Comparing Bayesian Posterior Passing with Meta-analysis

Joshua Pritsker

San Diego Mesa College

Purdue University

Submitted to *Meta-Psychology*. Click here to follow the fully transparent editorial process of this submission. Participate in open peer review by commenting through hypothes.is directly on this preprint.

Abstract

Brand, von der Post, Ounsley, and Morgan (2019) introduced Bayesian posterior passing as an alternative to traditional meta-analyses. In this commentary I relate their procedure to traditional meta-analysis, showing that posterior passing is equivalent to fixed effects meta-analysis. To overcome the limitations of simple posterior passing, I introduce improved posterior passing methods to account for heterogeneity and publication bias. Additionally, practical limitations of posterior passing and the role that it can play in future research are discussed.

Keywords: bayesian updating, posterior passing, meta-analysis

Comparing Bayesian Posterior Passing with Meta-analysis

The ability to accumulate evidence across studies is often said to be a great advantage of Bayesian inference. This point was recently discussed by Brand, von der Post, Ounsley, and Morgan (2019), who suggested that Bayesian posterior passing alleviates the need for traditional meta-analysis. They performed a simulation study comparing posterior passing to non-cumulative analysis and combined analysis of the data from all studies. However, they lacked a formal theoretical comparison of posterior passing to traditional methods of meta-analysis. This commentary relates posterior passing to traditional meta-analyses, allowing for one to determine the performance of posterior passing on the basis of how traditional meta-analytic methods are known to perform. To avoid some of the pitfalls of posterior passing, I suggest improved procedures that account for heterogeneity and publication bias. I address Brand et al. (2019)'s suggestion that posterior passing avoids some of the problems of traditional meta-analyses, and discuss practical limitations in using it as a replacement for traditional meta-analysis.

Posterior Passing as Meta-analysis

As Brand et al. (2019) suggest that posterior passing may replace traditional meta-analyses, one might wonder how posterior passing relates to traditional meta-analyses. Given that meta-analyses are typically based on proxy statistics and their standard errors instead of individual data points, one might consider posterior distributions generated the same way. We can do this by using the likelihood of a statistic instead of the likelihood of the full data. For instance, supposing that we have a parameter θ that we want to make inferences about, we might construct the likelihood function using the sampling distribution of an estimate of θ . In standard meta-analysis, studies are typically summarized by their estimates and standard errors. If an

estimate was derived by maximizing a likelihood that takes the form of gaussian function over θ , the estimate and its standard error fully summarize this likelihood function. Even in cases where the likelihood function is not fully described by its maximum likelihood estimate and standard error, using the estimate and standard error may be viewed as a second-order asymptotic approximation to the log-likelihood.

Taking this approach, with $\hat{\theta}_i$ being the estimate from study i of θ , we can get the posterior after study i , denoted as $f_i(\theta)$, by:

$$f_i(\theta) \propto \pi(\theta) f(x_i | \theta) \quad (1a)$$

$$= f_{i-1}(\theta) f(\hat{\theta}_i | \theta, SE_i^2) \quad (1b)$$

Where $\pi(\cdot)$ is our prior distribution, x_i refers to all the information of study i , and SE_i is the standard error for $\hat{\theta}_i$. Typically, $f(\hat{\theta}_i | \theta, SE_i^2) = \phi(\hat{\theta}_i | \theta, SE_i^2)$, where $\phi(\hat{\theta}_i | \theta, SE_i^2)$ is the density at $\hat{\theta}_i$ under a normal distribution with a mean of θ and variance of SE_i^2 . Now, we may expand (1) across k studies:

$$f(\theta) \propto \pi_0(\theta) \prod_{i=1}^k \phi(\hat{\theta}_i | \theta, SE_i^2) \quad (2)$$

Where $\pi_0(\theta)$ is the prior distribution used by the initial study. Hence, the posterior is proportional to the initial prior times the product of the likelihoods of the relevant studies. How does this compare to the posterior distribution given by a traditional meta-analysis? In a fixed-effects setting (3), they are identical, provided that the meta-analysis uses the same prior as the initial study. The standard fixed effects model is that each $\hat{\theta}_i$ follows a normal distribution with a universal mean value of θ , and a variance of SE_i^2 :

$$\hat{\theta}_i \sim \mathcal{N}(\theta, SE_i^2) \quad (3)$$

Then, the posterior density can be constructed by taking the product of the likelihoods of each estimate:

$$f_{fixed}(\theta | x) \propto \pi(\theta) \prod_{i=1}^k \phi(\hat{\theta}_i | \theta, SE_i^2) \quad (4)$$

This equivalence also answers some of the questions brought up by Brand et al. (2019), such as how posterior passing would perform under publication bias. Now, it is clear that we can use existing studies on traditional fixed-effects meta-analyses to answer these questions (e.g., Simonsohn, Nelson, & Simmons, 2014).

In a random effects setting, we assume that each study samples from a slightly different population, and these populations have their own parameter values, referred to as μ_i , which vary around the true θ value:

$$\mu_i \sim \mathcal{N}(\theta, \tau^2) \quad (5a)$$

$$\hat{\theta}_i \sim \mathcal{N}(\mu_i, SE_{\hat{\theta}_i}^2) \quad (5b)$$

Then, the likelihood for each estimate is given by marginalizing out μ_i :

$$f_{random}(\theta | x) \propto \pi(\theta, \tau) \prod_{i=1}^k \int_{\mu} \phi(\hat{\theta}_i | \mu_i, SE_i^2) \phi(\mu_i | \theta, \tau^2) d\mu_i \quad (6a)$$

$$= \pi(\theta, \tau) \prod_{i=1}^k \phi(\hat{\theta}_i | \theta, SE_i^2 + \tau^2) \quad (6b)$$

Where μ_i is the mean of the population that study i comes from, τ^2 is the variance across populations, and the prior distribution is now over both θ and τ . Hence, in a random-effects setting (4), posterior passing, like (3), will underestimate the posterior variance. This point was previously noted by Martin (2017). A random-effects model is typically preferable, as the fixed-effects assumption of no between-study variance is implausible in most settings (Borenstein, Hedges, Higgins, & Rothstein, 2010). The use of the full-data posterior by Brand et al. (2019) is

inconsequential to this result. Hence, although posterior passing will produce consistent point estimates, the posterior variance may be underestimated.

How Can we Improve Posterior Passing?

Incorporating Random Effects

An obvious question to ask at this point is if we can modify posterior passing to incorporate random effects. A Bayesian solution is to update τ along with θ , modeling their joint posterior distribution. Using statistic likelihoods as in the previous section, we can derive the joint posterior update by marginalizing out μ_i , just as in a standard random-effects meta-analysis:

$$f_i(\theta, \tau) \propto \pi(\theta, \tau) f(x_i | \theta, \tau) \quad (7a)$$

$$= f_{i-1}(\theta, \tau) f(x_i | \theta, \tau) \quad (7b)$$

$$= f_{i-1}(\theta, \tau) \int f(x_i | \mu_i) \phi(\mu_i | \theta, \tau^2) d\mu_i \quad (7c)$$

$$= f_{i-1}(\theta, \tau) \phi(\hat{\theta}_i | \theta, SE_i^2 + \tau^2) \quad (7d)$$

To get the marginal posterior distribution of θ , one may integrate out τ , and vice versa. Notably, the parameter τ must be tracked across studies to gain evidence about its value. However, even if previous studies hadn't used a random effects model, they can still be added using their likelihood functions:

$$f_i(\theta, \tau) \propto \pi(\theta, \tau) \phi(\hat{\theta}_{i-1} | \theta, SE_{i-1}^2 + \tau^2) \phi(\hat{\theta}_i | \theta, SE_i^2 + \tau^2) \quad (8)$$

Where study $i - 1$ is the study that hadn't used a random effects model. In the case of the current study being the first on a topic, no evidence will be gained about the value of τ , but incorporating it using only the information from a subjective prior distribution will nonetheless give a more realistic view of the uncertainty of $f_i(\theta)$.

Addressing Publication Bias

A second major issue with posterior passing is that Brand et al. (2019) provide no way to address publication bias. Without adjustment, posterior passing will perform identical to a fixed-effects meta-analysis that completely ignores publication bias. Unadjusted meta-analyses are known to perform poorly when publication bias is substantive, yielding potentially misleading results (Simonsohn, Nelson, & Simmons, 2014). This problem can be viewed as one of biased sampling, hence the likelihood function is given by a weighted distribution as follows

(Pfeffermann, Krieger, & Rinott, 1998):

$$f(x_i | \text{published}_i) = \frac{E_i[p | x_i] f(x_i)}{E[p]} = \frac{E_i[p | x_i] f(x_i)}{\int_x E[p] f(x_i) dx_i} \quad (9)$$

Where p are the probabilities of publication. Now, we need to model $E[p | x_i]$. Considering that studies are typically given a dichotomous interpretation (McShane & Gal, 2017), a realistic option is a simple step function:

$$E[p | x_i] = \begin{cases} \alpha & s_i(\hat{\theta}_i) = 0 \\ \beta & s_i(\hat{\theta}_i) = 1 \end{cases} \quad (10)$$

Where $s_i(x_i)$ is a function that gives the standard interpretation of $\hat{\theta}_i$ in a pass/fail manner, such as its p-value being below 0.05. However, it is unclear if dichotomization exists to the same extent in Bayesian studies as it does in frequentist studies. In such a context one may replace (7) with a smoother model, such as a logistic one:

$$E[p | x_i] = \text{logistic} \left(\alpha + \beta s_i(\hat{\theta}_i) \right) \quad (11)$$

Where $s_i(\hat{\theta}_i)$ could represent a Bayes factor cutoff or similar. For a review of other models that have been suggested, see Sutton, Song, Gilbody, and Abrams (2000). In any case, the posterior update is now as follows:

$$f_i(\theta, \tau, \alpha, \beta) \propto \pi(\theta, \tau, \alpha, \beta) f(x_i | \theta, \tau, \alpha, \beta) \quad (12a)$$

$$\propto \pi(\theta, \tau, \alpha, \beta) \frac{E[p | x_i, \theta, \tau, \alpha, \beta]}{E[p | \theta, \tau, \alpha, \beta]} f(x_i | \theta, \tau) \quad (12b)$$

$$= f_{i-1}(\theta, \tau, \alpha, \beta) \frac{E[p | x_i, \theta, \tau, \alpha, \beta]}{E[p | \theta, \tau, \alpha, \beta]} \int f(x_i | \mu_i) \phi(\mu_i | \theta, \tau^2) d\mu_i \quad (12c)$$

$$= f_{i-1}(\theta, \tau, \alpha, \beta) \frac{E[p | x_i, \theta, \tau, \alpha, \beta]}{E[p | \theta, \tau, \alpha, \beta]} \int_{\mu} \phi(\hat{\theta}_i | \theta, SE_i^2 + \tau^2) \quad (12d)$$

As with τ , multiple studies are needed to identify α , and β . Hence, inference in early studies will be highly dependent on the prior distribution for these parameters, which should not be uninformative. However, this problem dissipates as evidence for these parameters accumulates.

Including Studies Outside of the Posterior Passing Chain

In meta-analysis, extensive literature searches are conducted to avoid systematically excluding any studies. However, it may not be obvious how one can incorporate studies outside of a posterior passing ‘chain,’ a sequence of studies where each study’s prior is equal to the previous study’s posterior, into our prior distribution. The same problem occurs if two studies are done simultaneously, creating a fork in the chain. Brand et al. (2019) suggest that a normal prior with variance representing our uncertainty and with a mean at the estimate given by the study may be used. However, it can be difficult to determine such a variance, and this procedure only makes sense if we are the first study in a chain aiming to get information from an unchained study. A better answer is to include the study’s likelihood function in our posterior derivation. When there are simultaneous studies, simply take one’s posterior as the prior and use the likelihood from the other as if it were outside of the chain altogether. Switching back to the simple fixed effects procedure for conciseness, this gives us:

$$f_i(\theta) \propto \pi(\theta) f(x_{i-1} | \theta) f(x_i | \theta) \quad (13)$$

To include more studies, one simply needs to add more $f(x_n | \theta)$ functions. As a side note, the nature of this function yields an obvious option for Bayesian-style updating in a non-Bayesian framework:

$$L_i(\theta) \propto f(x_1, \dots, x_{i-1} | \theta) f(x_i | \theta) \quad (14)$$

This may simply be interpreted as the likelihood function for θ across all the included studies.

Inferences can then be made using standard frequentist sequential trial methods (cf. Wetterslev, Jakobsen, & Gluud, 2017).

Further Comments and Discussion

Is Posterior Passing a Practical Replacement for Meta-analysis?

Brand et al. (2019) suggest that posterior passing can replace traditional meta-analyses. Indeed, the improved posterior passing procedures introduced in the previous section can compete with traditional meta-analyses, but is it practical? To match the quality of traditional meta-analyses, one would have to meet the same conditions, including extensive literature search and the inclusion of all available studies. Furthermore, most studies are currently done in a frequentist manner, and splits can occur in a chain due to simultaneous studies, so in practice each study in the chain will have to do its own mini meta-analysis. At that point, it might be preferable to just do traditional meta-analyses.

Brand et al. (2019) also suggest that posterior passing can solve the problem of conflicting meta-analyses by updating evidence in real time. Meta-analyses can provide contradictory results due to a number of reasons, such as differing inclusion criteria and using different methods. However, posterior passing only appears to be able to mitigate differences occurring by meta-analyses occurring at different times rather than by different methods. In this

case though, one would simply go with the most recent meta-analysis. However, when multiple meta-analyses disagree, they often have differences in statistical methodology or inclusion criteria. Replacing meta-analyses with posterior passing would likely result in multiple posterior passing chains to reflect this disagreement. Perhaps this might be avoided if fields are sufficiently vigilant in preventing conflicting chains, but the same would be true of traditional meta-analytic conflicts. In fact, one might argue that for *any* criticism of a meta-analysis, one could seemingly make an equivalent criticism of a posterior passing chain. Instead of having disagreeing meta-analyses, we would instead simply have a number of individual studies in disagreement. It may be that posterior passing could help mitigate disagreements of this nature by the fact that inclusion criteria could change with any study in the chain. However, having such fuzzy inclusion criteria is clearly undesirable as it would lead to results with an unclear interpretation. Hence, posterior passing does not appear to avoid the problem of conflicting meta-analyses in a desirable manner. The same applies for all methodological limitations of standard meta-analyses, such as being impacted by publication bias, as posterior passing and meta-analysis are mathematically equivalent.

Alternative Roles for Posterior Passing

Even if posterior passing cannot generally replace traditional meta-analyses, it may nonetheless be useful. With the improvements suggested above, posterior passing can replace traditional meta-analysis in areas where meta-analyses are unlikely to produce conflicting results in the first place. An alternative to creating posterior passing chains that still utilizes the posterior passing mechanism is to use it in meta-analyses. This doesn't solve the issue of conflicting meta-analyses, but has practical advantages. By using the posterior distribution of

the last similar meta-analysis as a prior distribution, meta-analyses can be performed in chunks instead of having to redo the entire meta-analysis with each update. Similarly, instead of using posterior passing chains, studies can use posterior distributions from meta-analyses as their priors to get accurate net effect estimates within each study. This allows for broader conclusions than would otherwise be warranted by the study alone. A particularly relevant case for this is large-scale replication projects, where the prior can be gotten from the last meta-analysis. This yields a readily interpretable lower-bound on the extent to which a field's view on a topic should shift as the result of the replication effort, by providing the change in posterior assuming that previous studies had been conducted properly. Hence, although posterior passing may have problems as a replacement for meta-analysis, it can have utility regardless.

References

- Brand, C. O., Ounsley, J. P., van der Post, D. J., & Morgan, T. J. H. (2019). Cumulative Science via Bayesian Posterior Passing. *Meta-Psychology*, 3. doi:10.15626/MP.2017.840
- Borenstein, M., Hedges, L. V., Higgins, J. P., & Rothstein, H. R. (2010). A basic introduction to fixed-effect and random-effects models for meta-analysis. *Research synthesis methods*, 1(2), 97-111. doi:10.1002/jrsm.12
- Martin, S. (2017). Open Peer Review by Stephen Martin. *Decision Letter for Brand et al.*
- McShane, B. B., & Gal, D. (2017). Statistical significance and the dichotomization of evidence. *Journal of the American Statistical Association*, 112(519), 885-895. doi:10.1080/01621459.2017.1289846
- Pfeffermann, D., Krieger, A. M., & Rinott, Y. (1998). Parametric distributions of complex survey data under informative probability sampling. *Statistica Sinica*, 1087-1114.
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). p-Curve and effect size: Correcting for publication bias using only significant results. *Perspectives on Psychological Science*, 9(6), 666-681.
- Sutton, A. J., Song, F., Gilbody, S. M., & Abrams, K. R. (2000). Modelling publication bias in meta-analysis: a review. *Statistical Methods in Medical Research*, 9(5), 421-445. doi:10.1177/096228020000900503
- Wetterslev, J., Jakobsen, J. C., & Gluud, C. (2017). Trial Sequential Analysis in systematic reviews with meta-analysis. *BMC medical research methodology*, 17(1), 39. doi:10.1186/s12874-017-0315-7