $g = 2$, then the prior has half as much importance as the sample in the inference. This pattern continues; if $g = n$, then the prior has $1/n^{\text{th}}$ as much importance as the sample in the inference.

Figure 1. Information paradox: with fixed data under $\mathcal{H}_1$, the Gönen Bayes factor is bounded with increasing t

The g -prior Bayes factor can be computed as

$$\text{BF}_{10} = (1 + g)^{\frac{n-k-1}{2}} (1 + g(1 - R^2))^{-\frac{n-1}{2}} \quad (9)$$

where n equals the number of observations, k equals the number of predictors in the regression model, and R^2 equals the coefficient of determination for the regression model. The analyst must choose the value of g , as discussed above. Unfortunately, this leads to the same issues discussed with Gönen et al. (2005) above. First, assuming fixed n and k , increasing values of g in Equation 9 will cause the Bayes factor to approach 0 (i.e., Bartlett's paradox). Second, increasing values of R^2 from data drawn under $\mathcal{H}_1$ will only result in Bayes factors that approach an asymptote, $(1 + g)^{(n-k-1)/2}$ (i.e., the information paradox).

In light of these paradoxes (both the Gönen and Zellner approaches), a common approach for resolution is to place a prior distribution on the relevant scaling parameter (e.g., the g in the Zellner Bayes factor). Two solutions have been recently proposed: (1) the hyper- g approach (Liang et al., 2008); and (2) the JZS approach (Rouder et al., 2012). Both of these approaches place priors on the scale factor g and thus avoid the paradoxes described above. However, the drawback to both of these approaches is that the resulting Bayes factors have integral representation. That is, their computation requires evaluation of at least a one-dimensional integral. With

software, this is no problem, but these approaches are inaccessible to beginners (especially students) with no calculus background or programming knowledge.

4.3. Analytic Bayes factors without integral representation

Since its introduction by Kass and Raftery (1995), the BIC Bayes factor has been a popular method for estimating evidential value from empirical data. However, since it is based on a large sample approximation, its utility is limited, particularly in a small sample context. It is especially limited for use in a teaching context, where sample sizes for hand-worked examples tend to be small. Thus, it is desirable to find a calculation method to replace the BIC approximation. Specifically, we want a method which, like the BIC, removes the integral representation, but unlike the BIC, gives an exact value for the Bayes factor.

Recent work by Wang and Sun (2014) has made possible such a calculation. In what follows, we consider our data to be a balanced set of observations $\mathcal{D} = Y_{ij}$ where $i = 1, \dots, p$ represents the experimental unit (i.e., the treatment group) and $j = 1, \dots, r$ represents the specific replicate within each unit. This gives a total of $n = pr$ observations. On these data, we place a random effects linear model:

$$Y_{ij} = \mu + a_i + \varepsilon_{ij},$$

assuming that $a_i \sim \mathcal{N}(0, \sigma_a^2)$ and $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. In this context, the main goal of analysis of variance is to test whether the random effects term a_i is identically 0. We do this by restricting the variability of the random effects term. That is, we want to test

$$\mathcal{H}_0 : \sigma_a^2 = 0 \text{ versus } \mathcal{H}_1 : \sigma_a^2 \neq 0.$$

Wang and Sun extended the approach of García-Donato and Sun (2007), who assumed common, noninformative priors on μ and σ in both $\mathcal{H}_0$ and $\mathcal{H}_1$ and considered a proper prior on the ratio of variance components $\tau = \sigma_a^2/\sigma^2$ under $\mathcal{H}_1$. Under such prior specification, Garcia-Donato and Sun showed

$$BF_{10} = \int_0^\infty (1 + \tau r)^{\frac{1-p}{2}} \left(1 - \frac{\tau r}{1 + \tau r} \cdot \frac{SSA}{SST} \right)^{\frac{1-n}{2}} \cdot \pi(\tau) d\tau \quad (10)$$

To use Equation 10, the user must place a prior distribution on the variance components τ . Motivated by the work of Maruyama (2009), Wang and Sun

used a Pearson Type VI distribution to serve as the prior for τ . The Pearson Type VI distribution has three parameters: two shape parameters $\alpha > -1$ and $\beta > -1$ and a scale parameter $\kappa > 0$. The density function for the Pearson Type VI prior is given by

$$\pi^{PT}(\tau) = \frac{\kappa(\kappa\tau)^\beta(1 + \kappa\tau)^{-\alpha-\beta-2}}{\mathcal{B}(\alpha + 1, \beta + 1)} I_{(0,\infty)}(\tau)$$

where $\mathcal{B}(x, y) = \int_0^1 t^{x-1}(1-t)^{y-1}dt$ is the standard Beta function. Wang and Sun further restricted π^{PT} to one parameter $\alpha \in [-\frac{1}{2}, 0]$ by taking $\kappa = r$ and $\beta = \frac{n-p}{2} - \alpha - 2$. A plot of this prior can be seen in Figure 2; in this figure, we take the values $r = 6$, $p = 3$, and $n = 18$ from Example 1, thus setting $\kappa = 6$ and $\beta = \frac{18-3}{2} - \alpha - 2$, where α takes specific values $-\frac{1}{2}, -\frac{1}{4}, -\frac{1}{10}, 0$. Further, Figure 2 shows the effect of varying α on the prior distribution for τ . As we can see, as α increases to 0, τ becomes more dispersed and less peaked around the mode. This places more prior mass on larger treatment effects than we would see for values of α closer to $-\frac{1}{2}$.

Figure 2. A Pearson Type VI prior for τ , plotted as a function of shape parameter α

With the above prior specification, Wang and Sun (2014) proved that the Bayes factor of Garcia-Donato and Sun (Equation 10) simplifies to an

analytic expression without integral representation:

$$\text{BF}_{10} = \frac{\Gamma\left(\frac{p}{2} + \alpha + \frac{1}{2}\right) \cdot \Gamma\left(\frac{n-p}{2}\right)}{\Gamma\left(\frac{n-1}{2}\right) \cdot \Gamma(\alpha + 1)} \cdot \left(\frac{SSR}{SST}\right)^{\alpha - \frac{n-p-2}{2}}. \quad (11)$$

With this equation, the Bayes factor can be computed exactly without any need for numerical integration routines. Note: though the gamma function is itself defined by an integral, its exact values can often be found, particularly for positive integers (i.e., $\Gamma(n) = (n-1)!$ for $n \in \mathbb{N}$). In other cases, extremely precise approximations can be had with most standard programming libraries).

Example 2. Let us now compute a Bayes factor using the Wang and Sun (2014) approach of Equation 11. We will consider the same data set given in Example 1, so we have $n = 18$ and $p = 3$. Equation 11 requires the user to specify $\alpha \in [-\frac{1}{2}, 0]$, which governs the shape of the Pearson Type VI prior on τ , the ratio of variance components σ_a^2/σ^2 . To see the effect of this choice on the resulting Bayes factor, we will consider two cases on the boundary: $\alpha = 0$ and $\alpha = -\frac{1}{2}$.

Case 1: Set $\alpha = 0$. Then by Equation 11

$$\begin{aligned} \text{BF}_{10} &= \frac{\Gamma\left(\frac{3}{2} + 0 + \frac{1}{2}\right) \cdot \Gamma\left(\frac{18-3}{2}\right)}{\Gamma\left(\frac{18-1}{2}\right) \cdot \Gamma(0+1)} \left(\frac{88}{172}\right)^{0 - \frac{18-3-2}{2}} \\ &= \frac{\Gamma(2) \cdot \Gamma\left(\frac{15}{2}\right)}{\Gamma\left(\frac{17}{2}\right) \cdot \Gamma(1)} \left(\frac{88}{172}\right)^{-\frac{13}{2}} = \frac{1 \cdot 1871.254}{14034.41 \cdot 1} \left(\frac{88}{172}\right)^{-\frac{13}{2}} = 10.393. \end{aligned}$$

Case 2: Set $\alpha = -\frac{1}{2}$. Then by Equation 11

$$\begin{aligned} \text{BF}_{10} &= \frac{\Gamma\left(\frac{3}{2} - \frac{1}{2} + \frac{1}{2}\right) \cdot \Gamma\left(\frac{18-3}{2}\right)}{\Gamma\left(\frac{18-1}{2}\right) \cdot \Gamma\left(-\frac{1}{2} + 1\right)} \left(\frac{88}{172}\right)^{-\frac{1}{2} - \frac{18-3-2}{2}} \\ &= \frac{\Gamma\left(\frac{3}{2}\right) \cdot \Gamma\left(\frac{15}{2}\right)}{\Gamma\left(\frac{17}{2}\right) \cdot \Gamma\left(\frac{1}{2}\right)} \left(\frac{88}{172}\right)^{-7} = \frac{0.8862269 \cdot 1871.254}{14034.41 \cdot 1.772454} \left(\frac{88}{172}\right)^{-7} \\ &= 7.265. \end{aligned}$$

Notice that for both values of α , the obtained Bayes factor is quite different from the value obtained from the BIC approximation in Example 1. However, both are well under the Sellke et al. bound of 11.10 (which we calculated earlier).

5. Analytic Bayes factor from summary statistics

We are now ready to prove the primary new result in this paper; namely, the development of the Pearson Bayes factor, which is summarized in Proposition 1. The Pearson Bayes factor provides researchers with an analytic Bayes factor (i.e., without integral representation) that can be obtained directly from the summary statistics typically reported from a t -test or analysis of variance. To do this, we simply need to start with Equation 11 of Wang and Sun (2014). As we previously saw, their Bayes factor requires knowing the residual sum of squares term SSR , the total sum of squares SST , the total number of observations n , and the number of treatment groups p . Additionally, the user must specify a hyperparameter α , which controls the shape of the prior distribution for $\tau = \sigma_a^2/\sigma^2$, the ratio of variance components.

To begin the proof, let us first consider two facts. First, x is the between-treatments degrees of freedom, so $x = p - 1$. This implies $p = x + 1$. Second, y is the residual (within-treatments) degrees of freedom, so $y = n - p$. Putting these two facts together, we have $x + y = (p - 1) + (n - p) = n - 1$. Substituting these into Equation 11 gives us

$$\begin{aligned} \text{BF}_{10} &= \frac{\Gamma\left(\frac{p}{2} + \alpha + \frac{1}{2}\right) \cdot \Gamma\left(\frac{n-p}{2}\right)}{\Gamma\left(\frac{n-1}{2}\right) \cdot \Gamma(\alpha + 1)} \cdot \left(\frac{SSR}{SST}\right)^{\alpha - \frac{n-p-2}{2}} \\ &= \frac{\Gamma\left(\frac{x+1}{2} + \alpha + \frac{1}{2}\right) \cdot \Gamma\left(\frac{y}{2}\right)}{\Gamma\left(\frac{x+y}{2}\right) \cdot \Gamma(\alpha + 1)} \cdot \left(\frac{SSR}{SST}\right)^{\alpha - \frac{y-2}{2}} \\ &= \frac{\Gamma\left(\frac{x}{2} + \alpha + 1\right) \cdot \Gamma\left(\frac{y}{2}\right)}{\Gamma\left(\frac{x+y}{2}\right) \cdot \Gamma(\alpha + 1)} \cdot \left(\frac{SSR}{SST}\right)^{\alpha - \frac{y}{2} + 1}. \end{aligned}$$

The last step is to express the fraction SSR/SST in terms of F , x , and y . To this end, we note that the reciprocal can be written as

$$\frac{SST}{SSR} = \frac{SSR + SSA}{SSR} = 1 + \frac{SSA}{SSR}.$$

By definition, the F ratio is written as $F = \frac{SSA/x}{SSR/y} = \frac{SSA}{SSR} \cdot \frac{y}{x}$. Combining these gives us

$$\frac{SST}{SSR} = 1 + \frac{xF}{y} = \frac{y + xF}{y}.$$

Reciprocating and substituting back into the equation above puts the Bayes factor into the form of Proposition 1, thus completing the proof.

Example 3. Again, we consider the same data set as Example 1, where we found $F(2, 15) = 7.16$. From these summary statistics, we apply Proposition 1 using $F = 7.16$, $x = 2$, and $y = 15$. As before, we consider the two boundary cases for the value of α :

Case 1: $\alpha = 0$. By Proposition 1, we have

$$\begin{aligned} \text{PBF}_{10} &= \frac{\Gamma\left(\frac{2}{2} + 0 + 1\right) \cdot \Gamma\left(\frac{15}{2}\right)}{\Gamma\left(\frac{2+15}{2}\right) \cdot \Gamma(0 + 1)} \left(\frac{15}{15 + 2 \cdot 7.16}\right)^{0 - \frac{15}{2} + 1} \\ &= \frac{1 \cdot 1871.254}{14034.41 \cdot 1} (0.5116)^{-6.5} \\ &= 10.397 \end{aligned}$$

Case 2: $\alpha = -1/2$. Here, we have

$$\begin{aligned} \text{PBF}_{10} &= \frac{\Gamma\left(\frac{2}{2} - \frac{1}{2} + 1\right) \cdot \Gamma\left(\frac{15}{2}\right)}{\Gamma\left(\frac{2+15}{2}\right) \cdot \Gamma\left(-\frac{1}{2} + 1\right)} \left(\frac{15}{15 + 2 \cdot 7.16}\right)^{-\frac{1}{2} - \frac{15}{2} + 1} \\ &= \frac{0.8862269 \cdot 1871.254}{14034.41 \cdot 1.772454} (0.5116)^{-7} \\ &= 7.268. \end{aligned}$$

In both cases, Proposition 1 produces Bayes factors reasonably close to the answers we produced using the Wang and Sun formula (Equation 11) in Example 2. Indeed, the inaccuracy results from the inevitable two-digit rounding that we often see with summary statistics when reported in papers. If we had used $F = 7.159$ instead of $F = 7.16$, our answers would have matched exactly.

Finally, let us consider the case where $p = 2$; that is, we have two treatment conditions. In this case, the between-treatments degrees of freedom is equal to 1. Thus, if we apply Proposition 1 with $x = 1$, we have the following:

$$\begin{aligned} \text{PBF}_{10} &= \frac{\Gamma\left(\frac{x}{2} + \alpha + 1\right) \cdot \Gamma\left(\frac{y}{2}\right)}{\Gamma\left(\frac{x+y}{2}\right) \cdot \Gamma(\alpha + 1)} \left(\frac{y}{y + xF}\right)^{\alpha - \frac{y}{2} + 1} \\ &= \frac{\Gamma\left(\frac{1}{2} + \alpha + 1\right) \cdot \Gamma\left(\frac{y}{2}\right)}{\Gamma\left(\frac{1+y}{2}\right) \cdot \Gamma(\alpha + 1)} \left(\frac{y}{y + F}\right)^{\alpha - \frac{y}{2} + 1} \\ &= \frac{\Gamma\left(\alpha + \frac{3}{2}\right) \cdot \Gamma\left(\frac{y}{2}\right)}{\Gamma\left(\frac{1+y}{2}\right) \cdot \Gamma(\alpha + 1)} \left(\frac{y}{y + F}\right)^{(2\alpha - y + 2)/2}. \end{aligned}$$

If we substitute the identity $F = t^2$ and replace y with ν (the typical symbol used to represent degrees of freedom in a t -test), we get the following:

$$\begin{aligned} \text{PBF}_{10} &= \frac{\Gamma\left(\frac{\nu}{2}\right) \cdot \Gamma\left(\alpha + \frac{3}{2}\right)}{\Gamma\left(\frac{\nu+1}{2}\right) \cdot \Gamma(\alpha + 1)} \left(\frac{\nu}{\nu + t^2}\right)^{(2\alpha - \nu + 2)/2} \\ &= \frac{\Gamma\left(\frac{\nu}{2}\right) \cdot \Gamma\left(\alpha + \frac{3}{2}\right)}{\Gamma\left(\frac{\nu+1}{2}\right) \cdot \Gamma(\alpha + 1)} \left(\frac{\nu + t^2}{\nu}\right)^{(-2\alpha + \nu - 2)/2} \\ &= \frac{\Gamma\left(\frac{\nu}{2}\right) \cdot \Gamma\left(\alpha + \frac{3}{2}\right)}{\Gamma\left(\frac{\nu+1}{2}\right) \cdot \Gamma(\alpha + 1)} \left(1 + \frac{t^2}{\nu}\right)^{(\nu - 2\alpha - 2)/2} \end{aligned}$$

which exactly matches the Pearson Bayes factor derived by Wang and Liu (2016) for the t -test. Thus, our proposed Pearson Bayes factor formula from Proposition 1 serves as a generalization that works for summary statistics from both the t -test and analysis of variance.

6. Simulation study

In this final section, I report the results of a simulation study designed to benchmark the performance of the Pearson Bayes factor against other popular Bayes factors. In the simulation, I generated random datasets that each

reflected the balanced one-factor designs that we have discussed throughout this paper. Specifically, data were generated as

$$Y_{ij} = \mu + a_j + \varepsilon_{ij}; \quad i = 1, \dots, r; \quad j = 1, \dots, p,$$

where μ represents a grand mean, $a_j \sim \mathcal{N}(0, \sigma_a)$ represent each of the p randomly drawn treatment effects, and $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$ represent the normally distributed error terms. For convenience I set $p = 3$, though similar results were obtained with other values of p (not reported here). Also, without loss of generality I set $\mu = 0$ and $\sigma = 1$. I then systematically varied the following three components of the model:

1. The number of replicates per treatment condition r was set to either $n = 10$, $n = 30$, or $n = 80$;
2. The size of the overall treatment effect was manipulated by setting $\tau = \sigma_a^2 / \sigma^2$ to be either $\tau = 0$, $\tau = 0.5$, or $\tau = 1$. Note that data generated under the condition $\tau = 0$ reflect the null model $\mathcal{H}_0$, whereas data generated under $\tau > 0$ reflect the alternative model $\mathcal{H}_1$.
3. The prior setting α used in the computation of the Pearson Bayes factor was set at both ends of its consistency range; that is, $\alpha = -1/2$ and $\alpha = 0$.

For each combination of number of replicates ($r = 10, 30, 80$), treatment effect ($\tau = 0, 0.5, 1.0$), and prior setting α ($\alpha = -0.5, 0$), I generated 1000 simulated datasets. For each of the datasets, I computed three different Bayes factors: the Pearson Bayes factor proposed in this paper, the JZS Bayes Factor (Rouder et al., 2012), and the BIC Bayes factor (Faulkenberry, 2018). These Bayes factors were then converted to posterior probabilities via Equation 3. To compare the performance of the three Bayes factors in the simulation, I considered three analyses for each simulated dataset: (1) a visualization of the distribution of posterior probabilities $P(\mathcal{H}_0 \mid \mathcal{D})$; (2) a calculation of the proportion of simulated trials for which the correct model was chosen (i.e., model choice accuracy); and (3) a calculation of the proportion of simulated trials for which the Pearson Bayes factor chose the same model as the JZS Bayes factor (i.e., model choice consistency).

To begin, let us consider the distribution of posterior probabilities $P(\mathcal{H}_0 \mid \mathcal{D})$. Boxplots of these posterior probabilities can be seen in Figure 3. From the figure, we can see that the Pearson Bayes factor performs well compared

to the JZS and BIC Bayes factors, but there are some distinct differences. First, when $\tau = 0$ (and thus, $\mathcal{H}_0$ is the correct model) both the Pearson and BIC Bayes factors produce larger posterior probabilities for $\mathcal{H}_0$ than the JZS Bayes factor. When the prior parameter α is set to $-1/2$, the Pearson Bayes factor produce a smaller range of posterior probabilities than the BIC Bayes factor. When $\tau > 0$ ($\mathcal{H}_1$ is the correct model), the overall advantage goes to the JZS Bayes factor, though the distribution of posterior probabilities becomes indistinguishable across methods as τ increases and r increases. Note that in the $\tau = 0.5$ case (where treatment effect variance is $1/2$ of residual variance), model selection is a bit more scattered. In this case, the Pearson Bayes factor outperforms (i.e., has smaller range and smaller posterior probability for $\mathcal{H}_0$) the BIC Bayes factor when the prior parameter α is set to 0. This is not surprising, as the Pearson Type VI prior for τ is more dispersed when $\alpha = 0$, allowing for a greater prior mass being placed on larger effects.

Though the distributions of posterior probabilities tend to follow similar patterns, it is not clear to what extent the three methods provide the user with an accurate inference. To compare the accuracy of the three methods, I calculated *model choice accuracy*, defined as the proportion of simulated datasets for which the correct model was chosen. Model choice was defined by considering $\mathcal{H}_0$ to be chosen whenever $\text{BF}_{01} > 1$ and $\mathcal{H}_1$ to be chosen whenever $\text{BF}_{01} < 1$. The results are displayed in Table 1. From the table, we can see confirmation of the intuition gained from the distribution of posterior probabilities in Figure 3. Again, for null effects, both the PBF and the BIC consistently outperform the JZS Bayes factor (particularly when the prior shape parameter is set to $\alpha = -1/2$). When $\tau > 0$, the JZS Bayes factor chooses the correct model ($\mathcal{H}_1$) more often than either the PBF or the BIC. Again, the PBF outperforms the BIC when setting $\alpha = 0$. The accuracy of both PBF and BIC increases as the number of replicates per treatment condition increases.

Finally, I assessed the degree to which the model chosen by the Pearson Bayes factor matches the model chosen by the JZS Bayes factor. This model choice consistency is displayed in Table 2. Overall, model choice consistency is very good (all above 84%). Certainly, consistency improves as the number of replicates per treatment condition increases. However, it is also notable that consistency improves across the board when α is set to 0.

Figure 3. Results from our simulation. Each boxplot depicts the distribution of the posterior probability $P(\mathcal{H}_0 \mid \mathcal{D})$ for 1000 Monte Carlo simulations. White boxes represent posterior probabilities derived from the Pearson Bayes factor. Medium gray boxes represent posterior probabilities derived from the JZS Bayes factor. Dark gray boxes represent posterior probabilities derived from the BIC Bayes factor.

Table 1. Model choice accuracies for the Pearson Bayes factor (PBF), JZS Bayes factor (JZS), and BIC Bayes factor (BIC), calculated as the proportion of simulated datasets for which the correct model was chosen

	$\alpha = -1/2$			$\alpha = 0$		
	PBF	JZS	BIC	PBF	JZS	BIC
$\tau = 0$:						
$r = 10$	.980	.913	.965	.931	.886	.978
$r = 30$	.983	.933	.982	.978	.930	.982
$r = 80$	.995	.973	.995	.987	.966	.994
$\tau = 0.5$:						
$r = 10$	.497	.655	.551	.584	.652	.555
$r = 30$	.732	.824	.742	.771	.820	.733
$r = 80$	.872	.914	.872	.892	.913	.874
$\tau = 1.0$:						
$r = 10$	.732	.830	.765	.751	.790	.727
$r = 30$	.855	.908	.859	.868	.900	.856
$r = 80$	.932	.956	.933	.945	.959	.938

Table 2. Model choice consistency for the PBF Bayes factor and the JZS Bayes factor, calculated as the proportion of simulated datasets for which both methods chose the same model

	$\tau = 0$	$\tau = 0.5$	$\tau = 1.0$
$\alpha = -1/2$			
$r = 10$	.933	.842	.902
$r = 30$	.950	.908	.947
$r = 80$	.978	.958	.976
$\alpha = 0$			
$r = 10$	.955	.932	.961
$r = 30$	.952	.951	.968
$r = 80$	.979	.979	.986

7. Conclusion

In this paper, I have developed an analytic Bayes factor (the Pearson Bayes factor) which allows a researcher to obtain Bayes factors using the minimal summary statistics (e.g., the test statistic and the degrees of freedom) from both t -tests and analysis of variance designs. This formula improves upon the BIC Bayes factor formula of Faulkenberry (2018) not only by providing the user with an *exact* Bayes factor instead of an approximation, but also

by giving the user the ability to specify prior knowledge about the size of effects expected. Our simulation study shows that the Pearson Bayes factor performs comparably to the JZS Bayes factor of Rouder et al. (2012) (especially when the prior parameter α is set to 0). Also, the Pearson Bayes factor outperforms JZS on data generated under a null model. The Pearson Bayes factor did not perform as well as the JZS Bayes on data generated under the assumption of nonzero treatment effects, but this limitation is outweighed by the Pearson Bayes factor's unique ability to be computed from summary statistics alone and without the need to compute a multi-dimensional integral. For this reason, the Pearson Bayes factor is an excellent tool for researchers who wish to assess the evidential value of data from published studies in which raw data is not readily available. In all, the Pearson Bayes factor should be a valuable tool in the scientist's statistical toolbox.

REFERENCES

- Faulkenberry T.J. (2018): Computing Bayes factors to measure evidence from experiments: An extension of the BIC approximation. *Biometrical Letters* 55(1): 31–43.
- Faulkenberry T.J. (2019a): Estimating Bayes factors from minimal ANOVA summaries for repeated-measures designs. [arXiv:1905.05569](https://arxiv.org/abs/1905.05569).
- Faulkenberry T.J. (2019b): Estimating evidential value from analysis of variance summaries: A comment on Ly (2018). *Advances in Methods and Practices in Psychological Science* 2(4): 406–409.
- Fisher R.A. (1925): *Statistical Methods for Research Workers*. Oliver & Boyd, Edinburgh.
- García-Donato G., Sun D. (2007): Objective priors for hypothesis testing in one-way random effects models. *Canadian Journal of Statistics*, 35(2): 303–320.
- Gönen M., Johnson W.O., Lu Y., Westfall P.H. (2005): The Bayesian two-sample t test. *The American Statistician*, 59(3): 252–257.
- Jeffreys H. (1961): *The Theory of Probability* (3rd ed.). Oxford University Press, Oxford, UK.
- Kass R.E., Raftery A.E. (1995): Bayes factors. *Journal of the American Statistical Association*, 90(430): 773.
- Liang F., Paulo R., Molina G., Clyde M.A., Berger J.O. (2008): Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103(481): 410–423.
- Maruyama Y. (2009): A Bayes factor with reasonable model selection consistency for ANOVA model. [arXiv:0906.4329v2](https://arxiv.org/abs/0906.4329v2).
- Masson M.E.J. (2011): A tutorial on a practical Bayesian alternative to null-hypothesis significance testing. *Behavior Research Methods*, 43(3): 679–690.
- Raftery A.E. (1995): Bayesian model selection in social research. *Sociological Methodology*, 25: 111–163.

- Rouder J.N., Engelhardt C.R., McCabe S., Morey R.D. (2016): Model comparison in ANOVA. *Psychonomic Bulletin & Review*, 23(6): 1779–1786.
- Rouder J.N., Morey R.D., Speckman P.L., Province J.M. (2012): Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56(5): 356–374.
- Schwarz G. (1978): Estimating the dimension of a model. *The Annals of Statistics*, 6(2): 461–464.
- Sellke T., Bayarri M.J., Berger J.O. (2001): Calibration of p -values for testing precise null hypotheses. *The American Statistician*, 55(1): 62–71.
- Student (1908): The probable error of a mean. *Biometrika*, 6(1):1.
- Thiele J.E., Haaf J.M., Rouder J.N. (2017): Is there variation across individuals in processing? Bayesian analysis for systems factorial technology. *Journal of Mathematical Psychology*, 81: 40–54.
- Wagenmakers E.J. (2007): A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5): 779–804.
- Wang M., Liu G. (2016): A simple two-sample Bayesian t -test for hypothesis testing. *The American Statistician*, 70(2): 195–201.
- Wang M., Sun X. (2014): Bayes factor consistency for one-way random effects model. *Communications in Statistics - Theory and Methods*, 43(23): 5072–5090.
- Wasserstein R.L., Lazar N.A. (2016): The ASA statement on p -values: Context, process, and purpose. *The American Statistician*, 70(2): 129–133.
- Zellner A. (1986): On assessing prior distributions and Bayesian regression analysis with g -prior distributions. In Goel, P. K. and Zellner, A., editors, *Bayesian inference and decision techniques: Essays in Honor of Bruno de Finetti*, pages 233–243. Elsevier.