Closed form Bayes factor techniques for measuring evidential value from analysis of variance models

Thomas J. Faulkenberry

Tarleton State University

JMM 2020

The goal of this talk is to describe some methods for evaluating evidential value of data in analysis of variance models.

By *evidential value*, I mean the factor by which the prior odds is updated after observing data:

$$\underbrace{\frac{p(\mathcal{H}_1 \mid \mathcal{D})}{p(\mathcal{H}_0 \mid \mathcal{D})}}_{\text{posterior odds}} = \underbrace{\frac{p(\mathcal{H}_1)}{p(\mathcal{H}_0)}}_{\text{prior odds}} \times \underbrace{\frac{p(\mathcal{D} \mid \mathcal{H}_1)}{p(\mathcal{D} \mid \mathcal{H}_0)}}_{\text{predictive updating factor}}$$

Kass and Raftery (1995) called this predictive updating factor the Bayes factor

Motivating example: consider test scores from students in three instructional treatments:

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

Typical question – are there differences among these condition means?

Classical approach - analysis of variance (ANOVA)

- model $Y_{ij} = \mu + \alpha_j + \varepsilon_{ij}$, where $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$
- assume "null hypothesis" $\mathcal{H}_0 : \alpha_j = 0$
- compute probability of observing data Y_{ij} under $\mathcal{H}_0$
- if data is $rare$ under $\mathcal{H}_0$, reject $\mathcal{H}_0$
- index "rareness" by computing F (a ratio of between-group and within-group variances)

ANOVA computations

source	SS	df	MS	F
between treatments	84	2	42	7.16
within treatments	88	15	5.87	
total	172	17		

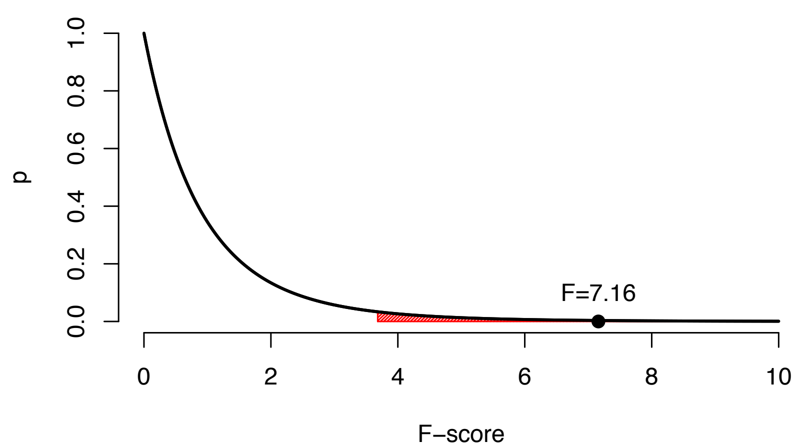

Since our data Y_{ij} is rare under $\mathcal{H}_0$, we reject $\mathcal{H}_0$ as an implausible model restriction.

This ignores predictive adequacy of $\mathcal{H}_1$.

The **Bayes factor**

$$B_{10} = \frac{p(\mathcal{D} \mid \mathcal{H}_1)}{p(\mathcal{D} \mid \mathcal{H}_0)}$$

tells us how much better $\mathcal{H}_1$ predicts our observed data compared to $\mathcal{H}_0$.

Example: suppose $B_{12} = 5$. This means that the observed data are 5 times more likely under $\mathcal{H}_1$ than $\mathcal{H}_0$.

Problem – computing Bayes factors is hard!

$$\begin{aligned} B_{10} &= \frac{p(\mathcal{D} \mid \mathcal{H}_1)}{p(\mathcal{D} \mid \mathcal{H}_0)} \\ &= \frac{\int f(\mathcal{D} \mid \theta_1, \mathcal{H}_1)\pi(\theta_1 \mid \mathcal{H}_1) d\theta_1}{\int f(\mathcal{D} \mid \theta_0, \mathcal{H}_0)\pi(\theta_0 \mid \mathcal{H}_0) d\theta_0} \end{aligned}$$

Today, I'll describe two approaches to making this computation easier

1. BIC approximation (Raftery, 1995; Wagenmakers, 2007; Masson, 2011)

Basic idea – if we construct 2nd order Taylor approximation of log-marginal likelihood of each $\mathcal{H}_i$, we get

$$BF_{01} \approx \exp\left(\frac{\mathrm{BIC}(\mathcal{H}_1) - \mathrm{BIC}(\mathcal{H}_0)}{2}\right),$$

where

$$\mathrm{BIC}(\mathcal{H}_i) = n \ln\left(\frac{SSR}{SST}\right) + k \ln n.$$

In 2018[1], I derived a simple formula that computes this BIC Bayes factor using only the ANOVA summary statistics:

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

where

- F is the observed F-ratio
- x, y are the numerator/denominator df, respectively
- n is the total number of observations.

[1] Faulkenberry, T.J. (2018). Computing Bayes factors to measure evidence from experiments: An extension of the BIC approximation. *Biometrical Letters, 55*, 31-43

Using the summary statistics from the ANOVA ($F = 7.16$, $x = 2$, $y = 15$, and $n = 18$), we get

$$
\begin{aligned}
BF_{01} &\approx \sqrt{n^x \left(1 + \frac{Fx}{y}\right)^{-n}} \\
&= \sqrt{18^2 \left(1 + \frac{7.16 \cdot 2}{15}\right)^{-18}} \\
&= 0.0432,
\end{aligned}
$$

Thus, $BF_{10} = 1/BF_{01} \approx 1/0.0432 = 23.15$

Bad approximation?

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

$$\begin{aligned} BF_{10} &\leq -\frac{1}{e \cdot p \ln(p)} \\ &= \leq -\frac{1}{e \cdot 0.0066 \cdot \ln(0.0066)} \\ &= 11.10. \end{aligned}$$

Thus, our BIC Bayes factor of 23.15 is quite an overestimate of the actual Bayes factor. **Can we compute an exact Bayes factor with only summary statistics?**

2. Pearson Type VI Bayes factor

In some new work[2], I derived the following $exact$ Bayes factor:

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

[2] based on Wang & Sun (2014); preprint forthcoming on ArXiV

Using our example, we have

$$
\begin{aligned}
BF_{10} &= \frac{\Gamma\left(\frac{2+1}{2}\right)\cdot\Gamma\left(\frac{15}{2}\right)}{\Gamma\left(\frac{2+15}{2}\right)\cdot\Gamma\left(\frac{1}{2}\right)}\left(\frac{15}{15+2\cdot 7.16}\right)^{\frac{1-15}{2}} \\
&= \frac{0.8662269\cdot 1871.254}{14034.41\cdot 1.772454}(0.5116)^{-7} \\
&= 7.268
\end{aligned}
$$

Some observations:

- the resulting Bayes factor is reasonable w.r.t. the Sellke bound
- the expression is *analytic*, but not *closed form*

Theorem 1. *Given an ANOVA summary reported in standard form $F(x, y)$ (i.e., where x equals the between-treatments degrees of freedom and y equals the residual degrees of freedom), the Bayes factor can be expressed in closed form as*

$$BF_{10} = C\sqrt{\left(\frac{y}{y + xF}\right)^{1-y}}$$

where C depends on the parity of x and y, as follows:

Case 1: if x and y are even, then

$$C = \frac{x!\left(\frac{y}{2} - 1\right)!}{2^x\left(\frac{x}{2}\right)!\left(\frac{x+y}{2} - 1\right)!}$$

Case 2: if x is even and y is odd, then

$$C = \frac{x!(y-1)!\left(\frac{x+y-1}{2}\right)!}{\left(\frac{x}{2}\right)!\left(\frac{y-1}{2}\right)!(x+y-1)!}$$

Case 3: if x is odd and y is even, then

$$C = \frac{2^{x+y-1}\left(\frac{x+y-1}{2}\right)!\left(\frac{y}{2}-1\right)!\left(\frac{x-1}{2}\right)!}{\pi(x+y-1)!}$$

Case 4: if x and y are odd, then

$$C = \frac{\left(\frac{x-1}{2}\right)!(y-1)!}{2^{y-1}\left(\frac{y-1}{2}\right)!\left(\frac{x+y}{2}-1\right)!}$$

Future goals

In a new paper[3], I describe a web app that computes Bayes factors from ANOVA summaries – https://tomfaulkenberry.shinyapps.io/anovaBFcalc

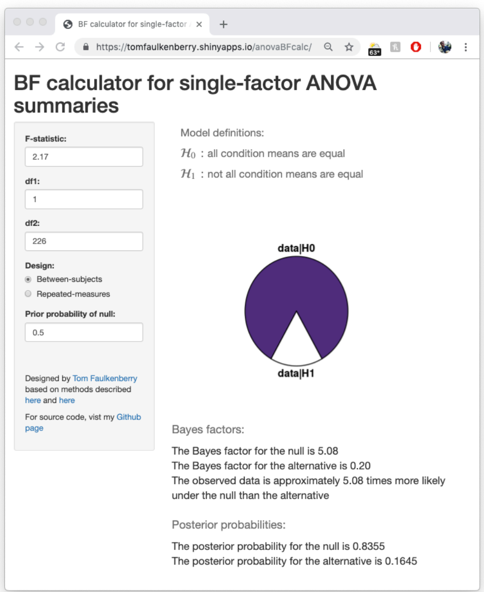

- "Shiny app" (based on R)
- gives Bayes factors and posterior probabilities
- currently uses BIC approximation
- working on integrating exact Bayes factors

[3] Faulkenberry, T. J. (2019). Estimating evidential value from analysis of variance summaries: A comment on Ly et al.(2018). *Advances in Methods and Practices in Psychological Science*, *2*(4), 406-409

Thank you!

- Twitter: `@tomfaulkenberry`
- Email: `faulkenberry@tarleton.edu`