SWPA Workshop

Getting started in Bayesian Statistics with JASP

Tom Faulkenberry

Tarleton State University
Stephenville, TX, USA
https://tomfaulkenberry.github.io
Twitter: @tomfaulkenberry

SWPA 2022

Outline:

- discuss differences between p-values & Bayes factors
- priors on models vs. priors on parameters
- correlation example using JASP, w/ reporting template.
- more resources!

* These slides can be downloaded from
https://tomfaulkenberry.github.io/talks.html

Suppose we are interested in the relationship between math anxiety and performance on a standardized assessment.

Define hypotheses about (population) correlation ρ

$\mathcal{H}_0 : \rho = 0$, $\mathcal{H}_1 : \rho \neq 0$

Collect data

Compute:

$p(\text{data} \mid \mathcal{H}_0)$

"p-value"

Interpretation:

If p is small, data is <u>rare</u> under $\mathcal{H}_0$, so we reject $\mathcal{H}_0$ in favor of $\mathcal{H}_1$.

Compute:

$$BF_{01} = \frac{p(\text{data} \mid \mathcal{H}_0)}{p(\text{data} \mid \mathcal{H}_1)}$$

"Bayes factor"

Interpretation:

if $BF_{01} > 1$, data more likely under $\mathcal{H}_0$.

if $BF_{01} < 1$, data more likely under $\mathcal{H}_1$.

p-value = $p(\text{data} \mid \mathcal{H}_0)$

1) only considers fit of $\mathcal{H}_0$ as a potential model for data
2) ignores fit of $\mathcal{H}_1$

Thus, "support" for $\mathcal{H}_1$ is only indirect

Bayes factor = $\dfrac{p(\text{data} \mid \mathcal{H}_0)}{p(\text{data} \mid \mathcal{H}_1)}$

1) considers relative adequacy of both models as predictors of data.

2) can directly index support for either $\mathcal{H}_0$ or $\mathcal{H}_1$.

Ex: $BF_{01} = 8$ → "The observed data are 8 times more likely under $\mathcal{H}_0$ than $\mathcal{H}_1$."

Jeffreys (1961):

BF	Evidence*
1 – 3	anecdotal
3 - 10	moderate
10 - 30	strong
30 - 100	very strong
> 100	extreme

* these are only guidelines!

How does Bayes work?

for single model $\mathcal{H}$:

$$\boxed{p(\mathcal{H} \mid \text{data}) = p(\mathcal{H}) \times \frac{p(\text{data} \mid \mathcal{H})}{p(\text{data})}}$$

posterior belief in $\mathcal{H}$ = prior belief in $\mathcal{H}$ × updating factor

for two models:

$$\boxed{\frac{p(\mathcal{H}_0 \mid \text{data})}{p(\mathcal{H}_1 \mid \text{data})} = \frac{p(\mathcal{H}_0)}{p(\mathcal{H}_1)} \times \frac{p(\text{data} \mid \mathcal{H}_0)}{p(\text{data} \mid \mathcal{H}_1)}}$$

posterior odds = prior odds × Bayes factor

What do we mean by prior?

Two types of "priors":

1) priors on models

2) priors on parameters within a given model

① Priors on models — before observing data, what is relative likelihood of competing models?

- common default: $p(\mathcal{H}_0) = p(\mathcal{H}_1) = 1/2$

  ↳ i.e., "1-1 prior odds"

- these prior model probabilities must add to 1

  ↳ $p(\mathcal{H}_0) + p(\mathcal{H}_1) = \frac{1}{2} + \frac{1}{2} = 1$

- prior model probabilities are updated after observing data:

$$\boxed{p(\mathcal{H}_0 \mid \text{data}) = \frac{BF_{01} \cdot p(\mathcal{H}_0)}{BF_{01} \cdot p(\mathcal{H}_0) + p(\mathcal{H}_1)}}$$

* Note: if $p(\mathcal{H}_0) = p(\mathcal{H}_1) = \frac{1}{2}$,

$$\boxed{p(\mathcal{H}_0 \mid \text{data}) = \frac{BF_{01}}{BF_{01} + 1}}$$

Example:

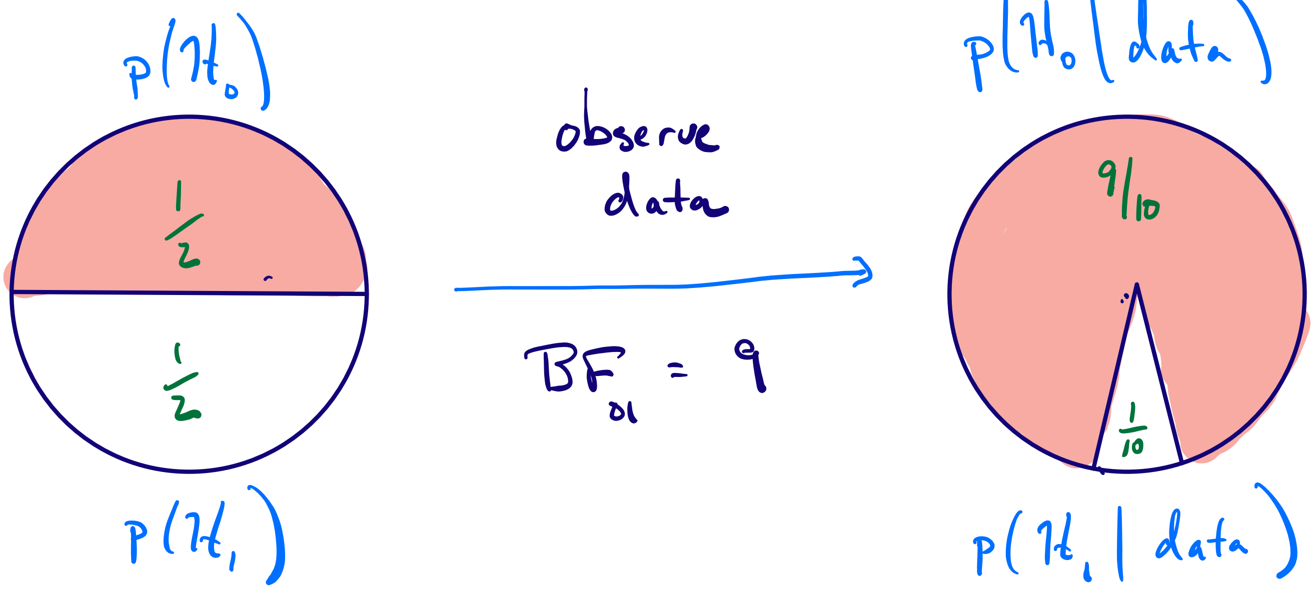

Prior odds = 1 : 1

Posterior odds = 9 : 1

$$\boxed{\begin{aligned} {}^{*}\,p(\mathcal{H}_0 \mid \text{data}) &= \frac{BF_{01}}{BF_{01}+1} \\ &= \frac{9}{9+1} \\ &= 0.9 \end{aligned}}$$

② priors on parameters within a given model.

- model definitions : $\mathcal{H}_0 : \rho = 0$

  $\mathcal{H}_1 : \rho \neq 0$ ← what exactly do we mean here?

- we quantify our uncertainty about the correlation ρ under $\mathcal{H}_1$ by placing a distribution on ρ

- suppose we have no idea what to expect. Here, we might believe any value of ρ is equally likely to occur.

  ↳ we say ρ is uniformly distributed on $(-1, 1)$

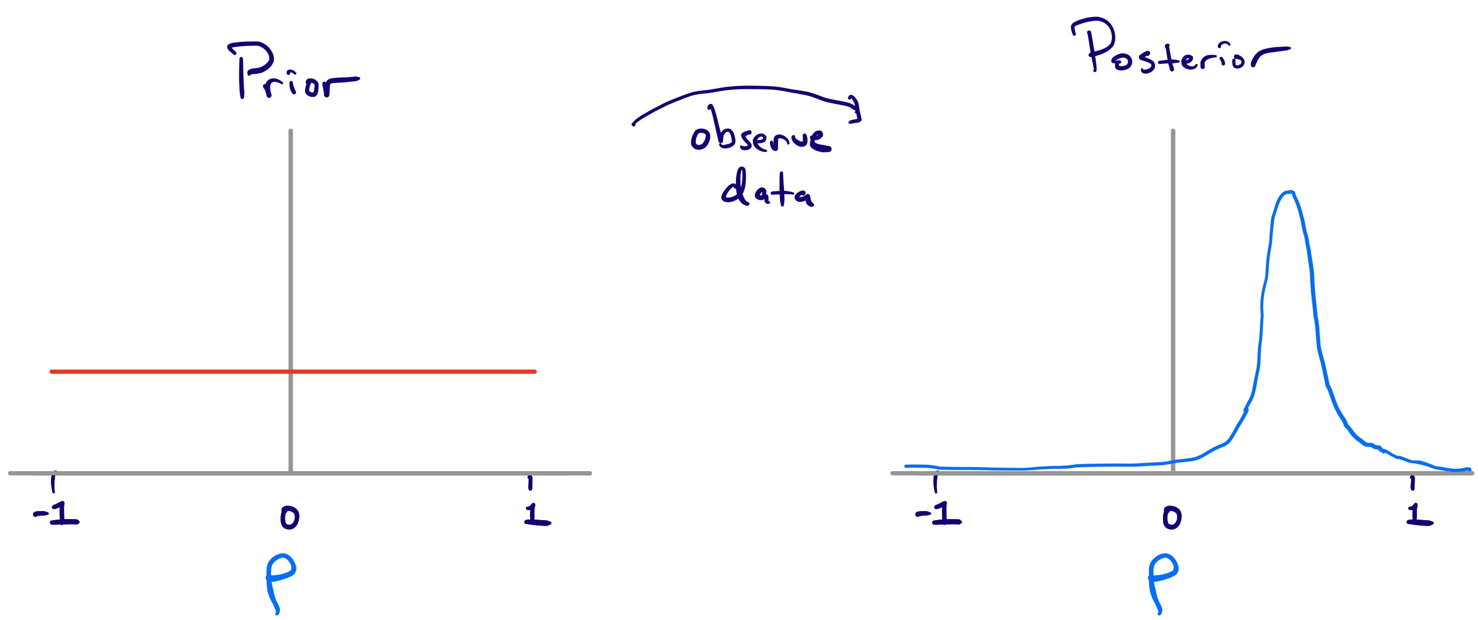

Let's continue our working example. Suppose we tested $N = 65$ participants and observed a correlation of $r = 0.37$.

- use JASP "Summary Statistics" module

Elements to report:

1. report results of hypothesis test
   - define $\mathcal{H}_0$, $\mathcal{H}_1$ and specify prior under $\mathcal{H}_1$.

     "Under the null hypothesis we expect a correlation of 0 between maths anxiety and performance. Thus, we define $\mathcal{H}_0: \rho = 0$. The alternative hypothesis is two-sided, $\mathcal{H}_1: \rho \neq 0$, and we assigned a uniform prior probability to all values of ρ between -1 and $+1$."

   - report and interpret Bayes factor

     "We found a Bayes factor of $BF_{10} = 13.93$, which means that the observed data are approximately 14 times more likely under $\mathcal{H}_1$ than $\mathcal{H}_0$. This result indicates strong evidence in favor of $\mathcal{H}_1$"

- (optional) calculate and report posterior model probability for preferred model.

  - from earlier,

$$p(\mathcal{H}_1 \mid \text{data}) = \frac{BF_{10}}{BF_{10} + 1}$$

$$= \frac{13.93}{13.93 + 1} = 0.93.$$

  - "Assuming prior odds of 1-1 for $\mathcal{H}_1$ and $\mathcal{H}_0$, our observed data updated these odds to 13.93-to-1 in favor of $\mathcal{H}_1$. This is equivalent to a posterior model probability of $p(\mathcal{H}_1 \mid \text{data}) = 0.93$."

2. report results of parameter estimation
  - only if $\mathcal{H}_1$ is the preferred model!
  - specify parameter of interest and remind reader of prior under $\mathcal{H}_1$

    - "Of interest is the posterior distribution for ρ, the population-level correlation between maths anxiety and performance. Under $\mathcal{H}_1$, ρ was assigned a uniform prior over the interval from -1 to $+1$."

- report the 95% credible interval.
  - "The posterior distribution for ρ had a median of 0.356, with a central 95% credible interval that ranges from 0.134 to 0.554."

Let's do another example. Suppose we tested $N = 175$ participants and observed a correlation of $r = 0.15$.

- ???
- frequentist p-value $\longrightarrow$ support for $\mathcal{H}_1$, but Bayes factor $\longrightarrow$ support for $\mathcal{H}_0$!

  ↳ Jeffreys–Lindley Paradox

BayesianSpectacles
Powered by JASP

MENU

Everything You Always Wanted to Know About the Jeffreys-Lindley Paradox But Were Afraid to Ask

JANUARY 24 - 2022

This post is a teaser for Wagenmakers, E.-J., & Ly, A. (2020). History and nature of the Jeffreys-Lindley paradox. Preprint available on ArXiv: https://arxiv.org/abs/2111.10191

More resources

* van Doorn et al. (2021). The JASP guidelines for conducting and reporting a Bayesian analysis. Psychonomic Bulletin & Review

* Faulkenberry, Ly, & Wagenmakers (2020). Bayesian inference in numerical cognition: A tutorial using JASP. Journal of Numerical Cognition

JASP

DOWNLOAD | FEATURES | SUPPORT | TEACHING | BLOG | DONATE

How to do Bayesian Linear Regression in JASP – A Case Study on Teaching Statistics

NOVEMBER 26 - 2020

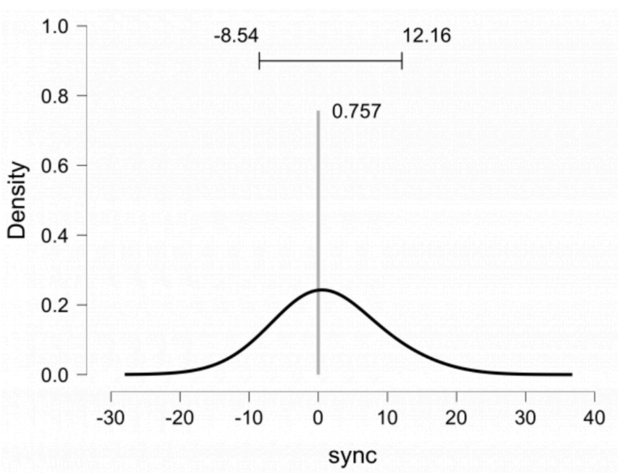

This is a guest post by Tom Faulkenberry (Tarleton State University). Click here to access the supplementary materials.

Before you go – a couple of shameless plugs!

1) My new book, covering the basics of psychological statistics from both the frequentist & Bayesian perspectives.

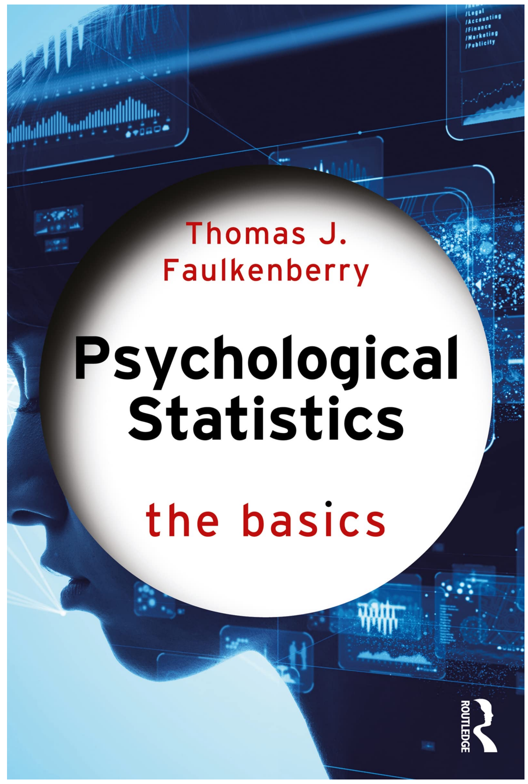

2) PsyStat – a free online Bayes factor calculator

https://tomfaulkenberry.shinyapps.io/psystat

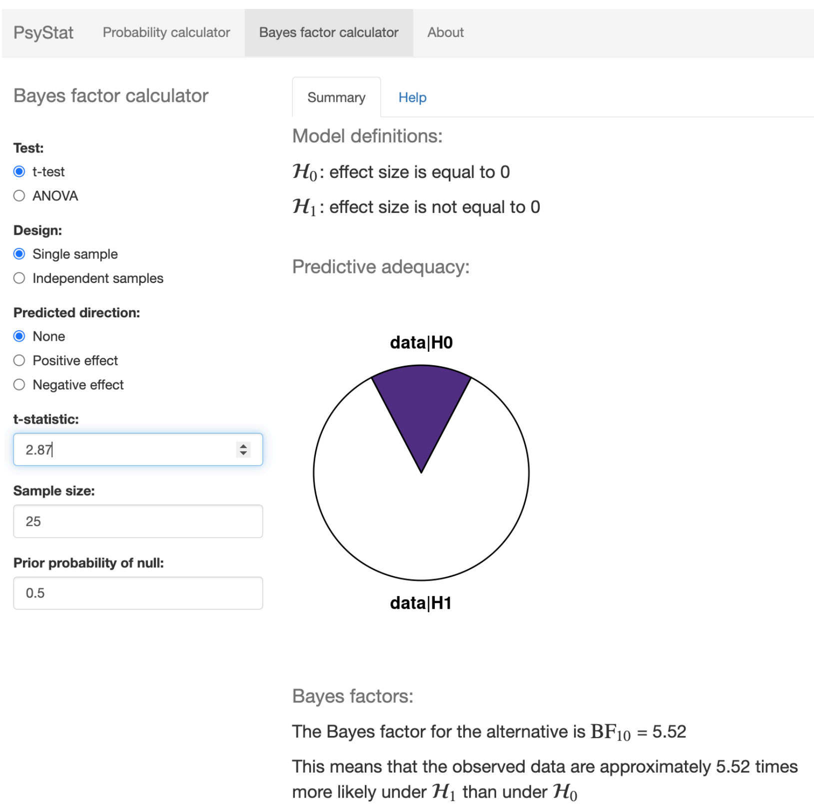

Take home points:

- Bayes is easy, especially with the right software.
- Bayes answers the questions you thought you were asking
- testing or estimation? No need to choose — Bayes gives you both!

More questions — contact me!

Email: faulkenberry@tarleton.edu

Twitter: @tomfaulkenberry

Web: https://tomfaulkenberry.github.io