Online Appendix for: "A Cautionary Note on Estimating Effect Size"

Don van den Bergh[1], Julia M. Haaf[1], Alexander Ly[1,2], Jeffrey N. Rouder[3], and Eric-Jan Wagenmakers[1]

[1]University of Amsterdam

[2]Centrum Wiskunde & Informatica

[3]University of California Irvine

Posterior Distribution for Effect Size under the Spike-and-Slab Model

The main text featured a paired samples t-test, both for the example and for the demonstration of regularities regarding the prior probability of the spike and the prior width of the slab. In this online Appendix we detail the prior distributions for this t-test and explain how the spike-and-slab shrinkage is related to Bayes factors. More generally, we show to derive the posterior distribution for effect size δ under the spike-and-slab model. We first derive the results for the slab and spike individually and combine them afterwards.

Following Rouder et al. (2018), we assume that the observed differences between the paired samples, denoted Z_i, are normally distributed with unknown mean δ and a variance of 1. As prior distribution for δ we use a normal distribution with mean 0 and variance σ^2. This implies $Z_i \sim \mathcal{N}(\delta, 1)$ for the data and $\delta \sim \mathcal{N}\left(0, \sigma^2\right)$ for the prior. The posterior distribution for δ is obtained through Bayes' theorem:

$$\underbrace{p\left(\delta \mid \boldsymbol{Z}\right)}_{\substack{\text{Posterior} \\ \text{distribution}}} = \underbrace{p\left(\delta\right)}_{\substack{\text{Prior} \\ \text{distribution}}} \times \frac{\overbrace{L(\boldsymbol{Z} \mid \delta)}^{\text{Likelihood}}}{\underbrace{p\left(\boldsymbol{Z}\right)}_{\substack{\text{Marginal} \\ \text{Likelihood}}}}.$$

The likelihood is given by:

$$\begin{aligned} L(\boldsymbol{Z} \mid \delta, \text{slab}) &= \prod_{i=1}^{N} \mathcal{N}\left(Z_i \mid \delta, 1\right) \\ &= (2\pi)^{-\frac{N}{2}} \exp\left(-\frac{N}{2}\left(\bar{\boldsymbol{Z}} + s_{\boldsymbol{Z}}^2 + \delta^2 - 2\bar{\boldsymbol{Z}}\delta\right)\right), \end{aligned}$$

where $\bar{\boldsymbol{Z}}$ and $s_{\boldsymbol{Z}}^2$ are the sample mean and sample variance of Z_i respectively. Next, we compute the marginal likelihood by integrating out the likelihood times prior with respect to δ:

$$\begin{aligned} p\left(\boldsymbol{Z} \mid \text{slab}\right) &= \int_{-\infty}^{\infty} L(\boldsymbol{Z} \mid \delta)\, p\left(\boldsymbol{Z}\right) \ \mathrm{d}\delta \\ &= (2\pi)^{-\frac{N+1}{2}} \exp\left(-\frac{N}{2}\left(\bar{\boldsymbol{Z}} + s_{\boldsymbol{Z}}^2\right)\right), \\ &\quad \times \int_{-\infty}^{\infty} \exp\left(-\frac{1}{2}\left(\delta^2\left(N + \frac{1}{\sigma^2}\right) - \delta\frac{2N\bar{\boldsymbol{Z}}}{\sigma^2}\right)\right) \ \mathrm{d}\delta. \end{aligned}$$

Here we may recognize a Gaussian integral and use the following identity:

$$\int_{-\infty}^{\infty} \exp\left(-ax^2 + bx + c\right) \ \mathrm{d}x = \sqrt{\frac{\pi}{a}} \exp\left(\frac{b^2}{4a} + c\right).$$

Filling in the identity and simplifying yields:

$$p\left(\boldsymbol{Z} \mid \text{slab}\right) = (2\pi)^{-\frac{N}{2}} \exp\left(-\frac{N}{2}\left(\bar{\boldsymbol{Z}} + s_{\boldsymbol{Z}}^2\right)\right) \frac{\exp\left(\frac{N^2\bar{\boldsymbol{Z}}^2}{2\left(N + \frac{1}{\sigma^2}\right)}\right)}{\sqrt{N + \frac{1}{\sigma^2}}}.$$

Next, we can obtain an expression for the posterior distribution. However, often it suffices to write out the expression for the likelihood times prior and then identify the result as a known distribution. This is particularly common in Gibbs sampling where one is interested in the conditional posterior distributions. We also do this here, as it reduces inference about the posterior distribution (e.g., what is the mean or variance) to inference about a known distribution, in this case a normal distribution:

$$
\begin{aligned}
p\left(\delta \mid \boldsymbol{Z}, \text{slab}\right) \propto (2\pi)^{-\frac{N}{2}} \exp\left(-\frac{N}{2}\left(\bar{\boldsymbol{Z}} + s^2_{\boldsymbol{Z}}\right)\right) \exp\left(-\frac{N}{2}\left(\delta^2 - 2\bar{\boldsymbol{Z}}\delta\right)\right) \\
\times (2\pi)^{-\frac{1}{2}} \exp\left(-\frac{1}{2\sigma^2}\delta^2\right) \\
\propto \exp\left(-\frac{1}{2}\left(\delta^2\left(N + \frac{1}{\sigma^2}\right) - \delta\frac{2N\bar{\boldsymbol{Z}}}{\sigma^2}\right)\right).
\end{aligned}
$$

We recognize a normal distribution with variance $\sigma_1^2 = \frac{1}{N+\frac{1}{\sigma^2}}$ and mean $\mu_1 = N\bar{\boldsymbol{Z}}\sigma_1^2$. Thus we have $p\left(\delta \mid \boldsymbol{Z}\right) \propto \mathcal{N}\left(\mu_1, \sigma_1^2\right)$.

Next we compute the same for the spike. The spike states that $Z_i \sim \mathcal{N}\left(0, 1\right)$ and contains no parameters to estimate. Thus there are no prior distributions to specify and all that needs to be computed is the marginal likelihood:

$$
p\left(\boldsymbol{Z} \mid \text{spike}\right) = (2\pi)^{-\frac{N}{2}} \exp\left(-\frac{N}{2}\left(\bar{\boldsymbol{Z}} + s^2_{\boldsymbol{Z}}\right)\right).
$$

Using both marginal likelihoods we can now obtain the Bayes factor in favor of the spike:

$$
\text{BF}_{01} = \frac{p\left(\boldsymbol{Z} \mid \text{spike}\right)}{p\left(\boldsymbol{Z} \mid \text{slab}\right)} = \frac{\sqrt{N + \frac{1}{\sigma^2}}}{\exp\left(\frac{N^2\bar{\boldsymbol{Z}}^2}{2\left(N+\frac{1}{\sigma^2}\right)}\right)}.
$$

The posterior probability of the slab then equals:

$$
\Pr\left(\text{slab} \mid \boldsymbol{Z}\right) = \frac{\Pr\left(\text{spike}\right)}{\Pr\left(\text{spike}\right) + \left(1 - \Pr\left(\text{spike}\right)\right)\text{BF}_{01}},
$$

and the posterior probability of the spike is the complement. It then follows that the cumulative distribution function for the spike-and-slab posterior is given by:

$$P(\delta \leq x \mid \boldsymbol{Z}) = \begin{cases} \Pr(\text{slab} \mid \boldsymbol{Z})\, \Phi(x; \mu_1, \sigma_1) & \text{if } x < 0, \\ \Pr(\text{spike} \mid \boldsymbol{Z}) + \Pr(\text{slab} \mid \boldsymbol{Z})\, \Phi(x; \mu_1, \sigma_1) & \text{if } x \geq 0, \end{cases}$$

where $\Phi(x; \mu_1, \sigma_1)$ is the cumulative normal distribution. Due to the discontinuity at $x = 0$ there is no useful closed form expression for the posterior density. Nevertheless, the posterior mean of the spike-and-slab model is available in closed form. Using the law of total probability, we have:

$$p(\delta \mid \boldsymbol{Z}) = \Pr(\text{spike} \mid \boldsymbol{Z})\, p(\delta \mid \text{spike}, \boldsymbol{Z}) + \Pr(\text{slab} \mid \boldsymbol{Z})\, p(\delta \mid \text{slab}, \boldsymbol{Z}).$$

Computing the mean of left hand side yields:

$$\begin{aligned} \int_{-\infty}^{\infty} \delta\, p(\delta \mid \boldsymbol{Z})\ \mathrm{d}\delta &= \Pr(\text{spike} \mid \boldsymbol{Z}) \int_{-\infty}^{\infty} \delta\, p(\delta \mid \text{spike}, \boldsymbol{Z})\ \mathrm{d}\delta, \\ &\quad + \Pr(\text{slab} \mid \boldsymbol{Z}) \int_{-\infty}^{\infty} \delta\, p(\delta \mid \text{slab}, \boldsymbol{Z})\ \mathrm{d}\delta, \\ &= 0 + \Pr(\text{slab} \mid \boldsymbol{Z})\, (\mu_\delta \mid \text{slab}, \boldsymbol{Z})\,. \end{aligned}$$

Here $(\mu_\delta \mid \text{slab}, \boldsymbol{Z})$ is the posterior mean of effect size under the slab. In a similar fashion, other statistics may be obtained. However, it is also possible to draw samples from marginal posterior distribution. To obtain a sample s, first draw u from a uniform distribution on $[0,\ 1]$. If $u < \Pr(\text{slab} \mid \boldsymbol{Z})$ draw s from $p(\delta \mid \text{slab}, \boldsymbol{Z})$, otherwise s is zero. This approach is often used when the integrals become too unwieldy to compute analytically. For example, the R package BAS uses this procedure to compute credible intervals (Clyde et al., 2011).

References

Clyde, M. A., Ghosh, J., & Littman, M. L. (2011). Bayesian adaptive sampling for variable selection and model averaging. *Journal of Computational and Graphical Statistics*, *20*, 80–101.

Rouder, J. N., Haaf, J. M., & Vandekerckhove, J. (2018). Bayesian inference for psychology, part IV: Parameter estimation and Bayes factors. *Psychonomic Bulletin & Review*, *25*, 102–113.