Putting Variation into Variance: Modeling Between-Study Heterogeneity in Meta-Analysis

Donald R. Williams
University of California, Davis

Josue E. Rodriguez
University of California, Davis

Paul-Christian Bürkner
University of Stuttgart, Germany

We shed much needed light upon a critical assumption that is oft-overlooked—or not considered at all—in random-effects meta-analysis. Namely, that between-study variance is constant across *all* studies which implies they are from the *same* population. Yet it is not hard to imagine a situation where there are several and not merely one population of studies, perhaps differing in their between-study variance (i.e., heteroskedasticity). The objective is to then make inference, given that there are variations in heterogeneity. There is an immediate problem, however, in that modeling heterogeneous variance components is not straightforward to do in a general way. To this end, we propose novel methodology, termed Bayesian location-scale meta-analysis, that can accommodate moderators for both the overall effect (location) and the between-study variance (scale). After introducing the model, we then extend heterogeneity statistics, prediction intervals, and hierarchical shrinkage, all of which customarily assume constant heterogeneity, to include variations therein. With these new tools in hand, we go to work demonstrating that quite literally *everything* changes when between-study variance is not constant across studies. The changes were not small and easily passed the interocular trauma test—the importance hits right between the eyes. Such examples include (but are not limited to) inference on the overall effect, a compromised predictive distribution, and improper shrinkage of the study-specific effects. Further, we provide an illustrative example where heterogeneity was not considered a mere nuisance to show that modeling variance for its own sake can provide unique inferences, in this case into discrimination across nine countries. The discussion includes several ideas for future research. We have implemented the proposed methodology in the R package **blsmeta**.

Keywords: meta-analysis, between-study heterogeneity, heterogeneous variance components, Bayesian

Meta-analysis plays an important role in psychological science. When summarizing the results from several studies, researchers can determine the average effect size, the extent to which the effects vary from study to study, and examine whether study characteristics moderate the effect size. As a result, meta-analysis has proven to be an indispensable tool, casting doubt on long-standing theories (e.g., ego-depletion, Hagger et al., 2016), offering a refined view of psychological interventions (van Agteren et al., 2021), and resolving possible explanations of replication failures (e.g., "hidden moderators", Klein et al., 2018)

In this work, we tackle the important topic of between-study heterogeneity or variance (τ^2). This variability in the

effects could be due to experimental design, age of the study subjects, or environmental context (e.g., varying cultures or country of origin). Because of these potential differences, the analytical decision is often posited as a choice between a fixed-effects or random-effects model (Aguinis, Gottfredson, & Wright, 2011). This is somewhat misguided, however, because both "provide perfectly valid inferences under heterogeneity" (see p. 4 in Viechtbauer, 2010). The key distinction is that a fixed-effects model provides *conditional* inference for the observed studies. On the other hand, the random-effects model allows for *unconditional* inference on the average effect in a population of studies, of which the observed are assumed to be a random sample (Hedges & Vevea, 1998). Additionally, the random-effects model provides an estimate of τ^2, which opens the door to explaining heterogeneity (S. G. Thompson & Sharp, 1999). For these reasons, we focus on the random-effects model in this work.

We have the lofty ambition of *completely* revising the typical view of heterogeneity in meta-analysis. Often the goal

is unconditional inference, that is, to generalize beyond the observed studies and to a "normally distributed superpopulation" of hypothetical studies (Bonett, 2010). The mean μ of this superpopulation is defined by the average of the *true* effects, with τ^2 defining their dispersion around μ (the overall/average effect). This assumption is made in almost all meta-analyses, without considering the implications for the observed effects. For example, as noted in Higgins, Thompson, and Spiegelhalter (2009):

> This variance [τ^2] explicitly describes the extent of the heterogeneity and has a crucial role in assessing the degree of consistency of effects across studies, which is an element of random-effects meta-analysis that often receives too little attention (p. 139).

While we full heartedly agree with the importance of heterogeneity, a critical assumption is that τ^2 is constant across studies, which implies that they are *all* sampled from the same population. However, it is not hard to imagine a situation in which several and not merely one population of studies is more than plausible. Here there would be between-study variances and not a between-study variance.

Our idea is to fully extend the notion of *unconditional* inference to heterogeneous superpopulations, that potentially differ in their between-study variances. This entails changing τ^2 to τ_i^2, wherein study i can be sampled from a normal distribution with relatively more (or less) heterogeneity (see Figure 1). In doing so, this (i) accounts for the possibility that the effects are from populations with different variances (i.e., heteroskedasticity), and (ii) opens the door to predicting τ_i^2 with study-level covariates.

Are Variations in Heterogeneity Meaningful?

At first, over and above accommodating unequal (between-study) variances, the importance of τ_i^2 may not be readily apparent. The typical view is that heterogeneity is to be explained (i.e., accounted for by moderators, S. G. Thompson & Sharp, 1999) or perhaps that variance is a mere nuisance to understanding the overall effect. The latter is reflected in the large literature on heterogeneity of variance and error rates (e.g., Blanca, Alarcón, Arnau, Bono, & Bendayan, 2018; Delacre, Lakens, & Leys, 2017; Ruscio & Roche, 2012). While certainly important, this perspective strikes us as rather coarse, perhaps the result of running from rather than embracing variation (see for example Gelman, 2019). The following three examples provide motivation for modeling τ_i^2.

First, variability can be substantively interesting—modeling variance for its own sake. In mixed-effects models, often the "error" or within-person variance is not considered error at all, but rather an indicator of intraindividual variability or inconsistency (Ram, Stollery, Rabbitt, & Nesselroade,
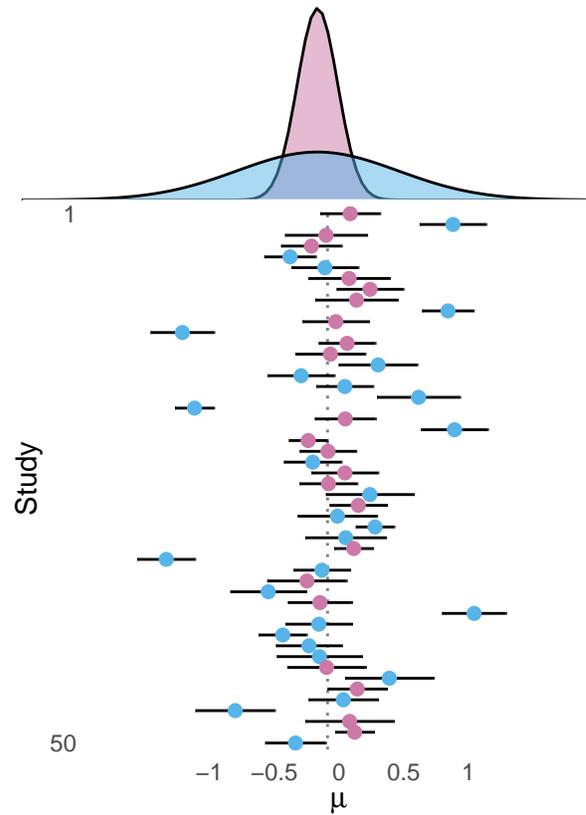

Figure 1. Illustrative example, where the studies are sampled from two superpopulations that differ in their variance (heteroskedasticity). Notice that the studies from the blue population are more "dispersed" around the overall effect. This implies that there are variations in heterogeneity.

2005). For example, in Kalin et al. (2014), inconsistency of response times in a cognitive task was suggested to be a marker of Alzheimer's disease. Translated to τ^2, this would be akin to investigating heterogeneity in relation to whether the studies included subjects with a major versus mild neurocognitive disorder. There is also a long tradition of studying variability in gerontology (e.g., MacDonald, Hultsch, & Dixon, 2003; Rast, 2011). In this case, modeling τ^2 could shed light upon age related heterogeneity in treatment effects.

Second, τ^2 is an important component of measures used to quantify heterogeneity in meta-analysis (Higgins & Thompson, 2002). For example, I^2 is defined as "the proportion of the variance in study estimates that is due to heterogeneity" (p. 2, Von Hippel, 2015), and it is reported in most meta-analyses. These measures are computed under the assumption that the studies are from one superpopulation. When τ_i^2 is determined to be a function of covariates, however, this implies that heterogeneity measures are a function of those same covariates. The importance of this cannot be under-

stated, relating again to the idea of heterogeneous populations. Given a study-level characteristic (e.g., age groups), that is, there could be varying amounts of heterogeneity that can then be explained or interpreted for its own sake. Further, it has been argued that meta-analyses should include prediction intervals, with the goal of assisting in "the clinical interpretation of the heterogeneity by estimating what true treatment effects can be expected in future settings" (p. 1, IntHout, Ioannidis, Rovers, & Goeman, 2016). Here, too, prediction intervals are computed under the assumption that τ^2 is constant across studies. In essence, by relaxing this assumption prediction intervals become a function of covariates as well.

Third, even when the focus is the overall effect and heterogeneity is an after thought, the magnitude and uncertainty of the overall effect is computed under the assumption of a common τ^2 (see Equations 15 and 16 in Hedges, 1992). In other words, that the studies were indeed sampled from a distribution with a constant between-study variance. As demonstrated below, when this is not the case the total variance in the effect size estimates is incorrect, thereby compromising statistical inference on the overall effect (i.e., not getting τ "right" affects μ). The objective is to then estimate the overall effect and its uncertainty, given that there are variations in τ_i^2.

Major Contributions

Importantly, the idea of varying τ_i^2 is not itself novel, as, for example, it was described in Bowater and Escarela (2013) and C. G. Thompson and Becker (2020). It is the full thrust of the implications therein that have the chance to alter the trajectory of meta-analysis, which, to date, have not been unpacked. Hence, by way of example, our first contribution is demonstrating the importance of considering τ_i^2 in meta-analysis—quite literally *everything* changes. There is an immediate challenge, however, in that modeling variance components is not straightforward to do in a general way. To this end, our second contribution is proposing a novel Bayesian model for location-scale meta-analysis. The location refers to mean (or effect size) whereas the scale is the between-study variance. A key advantage of our methodology is that it readily captures uncertainty in *all* measures of heterogeneity that are typically computed from τ^2 (e.g., I^2). Our third contribution is to compute them from τ_i^2 (resulting in I_i^2). To our knowledge, τ_i^2 has not been incorporated into heterogeneity statistics or prediction intervals. In addition, as our fourth contribution, we implemented the methods in the new R package **blsmeta**.

Overview

In what follows, we first describe the customary random-effects model. Here we focus on the implications of assuming τ^2 is across studies, when, in fact, there is variation in

between-study heterogeneity. In the next section, we introduce our Bayesian formulation for location-scale meta-analysis. This section includes demonstrating how our model naturally incorporates τ_i^2 into customary heterogeneity statistics and prediction intervals that are typically computed from τ^2. We then demonstrate the utility of our method in two illustrative examples, including unique insights that can be gleaned from directly modeling τ_i^2.

Motivating Example

We begin by revisiting the random-effects model estimated in what is considered a "frequentist" framework. The reason for this is rhetorical in nature, in that it readily allows for teasing apart the implications of τ_i^2 with closed form solutions. The basic idea also applies to Bayesian estimation, although the expressions are less familiar in the meta-analysis literature. To see this, we refer interested readers to Equation 17 in Röver (2020), as compared to those in Higgins and Thompson (2002, see section 3.1.1 therein).

The typical goal is to make inference on μ, given $i = 1, \ldots, k$ independent effect size estimates, y_i, and their sampling variances, σ_i^2, that are assumed to be known in advance (i.e., computed from the sample data within each study). In essence, the random-effects model is analogous to the customary normal-normal hierarchical model with known variances (Röver, 2020). Accordingly, at level-one, we have

$$y_i \sim \mathcal{N}(\theta_i, \sigma_i^2), \tag{1}$$

where θ_i is the corresponding true effect for the ith study. We use "true" because it is common in the literature (Borenstein, Hedges, Higgins, & Rothstein, 2010; Higgins et al., 2009), but reserve the possibility that true effects do not actually exist and are hypothetical themselves. Alternatively, θ_i can be considered the "study-specific" effects. From Equation (1), it follows that the level-one or within-study variance component is σ_i^2. To investigate additional variability, over and above the sampling variances, the true effects are assumed to be drawn from a normal distribution, that is,

$$\theta_i \sim \mathcal{N}(\mu, \tau^2), \tag{2}$$

where μ and τ^2 denote the average true effect and between-study variance, respectively. τ^2 is the level-two variance component that captures between-study heterogeneity in the true effects. To make sense of the hierarchical structure, it helps to consider that each effect can deviate from μ. It is the magnitude of those deviations, or their *dispersion* around the overall effect, that defines τ^2 (e.g., Figure 1). Accordingly, in the event that $\tau^2 = 0$, the study-specific effects are equal to the overall effect ($\theta_1 = \ldots = \theta_k = \mu$).

Notice in Equation (2) that τ^2 is assumed to be the *same* for each study. This translates into an assumption regarding the total variance in the observed effects. Under constant τ^2,

we have that $var(y_i) = \sigma_i^2 + \tau^2$ (see also p. 80 in Raudenbush & Bryk, 1985), which is readily seen when expressing the random-effects model in its marginal form. This is given by

$$y_i \sim \mathcal{N}(\mu, \sigma_i^2 + \tau^2). \qquad (3)$$

Equation (3) is often used when θ_i is not of direct interest. In this case, we use it to emphasize the underlying assumption of $var(y_i)$ that is made in most meta-analyses.

Inference on the Overall Effect

Although our primary focus is heterogeneity, it is important to note that inference on μ is inextricably linked to τ^2 in a random-effects meta-analysis. In the literature, this can be seen by the sheer number of estimators for τ^2 (see Table 3 in Veroniki et al., 2016). In part, this can be attributed to $\hat{\tau}^2$ affecting "the estimate of the mean effect and its confidence interval" (p. 182, Langan, Higgins, & Simmonds, 2017). This is because heterogeneity has a direct bearing upon $var(y_i)$, as seen in Equation (3). To our knowledge, however, there has been little work considering the implications of variation in τ_i^2 (but see Rubio-Aparicio et al., 2020). In what follows, we open Pandora's jar on this important topic.

With the observed effects, y_i, and the corresponding sampling variances, σ_i^2, the overall effect can be computed analytically. This is given by

$$\hat{\mu} = \frac{\sum\limits_{i=1}^{k} w_i^* \cdot y_i}{\sum\limits_{i=1}^{k} w_i^*}, \qquad (4)$$

which is effectively a weighted average. Further, an approximate variance of $\hat{\mu}$ under the random-effects model is $\sigma_\mu^2 = 1/(\sum_{i=1}^{k} w_i^*)$. It is the weights, w_i^*, that are of clear importance for making inference on μ and they are typically computed under the assumption that τ^2 is constant for each study, that is,

$$w_i^* = \frac{1}{\sigma_i^2 + \hat{\tau}^2}, \qquad (5)$$

where $\hat{\tau}^2$ can be obtained from any of the common estimators (see Table 3 in Veroniki et al., 2016). The basic idea of Equation (5) is that the weights correspond to the inverse of the *total* variance for each study (i.e., the precision), corresponding to "the within-study variance for study (i) plus the between-studies variance, tau-squared" (p. 16, Borenstein, Hedges, & Rothstein, 2007). Hence, when taking the reciprocal, the more precise effect size estimates necessarily receive a larger w^*, given that τ^2 is constant across the k studies. This, in turn,

means they have relatively more influence on the overall effect and its variance—the necessary ingredients for statistical inference.

However, if there are variations in heterogeneity, then Equation (5) would be incorrect. This is because it would have the incorrect variance in the denominator. To accommodate study-specific heterogeneity, τ_i^2, the marginal formulation given in Equation (3) can be redefined as

$$y_i \sim \mathcal{N}(\mu, \sigma_i^2 + \tau_i^2), \qquad (6)$$

where it is readily apparent that $var(y_i) = \sigma_i^2 + \tau_i^2$. As a result, the modified random-effects weights are

$$w_i^{**} = \frac{1}{\sigma_i^2 + \hat{\tau}_i^2}, \qquad (7)$$

which now reflects the total variance in y_i. These would then take the place of w_i^* in Equation (5) and Equation (4). Hence, the overall effect and its variance will change, now reflecting variations in heterogeneity. Of course, if there is no variation to speak of, the between-study variances are equal ($\tau_i^2 = \ldots = \tau_k^2 = \tau^2$). The weights in Equation (5) can therefore be viewed as a special case of those in Equation (7).

Random-Effects Weights. To highlight the importance of τ_i^2, we first examined w_i^* and w_i^{**} under the assumption that σ_i^2 and τ_i^2 ($i = 1, \ldots, 20$) were known in advance. Each sampling variance was set to $\sigma_i^2 \approx 0.021$, in essence a Fisher z-transformed correlation with $n = 50$. Heterogeneity for the first superpopulation was set to $\tau_{g1}^2 = 0.05^2$, whereas, for the second superpopulation, τ_{g2}^2 took on a range of values that spanned from 0.05^2 to 1. In this set up, each of the k studies belonged to either population, with the proportion "drawn" from the first ranging from $0.50, 0.60 \ldots, 0.90$. It follows that $w_i^{**} \approx 42$ is the true weight for studies from the first population and when assuming a constant variance the (incorrect) weights are $w_i^* = 1/(\sigma_i^2 + \frac{1}{k}\sum_{i=1}^{k} \tau_i^2)$.

Figure 2 (panel A) includes the computed weights. Notice the horizontal line at the true value, whereas the other lines were computed under the assumption that heterogeneity was constant across the studies (equal to the average). As heterogeneity increased in the second superpopulation (denoted $g2$), the weights approached zero. Said another way, given that the weights correspond to the inverse of the variance, the variance of the studies from $g1$ was overestimated as heterogeneity increased in $g2$. This was a function of the proportion of studies from each population. When there were fewer studies from $g1$, the (incorrect) weight more quickly approached zero. These findings can be understood in reference to the total variance of y_i (Equation 6), in that, again, it was overestimated for studies from the first population.
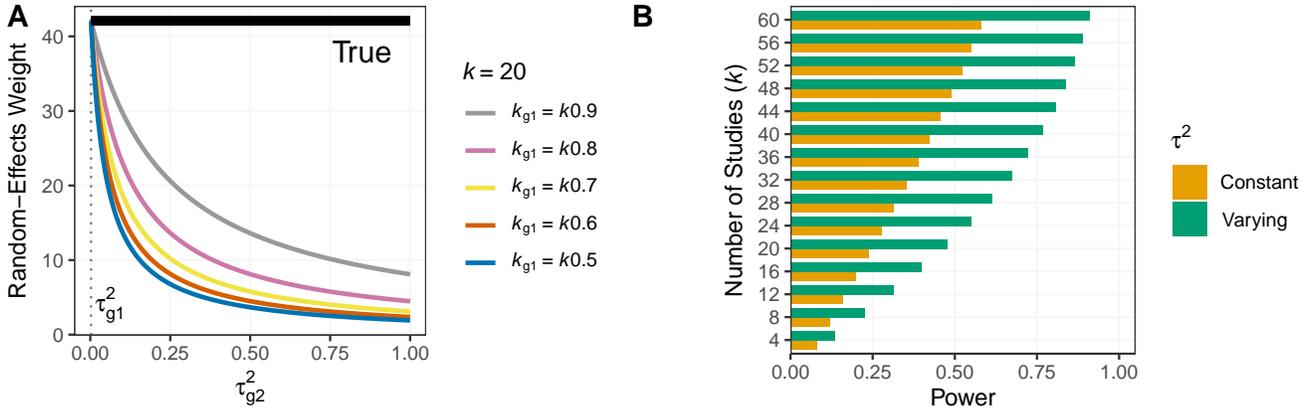
Figure 2. Panel A includes random-effects weights (inverse of the variance). Between-study variance (τ^2) for group two is on the x-axis, whereas the dotted line denotes τ^2 for group one. The *true* weight was computed with Equation (7), that accounts for variations in τ_i^2. The other (incorrect) weights were instead computed assuming that τ^2 was constant across studies (Equation 5). This indicates that the weights used to compute the overall effect (Equation 4) and its uncertainty are compromised when there are variations in heterogeneity, but assumed constant. Panel B includes statistical power to detect a small correlation ($r = 0.1$). There was an equal number of studies in each group ($\tau_{g1}^2 = 0.10^2$ and $\tau_{g2}^2 = 0.45^2$). The difference in power is a result of computing the weights when assuming constant versus varying heterogeneity.

Statistical Power. A natural question is then the extent to which getting the total variance wrong influences statistical inference. To this end, we investigated statistical power for detecting an effect size of 0.10 for $k = 4, 8, \ldots, 60$. The sampling variances were again set to $\sigma_i^2 = 0.021$ in two superpopulations ($\tau_{g1}^2 = 0.10^2$ and $\tau_{g2}^2 = 0.45^2$). These values, while certainly different from each other, were obtained from real data used in the illustrative examples. In this case, an equal number of studies belonged to each population. Because all of the variances are known, we computed power for a two-sided Z-test (see p. 212, Hedges & Pigott, 2001). The key distinction is whether each study had the correct variance, with the corresponding model denoted $\mathcal{M}_v$, or that computed under the assumption of a constant τ^2 (i.e., the average of τ_i^2, with the model denoted $\mathcal{M}_c$).

These results are presented in Figure 2 (panel B). There is a striking difference in statistical power. For example, power with $k = 20$ and the correct variance was nearly 0.50 ($\mathcal{M}_v$), but it was only 0.25 when assuming a constant between-study variance ($\mathcal{M}_c$). In fact, it was not until $k = 44$ that power reached 0.50 (over twice as many studies). Although not included in the results, the number of studies required to reach power of 0.80 was 100 for $\mathcal{M}_c$—this required only $k = 44$ for $\mathcal{M}_v$.

Summary

Together, although we made a slight modification to notation (i.e., merely adding i to τ^2) and presented an intentionally simplistic example, the importance of changing variance was readily apparent. As we demonstrated, when the weights are incorrect, this can result in overestimating the variance of $\hat{\mu}$. This translated into a drastic reduction in statistical power (Figure 2, panel B). We revisit this topic in the discussion, including several ideas to further explore the implications of τ_i^2 (e.g., type I error rate). In the next sections, we introduce the Bayesian location-scale meta-analysis and then illustrate the implications of modeling τ_i^2.

Bayesian Location-Scale Meta-Analysis

Our methodology is based upon a straightforward extension of the traditional mixed-effects approach in meta-analysis. To understand the innovation, it is important to first note that there are two kinds of variance in meta-analysis: (i) that which cannot be explained or the strictly unexplainable, corresponding to the sampling variances, and (ii) that which can be explained (potentially), corresponding to between-study heterogeneity. Our approach further partitions the latter by incorporating a sub-model to account for between-study variances (i.e., τ_i^2). Models for meta-analysis that can accommodate moderators (study-level variables) for both the location, or the overall effect, and the scale, or the between-study variance, are analogous to the mixed-effects *location-scale* model (Hedeker, Mermelstein, & Demirtas, 2008, 2012), which builds upon earlier work on variance heterogeneity (e.g., Aitkin, 1987; Cleveland, Denby, & Liu, 2003).[1]

---

[1]Location-scale modeling is also known as, for example, distributional regression (Bürkner, 2017) and doubly hierarchical models (Lee & Nelder, 2006)

Model Formulation

We now part ways from the traditional meta-analysis literature, and align our notation with that common to Bayesian mixed-effects (a.k.a., multilevel or hierarchical) modeling. The starting point is defining the likelihood for $i = 1, \ldots, k$ independent effect sizes, that is,

$$y_i \sim \mathcal{N}(\mu_i, \sigma_i^2) \tag{8}$$

which, in this case, is assumed to be Gaussian, with y_i and σ_i^2 denoting the effect size estimates and the known sampling variances, respectively.

Location Sub-Model. The mean of this distribution can be predicted, in essence a meta-regression. In full generality, this is given by

$$\mu_i = \mathbf{X}_i^l \boldsymbol{\beta} + u_i, \tag{9}$$

where $\mathbf{X}^l$ is the $k \times p$ design matrix including the study-level variables for the location sub-model (notice the l) and $\boldsymbol{\beta}$ is a $p \times 1$ vector including the regression coefficients. Note that first column in $\mathbf{X}^l$ typically includes 1's, corresponding to the intercept, with the moderators in the remaining columns. When the overall effect across all k studies is of interest (no moderator variables are considered), this corresponds to having $p = 1$ and thus fitting an "intercept only" model.

Scale Sub-Model. What remains from Equation (9) is u_i, the varying intercept for study i. Note that another term for a varying effect is a random effect. We prefer varying because it captures the fact that it varies from study to study, and avoids long-standing confusion surrounding what are and are not "random effects" (see section 6 in Gelman et al., 2005). In our formulation, this is given by

$$u_i \sim \mathcal{N}(0, \tau_i^2), \tag{10}$$

where u_i is the varying-effect and τ_i^2 is the between-study variance, both for the ith study. Recall that τ^2 is typically assumed constant for each study in the random-effects model (Equation 2), but here it has the subscript i.

At last, the sub-model for τ_i^2 is given by

$$\tau_i^2 = \exp(\mathbf{X}_i^s \boldsymbol{\gamma}) \tag{11}$$

where $\mathbf{X}^s$ is the $k \times p$ design matrix including the study-level variables for the scale sub-model (notice the s) and $\boldsymbol{\gamma}$ the $p \times 1$ vector including the regression weights. When the average heterogeneity across all k studies is of interest (no moderator variables are considered for the scale), this

corresponds to having $p = 1$ and thus fitting an "intercept only," resulting in the customary τ^2 in random-effects meta-analysis (e.g., Equation 2). In our formulation, heterogeneity is modeled on the log scale, that is, $\log(\tau_i^2) = \mathbf{X}_i^s \boldsymbol{\gamma}$, which is equivalent to Equation (11). This is motivated by needing an (inverse) link function to ensure τ_i is positive and that the "skewed, non-negative nature of the log-normal distribution makes it a reasonable choice for representing variances" (p. 629 Hedeker et al., 2008).

Prior Specification. What remains is the prior distributions (or simply "prior"). While certainly not the only option, we assign each coefficient a weakly informative prior (see discussion in Gelman, Simpson, & Betancourt, 2017). In the location sub-model, the assumed prior for the intercept is $\boldsymbol{\beta}_1 \sim \mathcal{N}(\bar{y}, 1)$, which is centered at the empirical mean of the effect size estimates $\frac{1}{k} \sum_{i=1}^k y_i$. This is reasonable choice because we center all predictors before fitting the model. Note that centering is often done in Bayesian analysis to ease sampling from the posterior distribution. When moderators are included in the model, the corresponding prior is $\boldsymbol{\beta}_{2:p} \sim \mathcal{N}(0, 1)$.

Importantly, the priors for the scale sub-model must be considered on the log scale. For the intercept, the assumed prior is $\boldsymbol{\gamma}_1 \sim \mathcal{N}(-2, 1)$, which mirrors the informed prior of Gronau et al. (2017, Figure 3 therein), that was based on 162 estimates of τ from the psychological literature (van Erp, Verhagen, Grasman, & Wagenmakers, 2013). Admittedly, thinking about priors on the log scale is far from intuitive. We recommend to rely on visualization when wanting to specify a different prior. In particular, samples can be drawn from the prior predictive distribution, which "makes use of simulations from the model rather than observed data" (p. 10, Gelman et al., 2020). This can then be used to refine the priors (see a proposed Bayesian workflow in Figure 1 of Gelman et al., 2020). When moderators are included in the scale sub-model, the corresponding prior is $\boldsymbol{\gamma}_{2:p} \sim \mathcal{N}(0, 1)$, which can be considered weakly informative. We return to the topic of prior specification in the discussion, especially when wanting to go beyond estimation (as in this work) and to Bayesian hypothesis testing with the Bayes factor.

Full Model

Having specified all the necessary ingredients, we can define the full model succinctly as

$$\begin{aligned} y_i &\sim \mathcal{N}(\mu_i, \sigma_i^2) \\ \mu_i &= \mathbf{X}_i^l \boldsymbol{\beta} + u_i \\ u_i &\sim \mathcal{N}(0, \tau_i^2) \\ \tau_i^2 &= \exp(\mathbf{X}_i^s \boldsymbol{\gamma}) \end{aligned}$$

This formulation readily accommodates the total variance in y_i, in the event that there are variations in τ_i^2. Additionally,

given that heterogeneity is now a function of covariates, *all* measures based on it are a function of those same covariates. We scratch the surface of considering τ_i^2 with the three examples, all of which are novel contributions to the meta-analysis literature.

Hierarchical Shrinkage

A key aspect of random-effects meta-analysis is that the study-specific effects, θ_i, are estimated by "borrowing strength" from the other studies. This can provide *shrinkage* towards the overall effect, μ, thereby reducing variability in the estimates. Although this is sometimes related to "Stein's Paradox" from the statistics literature (Efron & Morris, 1977; Stein, 1956), the basic idea, with derivations therein, was described in psychometrics nearly a century ago (p. 178, see eq. 22, Kelley, 1927).

An often overlooked—or not considered at all– assumption of hierarchical shrinkage in meta-analysis is that there is no variation in τ_i^2, that is,

$$\lambda_i = \frac{\tau^2}{\sigma_i^2 + \tau^2}, \tag{12}$$

where λ_i is sometimes called the shrinkage factor that determines the amount of shrinkage. Note that λ_i is the average score intraclass correlation coefficient or reliability (see p. 995 in Shieh, 2016).[2] In psychometrics, the numerator and denominator are thought to reflect "true score" and "observed score" variance (p. 86 in Raudenbush & Bryk, 1985), respectively.

In our formulation, due to relaxing the assumption of constant heterogeneity, the amount of shrinkage can be understood as

$$\lambda_i = \frac{\tau_i^2}{\sigma_i^2 + \tau_i^2} \tag{13}$$
$$= \frac{\exp(\mathbf{X}_i^s \boldsymbol{\gamma})}{\sigma_i^2 + \exp(\mathbf{X}_i^s \boldsymbol{\gamma})}.$$

When assuming that heterogeneity is constant (Equation 12), this corresponds to fitting an "intercept only" model to τ_i^2 in our formulation (Equation 13). As a result, if there are variations in τ_i^2, the customary formulation can result in "shrinkage of group-specific parameters [θ_i] by inappropriate amounts" (p. 147, Hoff, 2009).[3] This is due to assuming that the effects are from the same population, whereas, in our formulation, shrinkage can adapt to τ_i^2.

With λ_i in hand from Equation (13), the estimates of θ_i are given by

$$\theta_i = \lambda_i \cdot y_i + (1 - \lambda_i) \cdot \mu, \tag{14}$$

which is a weighted average for the ith effect size estimate and the overall effect (i.e., the estimated mean of the super-population). An approximate standard error of θ_i can be computed as $\sigma_i \cdot \sqrt{\lambda_i}$ (p. 146, Higgins et al., 2009). Accordingly, with $\lambda_i = 1$, this results in $\theta_i = y_i$ (no shrinkage). On the other hand, with $\lambda_i = 0$, this results in $\theta_i = \mu$ (complete shrinkage to the overall effect). Again, we emphasize that this derived under τ_i^2 and not τ^2, with the latter understood as a special case when $\tau_1^2 = \ldots = \tau_k^2 = \tau^2$.

Predicting Shrinkage. The estimates of θ_i naturally arise from our Bayesian formulation (obtained during model fitting). In a "frequentist" framework, they would be obtained after fitting the model. In addition to highlighting how τ_i^2 can influence the estimate of θ_i, the scale sub-model allows for predicting λ_i (the amount of shrinkage), say, at values of the moderators included in $\mathbf{X}^s$. This is given by

$$\hat{\lambda}_i = \frac{\hat{\tau}_i^2}{\sigma_i^2 + \hat{\tau}_i^2} \tag{15}$$
$$= \frac{\exp(\mathbf{X}_{new,i}^s \hat{\boldsymbol{\gamma}})}{\sigma_i^2 + \exp(\mathbf{X}_{new,i}^s \hat{\boldsymbol{\gamma}})},$$

where *new* denotes values of the moderators for which τ_i^2 is to be predicted. In words, Equation (15) effectively allows for predicting the shrinkage factor. Equivalently, when thinking of psychometrics and measurement, we are in essence predicting the reliability of effect size estimates. Hence, given values of y_i, σ_i^2, and μ, the predicted $\hat{\theta}_i$ can be obtained by plugging the numbers into Equation (14). The utility of this approach is that the predictions are obtained at any value of the moderator(s), which is akin customary regression modeling of the mean (location), but in this case for the between-study variance (scale). This approach is demonstrated below.

Heterogeneity Statistics

Quantifying heterogeneity is an important aspect of meta-analysis. Over and above reporting τ^2, this is often done by expressing heterogeneity as a proportion of the total variability, known as I^2 (Higgins & Thompson, 2002). Importantly, the derivations for I^2 are based on τ^2 and not τ_i^2 (e.g., Equation 2 in Von Hippel, 2015). To incorporate variations in τ_i^2, this entails making a modification to I^2, that is,

---

[2] In mixed-effects models more generally (of which meta-analysis is a special case), *both* σ^2 and τ^2 are typically assumed to be constant across groups.

[3] In Hoff (2009), the focus was on assuming a common σ^2 for all groups. The same logic extends to assuming a common between-study variance.

$$I_i^2 = \frac{\tau_i^2}{s^2 + \tau_i^2} \quad (16)$$

$$= \frac{\exp(\mathbf{X}_i^s \boldsymbol{\gamma})}{s^2 + \exp(\mathbf{X}_i^s \boldsymbol{\gamma})},$$

where s^2 is the so-called "'typical' within-study variance" (p. 1546, Higgins & Thompson, 2002). Note that the constant counterpart, I^2, would instead have τ^2 in Equation (16). This raises the question of choosing a value for s^2, with recent proposals including the harmonic or arithmetic mean (Table 1, Lin, 2020). The most common estimator is that of Higgins and Thompson (2002), which was "originally motivated by the expectation of the Q statistic" (p, 378 Lin, 2020). For simplicity, we also adopt that approach (Equation 9 in Higgins & Thompson, 2002), although, in our view, there could be several reasonable choices for s^2. Our proposal results in I_i^2 for each of the k studies, with the customary I^2 being a special case when $\tau_1^2 = \ldots = \tau_k^2 = \tau^2$. Further, it would be possible to obtain predictions of I_i^2, given an assumed value for s^2 and moderators in the scale sub-model, as described in Predicting Shrinkage.

Prediction Intervals

Our final extension is to prediction intervals, defined as "the range of the predicted true treatment effect in a new [future] study, given the realized (past) studies" (p. 1, Nagashima, Noma, & Furukawa, 2019). A key aspect of prediction intervals is that they incorporate heterogeneity into the uncertainty of the overall effect. Often, their variance is computed as $\tau^2 + \text{SE}(\mu)^2$, where $\text{SE}(\mu)$ is the standard error of the overall effect (see Equation 12 in Higgins et al., 2009). Hence, when $\tau^2 > 0$, they will be wider than the corresponding confidence interval. In turn, they reduce to the confidence interval for μ when $\tau^2 = 0$.

In a Bayesian framework, a future study, θ_{new}, is obtained by drawing samples from the posterior predictive distribution (Gelman et al., 1996; Meng, 1994). Typically, as described in Higgins et al. (2009), θ_{new} is sampled from

$$\theta_{\text{new}} \sim \mathcal{N}(\mu, \tau^2), \quad (17)$$

where it is clear the underlying assumption is constant heterogeneity. Relating again to the idea of having effect size estimates from (potentially) several populations, our proposal is to instead draw samples from the following distribution

$$\theta_{new,i} \sim \mathcal{N}(\mu, \tau_i^2) \quad (18)$$
$$\tau_i^2 = \exp(\mathbf{X}_i^s \boldsymbol{\gamma}),$$

which accommodates our scale sub-model for the between-study variance. Further, the basic idea of predicting λ_i and I_i^2 with moderators similarly applies to $\theta_{new,i}$, resulting in prediction intervals themselves being a function of covariates. This extension furthers Riley, Higgins, and Deeks (2011), where it was noted that reporting prediction intervals (in addition to confidence intervals) based on τ^2 "can provide a more complete summary of a random effects meta-analysis" (p. 964). We are taking this one step further, in that Equation (18) can provide an even more complete summary by accommodating heterogeneous populations.

Estimation and Software

We developed the R package **blsmeta** (Bayesian Location-Scale Meta-Analysis) specifically to (i) fit the proposed model, and (ii) compute our novel extensions to shrinkage, heterogeneity measures (e.g., I^2), and prediction intervals. Under the hood of **blsmeta**, the models are estimated with the popular Bayesian software JAGS (Plummer, 2013).

In the illustrative examples, the fitted models include four chains of 25,000 iterations after discarding a burn-in of 5,000 iterations (the default in **blsmeta**). This resulted in a total of 100,000 samples from the posterior distribution. This number of samples provided a good quality of the parameter estimates in which the models converged with potential scale reduction factors $\hat{R}$ smaller than 1.1 (Brooks & Gelman, 1998). Note that **blsmeta** provides $\hat{R}$ in the summary output, and there are diagnostics to assess convergence of the Markov chain Monte Carlo algorithm. The posterior distributions are summarized by their means and 90% credible intervals (CrI). The R code to reproduce the analyses is provided online.

Illustrative Examples

In this section, we further demonstrate the utility of modeling τ_i^2. The overall effect, μ, is not a point of emphasis. Of course, as shown in Figure 2, inference on μ can be compromised as well. Those results followed directly from getting the total variance in y_i incorrect and thus they were (perhaps) not too surprising. Further, in our view, long-run error rates are not tangible, leaving much to be desired. Our intention here is to instead highlight the importance of τ_i^2 with examples in real data.

Example 1: The Color Red and Cognitive Performance

An influential finding in psychology suggested that presenting red stimuli in an achievement context, as opposed to another color (e.g., green), resulted in worse performance (Elliot, Maier, Moller, Friedman, & Meinhardt, 2007). Although this finding has emerged elsewhere (e.g., Gnambs, Appel, & Kaspar, 2015), others have not detected the effect (e.g., Larsson & Von Stumm, 2015). In light of these conflicting results, Gnambs (2020) performed a meta-analysis of

studies examining whether the color red indeed affects cognitive performance in achievement situations. After exclusion criteria were applied, 67 effect sizes (Cohen's d) from 22 studies were included in the meta-analysis. Because some of them stemmed from the same sample, we cannot assume independence. As such, we randomly selected one effect size from each study, resulting in 22 independent effect sizes for our analysis.

Model. In this example, our goal is to further clarify the idea of modeling between-study heterogeneity. We begin by defining the location sub-model for the $i = 1, 2, \ldots, 22$ independent effect sizes, that is,

$$\mu_i = \beta_0 + u_i \tag{19}$$
$$u_i \sim \mathcal{N}(0, \tau_i^2),$$

where β_0 is the overall effect and u_i is the varying effect therein, with $\beta_0 + u_1$ denoting the hierarchical effect for the first study (i.e., $i = 1$). Note that the likelihood is given in Equation (8). In this example, we predict τ_i^2 with the sample size of each study. The idea here is to investigate whether the deviations from β_0 are a function of sample size. This was inspired by Bowater and Escarela (2013), where it was argued that several factors are "likely to cause the heterogeneity between small [clinical] trials to be greater than that between large trials" (p. 6). Our scale sub-model is given by

$$\tau_i^2 = \exp(\gamma_0 + \gamma_1 X_i^s), \tag{20}$$

where X_i is the sample size for study i. Note that centering can be very important in scale modeling, for both estimating the model and interpreting γ_0. In this case, it is not reasonable to estimate τ_i^2 when a study had zero observations ($n = 0$). To make γ_0 more interpretable, we subtracted 20 from each sample size (i.e., $X_i = n_i - 20$). In Equation (20), it follows that γ_1 is interpreted as the change in τ_i^2 on the log scale for every additional observation or per one unit increase in n. The assumed priors are described above (section Prior Specification).

For comparison purposes, we also estimated the customary random-effects model. In reference to Equation (20), this is accomplished with an intercept only model, that is, $\tau_i^2 = \exp(\gamma_0)$, which results in $\hat{\tau}_1^2 = \ldots = \hat{\tau}_{22}^2 = \hat{\tau}^2$. Herein, the fitted models are abbreviated $\mathcal{M}_c$ (constant heterogeneity) and $\mathcal{M}_v$ (varying heterogeneity). The priors were identical for parameters included in both models.

Results. Figure 3 (panel A) includes the observed effect sizes on the y-axis and the corresponding sample sizes on the x-axis. The idea here is that the deviations from the average (dotted line) can be seen to the eye. As described above, the dispersion around the overall effect is what defines τ^2. Notice that there is discernible pattern, in that the dispersion (or

"spread") of the observed effects is largest for small sample sizes and then reduces with larger sample sizes. This is what the scale sub-model in Equation (20) is predicting.

We have been focusing on τ^2, but we report τ because it is more readily interpretable (due to being on the scale of the effect sizes rather than on a squared scale). In $\mathcal{M}_v$, there was an inverse relation between n and heterogeneity ($\gamma_1 = -0.03$, 90% CrI = $[-0.05, -0.01]$). This was not captured in $\mathcal{M}_c$, where τ was assumed to be constant across all values of n ($\tau = 0.11$, 90% CrI = $[0.02, 0.28]$). To emphasize this point, we predicted τ for a range of sample sizes for both models ($n = 40, 41, \ldots, 300$). Note that computing conditional variance is a novel aspect of our methodology (see the numerator in Equation 16). Panel E includes the predicted values, $\hat{\tau}$, including 90% CrIs capturing the uncertainty. Here the difference between models is readily discernible, in that, when assuming heterogeneity is constant ($\mathcal{M}_c$), τ is the same for each sample size. On the other hand, with increasing n, $\hat{\tau}$ approached zero for $\mathcal{M}_v$. This parallels what can be seen in the observed data (panel A).

Although not our focus, it is interesting to consider the overall effect. As noted in Equation 4, with variations in τ, this can affect μ and its uncertainty. In this case, the overall effect went from "significant" in $\mathcal{M}_c$ ($\hat{\mu}$ = -0.10, 90% CrI = [-0.20, -0.02]) to "non-significant" in $\mathcal{M}_v$ ($\hat{\mu}$ = -0.07, 90% CrI = [-0.15, 0.01]). On closer inspection, however, the results were not much different when considering the entire posterior distribution—the most probable values for the true effect size are small and negative.

Hierarchical Shrinkage. We now shift gears to the study-specific effects (θ_i). Recall that these relations are obtained by "pooling of information" from the group-level. Although this is typically considered a good thing, an important assumption underlying shrinkage is that τ_i^2 is constant across all studies. In this case, because it is not constant (i.e., it is function of n), we can readily observe what happens when there is improper pooling of information.

Figure 3 (panel B) includes the study-specific effects. There is a striking difference between models, both with respect to the posterior means and uncertainty. For example, many of the 90% CrIs from $\mathcal{M}_v$ were more than three times the width of the corresponding interval from $\mathcal{M}_c$. Yet, on the other hand, several of the 90% CrIs were narrower for $\mathcal{M}_v$. This is because, in $\mathcal{M}_c$, only the sampling variances are contributing to differences in uncertainty (see Equation 13). By including the scale model, however, the estimates reflect variations in τ_i^2.

To further shed light upon shrinkage, we narrowed the focus onto the posterior means of the hierarchical estimates in relation to the observed effect size estimates. These results are provided Figure 3 (panel C). The difference between models was again striking. In $\mathcal{M}_c$, the hierarchical estimates were aggressively shrunk towards the overall effect. In $\mathcal{M}_v$,
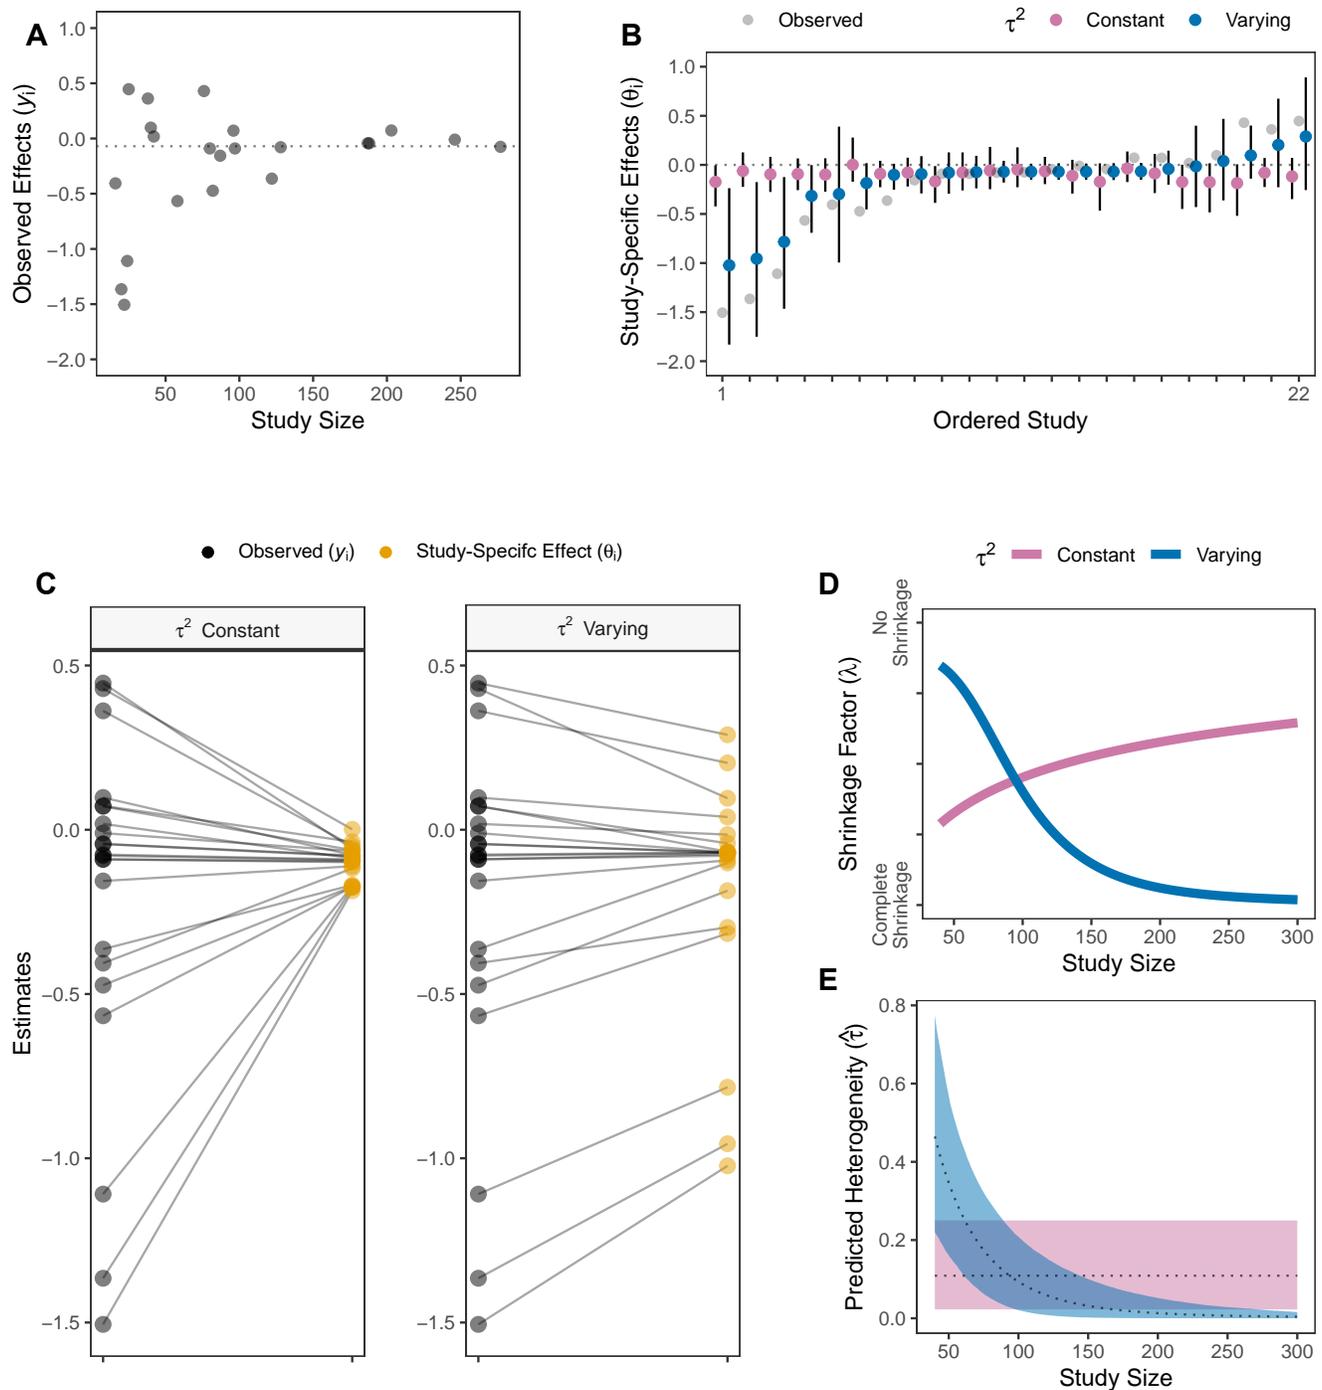
Figure 3. Panel A includes the observed effects. Notice that their dispersion around the average (dotted line) decreases when going from smaller to larger studies. The scale sub-model in Equation (19) predicted that dispersion or between-study variance (τ^2) with study size. Panel B includes the study-specific (hierarchical) effects from two models. The only difference is that τ^2 was either varying ($\mathcal{M}_v$), because it was predicted with study size, or assumed to be constant ($\mathcal{M}_c$). In panel C, the estimates reflect inappropriate shrinkage in $\mathcal{M}_c$, whereas shrinkage from $\mathcal{M}_v$ adapted to variations in τ_i^2. Panel D includes the predicted shrinkage factors (Equation 13). As expected (from panel E), in $\mathcal{M}_v$, there was less shrinkage when τ^2 was large (smaller studies) and complete shrinkage when τ^2 was very small (larger studies). On the other hand, the predictions from $\mathcal{M}_c$ showed the complete opposite pattern. Panel E includes conditional heterogeneity as a function of study size. In panels B and E, the error bars and ribbon correspond to 90% credible intervals.

while there was also shrinkage, it was much less severe. For example, the largest observed effect ($y_i = 6$) went from -1.5 to -1.05 in $\mathcal{M}_v$. The same estimate was shrunk all the way to -0.18 in $\mathcal{M}_c$. Further, in reference to panel B, the models could lead to different inferences, in that the 90% CrI excluded zero for several relations in $\mathcal{M}_v$ but not in $\mathcal{M}_c$. This was a result of some intervals being much different between models.

To explain these (huge) differences, we next predicted the shrinkage factor, $\hat{\lambda}_i$, following Equation (15) with a range of sample sizes ($n = 40, 41, \ldots, 300$). Recall that λ_i determines the amount of shrinkage towards the overall effect, with $\lambda_i = 0$ and $\lambda_i = 1$ corresponding to complete and complete shrinkage, respectively. For demonstrative purposes, we assumed a Fisher-z sampling variance $\sigma_i^2 = 1/(n-3)$. Because τ decreases with larger n (panel E), this implies that λ_i should approach zero—complete shrinkage towards the overall effect (i.e., $\theta_i = \mu$). In reference to Equation (13), this is because the numerator approaches zero. Indeed, with increasing n, the predicted values of $\hat{\lambda}_i$ approached zero for $\mathcal{M}_v$ (Figure 3, panel D). Accordingly, in large sample sizes the observed effects are predicted to converge upon the overall effect, as a result of complete shrinkage. In $\mathcal{M}_c$, however, the *exact opposite* pattern emerged. With large sample sizes, the model predicted less shrinkage towards the overall effect (perhaps no shrinkage at all with a large enough sample size). Together, this indicates that models differing in their between-study variance structures can provide much different estimates (e.g., as seen in panels B and C)

Example 2: Cross-national Comparisons of Racial Discrimination in Hiring

In this example, we highlight the utility of our methodology for answering novel research questions. This is juxtaposed to the previous example, where predicting τ_i^2 with n was mainly for pedagogical purposes, in that it was readily apparent when simply looking at the data and motivated by the meta-analysis literature (Bowater & Escarela, 2013). The following analysis builds directly upon the questions asked in the primary source (Quillian et al., 2019).

In an effort to better understand disparities across racial and ethnic groups play out in hiring practices, Quillian et al. (2019) conducted a meta-analysis of 97 studies aimed at determining variation in hiring discrimination across nine countries in Europe and North America. The included studies had comparable designs in that majority and minority groups were contrasted using applications for fictional job-seekers that were equivalent in their labor-market relevant characteristics. The resulting 159 effect sizes from these studies were coded as discrimination ratios (a.k.a "relative risk" or "risk ratio") on the log scale. The discrimination ratio is defined as the proportion of call-backs that the majority group received divided by the proportion of call-backs

received by a minority group. Intuitively, a discrimination ratio greater than 1 (or 0 on the log scale) indicates that the majority group received more positive responses than the minority group.

It is important to note that some studies in these data provide multiple effect sizes and are thus non-independent. To account for these dependencies, (Quillian et al., 2019) used robust standard errors for the fixed-effects, but a traditional random-effects model to estimate τ^2. As in the previous example, we eliminate these dependencies by randomly selecting one effect size from each study. We fit a random-effects model to the remaining 97 effects, but as we have done throughout this paper, allow for variations in heterogeneity.

Model 1. We estimated two location-scale models, differing in their mean structures (with and without a moderator). Because both have the same likelihood as Equation (8), we jump right to the model formulations.

For the $i = 1, 2, \ldots, 97$ independent effect sizes, the mean structure for model 1, ($\mathcal{M}_{1v}$), is given by

$$\mu_i = \beta_0 + \sum_{j=1}^{8} \beta_j X_{ij}^l + u_i \qquad (21)$$
$$u_i \sim \mathcal{N}(0, \tau_i^2)$$

where X is the dummy coded design matrix that includes indicators for each country. Hence, β_0 is the overall effect for the reference category (Belgium). We are not interested in the regression coefficients themselves, but instead the country-specific effects. From Equation (21), we thus obtain $\hat{\mu}$ for each of the nine countries. For example, $\beta_0 + \beta_j$ ($j = 1, 2, \ldots, 8$) corresponds to the overall effect for each country.

We then predict τ_i^2 with country as well, that is,

$$\tau_i^2 = \exp\left(\gamma_0 + \sum_{j=1}^{8} \gamma_j X_{ij}^s\right) \qquad (22)$$

In this case, γ_0 corresponds to heterogeneity on the log-scale for the reference category (Belgium). We are again not interested in the contrasts from the reference category, but instead the country-specific estimates of τ^2. For example, $\gamma_0 + \gamma_j$ ($j = 1, 2, \ldots, 8$) corresponds to the heterogeneity for each country. Together, this *baseline* model effectively allows each country to have their own overall effect and heterogeneity, whereas the latter was assumed constant for each country in Quillian et al. (2019).

Model 2. It is common practice to fit a baseline model without moderators for the overall effect and then compare it to a model that includes moderators. The idea is to determine the amount of heterogeneity explained by the moderator, resulting in pseudo R^2. We pursue that approach here, but with

an important extension. Due to the scale sub-model in Equation (22), each country is permitted to have their own variance explained (i.e., R_i^2). Accordingly, by adding a predictor to Equation (21), we can determine how much variance is explained for each country—perhaps a moderator has explanatory power for some but not other countries—which can be answered with our methodology.

In model 2, $\mathcal{M}_{2v}$, the one modification to $\mathcal{M}_{1v}$ is the location sub-model in Equation (21). We included an additional predictor that corresponds to the proportion of the local (where the study took place) population that was immigrants, with the coefficient denoted β_9. The models were otherwise identical. As a result, when computing R_i^2 for each country, we can determine how much variance in the log discrimination ratios is explained by how many immigrants there are in the local area. Again, because we are allowing for variation in τ_i^2, this implies that there can be varying variance explained, in this case perhaps as a function of country. This is but one example of the kinds of substantive questions that can be answered with location-scale modeling.

For comparison purposes, we also estimated the same models as above, but assuming that τ^2 was constant across the studies (one value for all countries). These are denoted by $\mathcal{M}_{1c}$ and $\mathcal{M}_{2c}$, respectively. These comparison are again qualitative in nature and are meant to highlight the importance of thinking about variations in τ_i^2. Later on, we discuss the topic of formal model comparison and selection. The priors were identical in all models (section Prior Specification). The estimation information is provide above (section Estimation and Software).

Results. We start by looking at the observed effects, y_i (x-axis), in relation to country (y-axis). The idea is to notice the dispersion ("spread") of y_i within each country (Figure 4, panel A). To the eye, there are striking differences, with, for example, those effects from Germany spanning from 1 to 1.5, whereas those from France had larger variation, spanning from 1 to 5. This goes back to the notion of having studies from superpopulations that differ in their scale (i.e., unequal between-study variance). In this case, it appears the studies are not sampled from a common population.

Of course, we cannot rely on visual inspection alone to infer differences between countries in the dispersion of their effects. Figure 4 (panel B) includes the country-specific estimates of τ, in addition to that obtained from $\mathcal{M}_{1c}$ (constant across studies). While not shown here, there was "significantly" more heterogeneity in those effects from France than several of the other countries.[4] This indicates that there are variations in τ_i^2 that would be masked by assuming constant heterogeneity. As shown earlier (e.g., Figure 2), this can impact inference on the overall effect, which follows directly from considering the total variance of y_i (Equations 3 and 6): if τ^2 is underestimated (overestimated) for a given country, then the interval for the overall effect will be too narrow (wide).

Figure 4 (panel C) includes I_i^2 that corresponds to the proportion of variability that is due to heterogeneity. This was computed with Equation (16), where each country had their own "typical" within-study variance (s^2). This is effectively what would be obtained if separate models were estimated for each country. The results reveal that a common I^2 would not adequately describe each country. Further, when taking differences, I^2 for both Sweden and France was notably larger than I^2 from some of the other countries.

The next question is how variation in τ_i^2, when ignored, affects inference for the overall effect of discrimination. Here we focus on predictive inference. Figure 4 (panel D) includes 90% prediction intervals for $\mathcal{M}_{1v}$ and $\mathcal{M}_{1c}$. Recall from above that prediction intervals include additional uncertainty due to heterogeneity. In this case, if an interval excludes zero there is a 95% posterior probability (conditional on the model) that a future study will have a positive effect. There are notable differences between models, especially when looking to reject zero (not an approach we recommend in practice). For example, in three countries a different conclusion would be reached altogether. These discrepant inferences map directly onto the estimate of τ^2 in Figure 4 (panel B). In countries where τ^2 was underestimated in $\mathcal{M}_{1c}$ (e.g., France), that is, the corresponding prediction intervals are narrower compared to those from $\mathcal{M}_{1v}$. This again follows directly from what is assumed to be the total variance of y_i (e.g., Equation 6).

Variance Explained. To answer a novel research question, we added a moderator to the location sub-model of $\mathcal{M}_{1v}$. The idea is to determine country-specific variance explained due to the proportion of the local population that was immigrants. This example therefore demonstrates how location-scale modeling can provide unique insights, in this case into discrimination across nine countries in Europe and North America.

Figure 4 (panel E) includes the posterior distribution for the moderator effect (denoted β_9). The models could again lead to different inferences, in that the posterior means and credible intervals were *much* different ($\mathcal{M}_{2v}$: β_9 = 0.98, 90% CrI = [0.39, 1.53] versus $\mathcal{M}_{2c}$: β_9 = 0.68, 90% CrI = [−0.17, 1.52]). This seems to beg the question of which model is "better." We intentionally forgo model selection, because, in our view, this topic deserves more attention than we can give it here (see the Discussion).

Figure 4 (panel F) includes the estimates of pseudo R^2. The results are quite intriguing. Notice that R^2 computed from $\mathcal{M}_{2c}$ was small ($\approx$ 0.02), perhaps reflecting that β_9 was not "significant" when assuming constant heterogeneity (panel E). In $\mathcal{M}_{2v}$, a much different picture emerged. Here variance explained ranged from quite a bit (Norway and

---

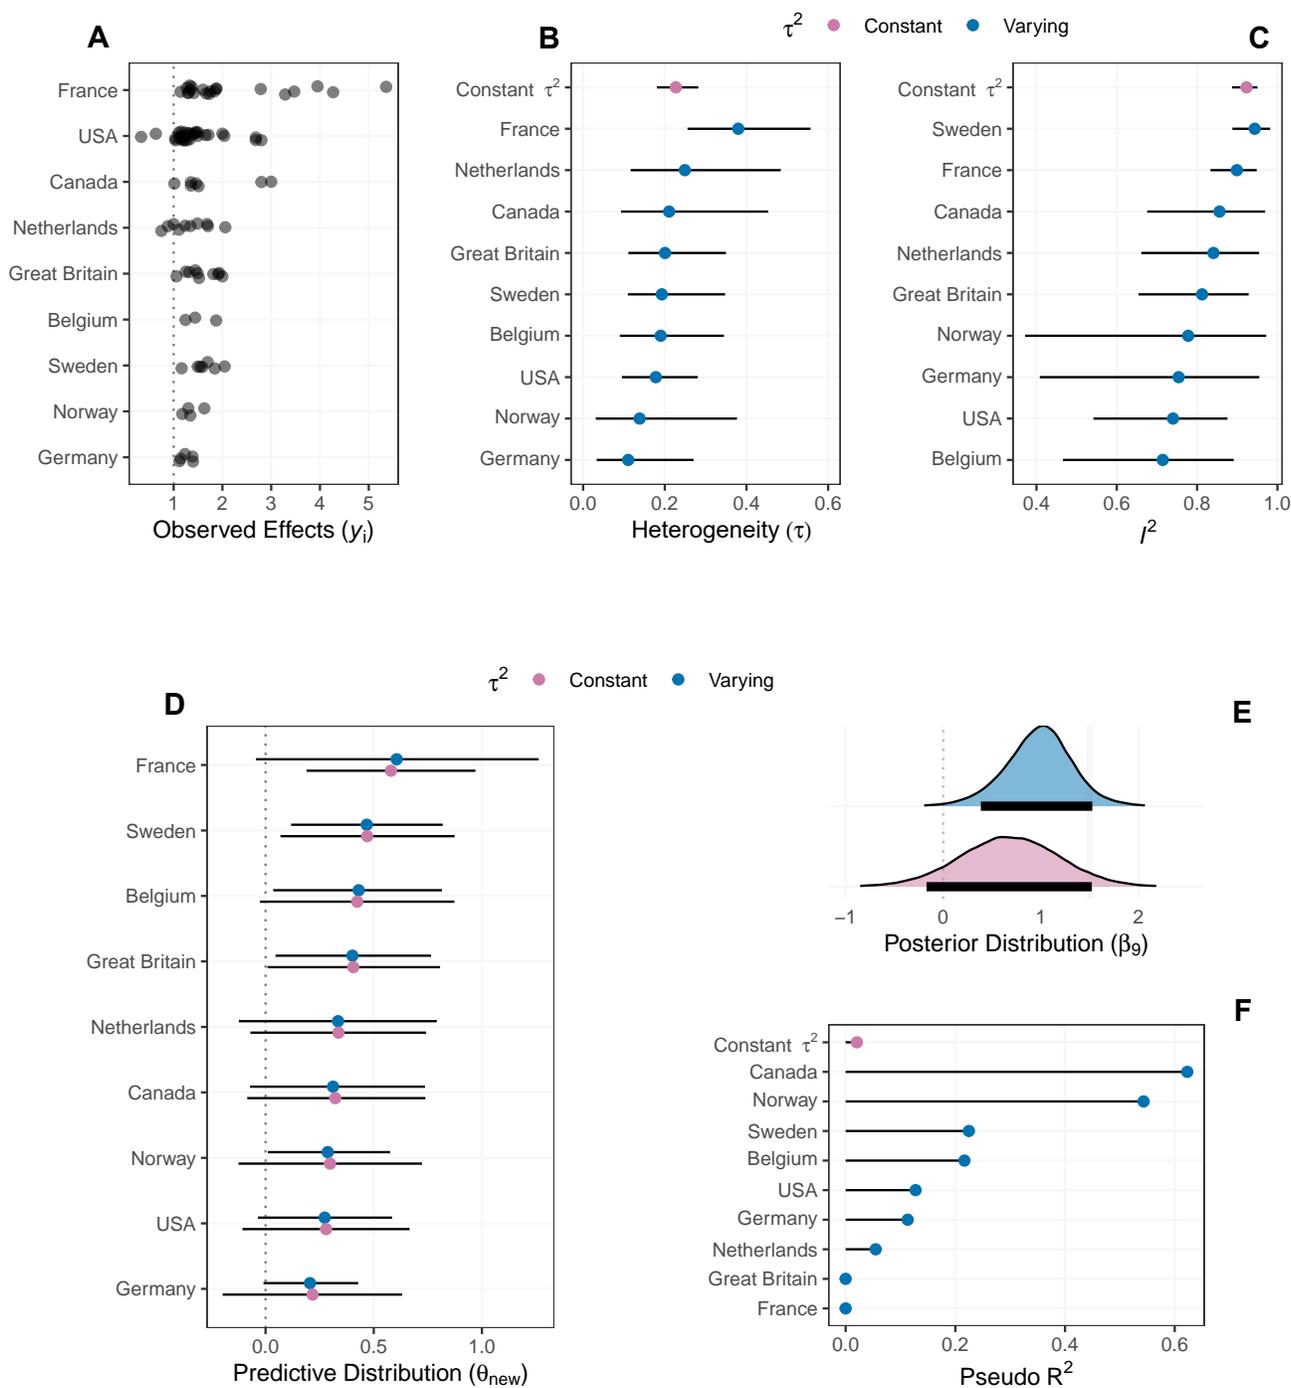[4]This was computed as the difference between the posterior distributions of τ^2.

Figure 4. Panel A includes the observed effects (discrimination ratio). Notice that the effects from France appear more heterogeneous than those from Germany. The scale sub-model in Equation (22) predicted that heterogeneity (τ^2) with country. Panels B and C include the estimates of τ and I^2 (Equation 16) from two models. The only difference is that τ^2 was either varying ($\mathcal{M}_v$), because it was predicted with country, or assumed to be constant across countries ($\mathcal{M}_c$). Panel D includes prediction intervals, where the models would lead to different inferences for several countries, assuming that the goal was to reject zero. In panels E and F, each model included a moderator in the location sub-model (the proportion of the local areas population that was immigrants). The posterior distributions for that parameter, β_9, are in panel E, again with the models leading to different inferences. Panel F includes variance explained by the moderator for each country, which apparently explains a lot of variance for some countries and none for others. The error bars correspond to 90% credible intervals.

Canada) to none at all (Great Britain and France). There is a caveat, in that we are relying on mere point estimates which means we cannot infer which countries are actually different from one another. Further, it is well known that computing uncertainty for pseudo R^2 is non-trivial. What does seem reasonable, however, is that there are variations in R_i^2 that would be missed by assuming a constant between-study variance for each country.

Discussion

We set out to demonstrate the importance of considering variations between-study variance. At the outset, we suggested that going beyond τ^2 and to τ_i^2 has serious implications for meta-analytic practice in psychological science (and beyond). In a series of examples, we demonstrated that *everything* changes when τ_i^2 is not constant across studies. This makes sense, because it implies the studies are sampled from various populations that differ in their heterogeneity. The goal is to then make inference, given that there are variations in τ_i^2, which has not been considered in the psychological literature. The changes were not small and passed the interocular trauma test—the importance hits right between the eyes.

To summarize, the following points towards the importance of modeling τ_i^2:

1. Heterogeneity can be a function of the study size (Figure 3), as argued in Bowater and Escarela (2013). While not reported here, we found this to be the case in many other datasets. This suggest that meta-analyses may routinely have the total variance in y_i incorrect, which can compromise statistical inference (Figure 2).

2. Assuming a constant τ^2, when this is not the case, can have a substantial impact on the study-specific estimates and their uncertainty (Figure 3, panel B). Further, this can result in improper "borrowing of information," given that a critical assumption of hierarchical shrinkage is constant heterogeneity (although this can be relaxed; Figure 3, panel C).

3. Variations in τ_i^2 can be incorporated into heterogeneity statistics (Figure 4, panel C) and prediction intervals (Figure 4, panel D). In essence, when predicting τ_i^2 with moderators, this implies that all quantities based on heterogeneity are then a function of those same moderators.

4. Computing conditional between-study variance opens the door to investigating consistency of the effects with an uncanny level of detail. This was demonstrated in Figure 3 (panel E), where $\hat{\tau}_i^2$ was predicted to be large in small studies and essentially zero in large studies.

5. Although not our focus, scale modeling has an appreciable effect on the location model (e.g., moderators for the effect size). In each example, assuming the goal was to reject zero,[5] several inferences would have changed altogether when predicting τ_i^2 (e.g.. Figure 4, panel D and E).

The First (and not the) Last Word

This work is the first word on the implications τ_i^2. Our results do, however, provide the foundation for future research on this topic (including simulation studies). Here we offer one suggestion. Often, at the end of simulation-based papers, the take home is that the "true" model is the winner, assuming that there is enough data to estimate it. This is obvious, in that a model accounting for unequal variances will typically be better (e.g., having the nominal type one error rate) than one that does not. This applies to model misspecification more generally. More useful, in our view, is to investigate situations in which the scale model can *always* be included (whether it is "true" or not). For example, when investigating whether the effect size is moderated, perhaps that same moderator can be included for the scale. This suggestion is analogous to using the unequal variance *t*-test by default (Delacre et al., 2017).

Additionally, an open question is the number of studies required for location-scale meta-analysis. Here the goal at hand should be considered. For example, if the goal is to detect differences in heterogeneity, then many studies will likely be needed. This is similar to moderators for the effect size, in that those tests are often under powered (due to few studies). Alternatively, when the goal is to account for variations in τ_i^2, then power is less of a concern. This is similar in spirit to the recommendation for choosing a random versus fixed-effects meta-analysis, as noted in Viechtbauer (2010), which should be "the type of inference desired before examining the data" (p. 5).

We intentionally avoided formal model comparison. This decision was made because we wanted focus on introducing the method. Here, too, there is much opportunity for future research. The presented models can be compared with approximate leave-one-out cross-validation (Gelman, Hwang, & Vehtari, 2014; Vehtari, Gelman, & Gabry, 2016) or the Bayes factor (Jeffreys, 1961). In the future, both options will be implemented in **blsmeta**.

Prior Specification

We anticipate that the adopted priors will work well in a variety of situations. An issue can arise when there are few studies. This is not particular to Bayesian methods, as estimating heterogeneity is difficult in these situations for "frequentist" methods as well (Veroniki et al., 2016). Pragmati-

---

[5]We do not recommend this approach to statistical inference.

cally, one solution is to forgo unconditional inference and to simply use the fixed-effects model. Although the inference does not generalize to a (hypothetical) population of studies, perhaps this is a bridge too far with few studies anyhow. Alternatively, we refer interested readers to Williams, Rast, and Bürkner (2018), who investigated the properties of weakly informative prior distributions in meta-analysis.

Often times Bayesian analysis is framed as estimation (Kruschke & Liddell, 2017) or hypothesis testing (but see Rouder, Haaf, & Vandekerckhove, 2018). This work is in the traditions of estimation. Here the priors are important to consider, but typically the estimates will be similar to "frequentist " methods when using weakly informative priors, particularly when the number of studies is large. On the other hand, in Bayesian hypothesis testing, the prior can play a much more important role, sometimes leading to a different conclusion than a null hypothesis significance test (Lindley, 1957). This is not to say it should be avoided altogether. Rather, it should be noted that formulating meaningful priors for parameters corresponding to changes in variance (on the log scale) is a difficult task. Going forward, we plan to develop a framework for Bayesian testing in location-scale meta-analysis.

Conclusion

Between-study heterogeneity has traditionally been considered constant across the studies included in a random-effects meta-analysis. This implies that the studies are sampled from the *same* population. Our Bayesian location-scale model relaxes that assumption. The illustrative examples highlighted that there can be substantial variations in heterogeneity, which necessarily influences the parameter estimates, can be incorporated into commonly used heterogeneity statistics, and leads to unique inferences. The methodology that made these insights possible is implemented in the new R package **blsmeta**.

References

Aguinis, H., Gottfredson, R. K., & Wright, T. A. (2011). Best-practice recommendations for estimating interaction effects using meta-analysis. *Journal of Organizational Behavior*, *32*(8), 1033–1043.

Aitkin, M. (1987, 6). Modelling Variance Heterogeneity in Normal Regression Using GLIM. *Applied Statistics*, *36*(3), 332. doi: 10.2307/2347792

Blanca, M. J., Alarcón, R., Arnau, J., Bono, R., & Bendayan, R. (2018, 6). Effect of variance ratio on ANOVA robustness: Might 1.5 be the limit? *Behavior Research Methods*, *50*(3), 937–962. doi: 10.3758/s13428-017-0918-2

Bonett, D. G. (2010). Varying coefficient meta-analytic methods for alpha reliability. *Psychological methods*, *15*(4), 368.

Borenstein, M., Hedges, L., & Rothstein, H. (2007). Meta-analysis: Fixed effect vs. random effects. *Meta-analysis. com*.

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

Bowater, R. J., & Escarela, G. (2013). Heterogeneity and study size in random-effects meta-analysis. *Journal of Applied Statistics*, *40*(1), 2–16.

Brooks, S. P. B., & Gelman, A. G. (1998). General methods for monitoring convergence of iterative simulations. *Journal of computational and graphical statistics*, *7*(4), 434–455. doi: 10.2307/1390675

Bürkner, P.-C. (2017). brms : An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, *80*(1). doi: 10.18637/jss.v080.i01

Cleveland, W. S., Denby, L., & Liu, C. (2003). Random scale effects. (2), 33. Retrieved from stat.bell-labs.com

Delacre, M., Lakens, D., & Leys, C. (2017). Why Psychologists Should by Default Use Welch's t-test Instead of Student's t-test. *International Review of Social Psychology*, *30*(1), 92–101. doi: 10.5334/irsp.82

Efron, B., & Morris, C. (1977). Stein's paradox in statistics. *Scientific American*, *236*(5), 119–127.

Elliot, A. J., Maier, M. A., Moller, A. C., Friedman, R., & Meinhardt, J. (2007). Color and psychological functioning: the effect of red on performance attainment. *Journal of experimental psychology: General*, *136*(1), 154.

Gelman, A. (2019). Ethics in statistical practice and communication: Five recommendations. In *The best writing on mathematics 2019* (pp. 214–223). Princeton University Press.

Gelman, A., Gelman, A., Meng, X.-L., Meng, X.-L., Stern, H., & Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. Vol.6, No.4. *Statistica Sinica*, *6*(4), 733–807. doi: 10.1.1.142.9951

Gelman, A., Hwang, J., & Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, *24*(6), 997–1016. doi: 10.1007/s11222-013 -9416-2

Gelman, A., Simpson, D., & Betancourt, M. (2017). The prior can generally only be understood in the context of the likelihood. *ArXiv preprint*.

Gelman, A., Tjur, T., McCullagh, P., Hox, J., Hoijtink, H., & Zaslavsky, A. M. (2005). Discussion paper analysis of variance - Why it is more important than ever. *Annals of Statistics*, *33*(1), 1–53. doi: 10.1214/009053604000001048

Gelman, A., Vehtari, A., Simpson, D., Margossian, C. C., Carpenter, B., Yao, Y., . . . Modrák, M. (2020). Bayesian workflow. *arXiv preprint arXiv:2011.01808*.

Gnambs, T. (2020). Limited evidence for the effect of red color on cognitive performance: A meta-analysis. *Psychonomic bulletin & review*, *27*(6), 1374–1382.

Gnambs, T., Appel, M., & Kaspar, K. (2015). The effect of the color red on encoding and retrieval of declarative knowledge. *Learning and Individual Differences*, *42*, 90–96.

Gronau, Q. F., Van Erp, S., Heck, D. W., Cesario, J., Jonas, K. J., & Wagenmakers, E.-J. (2017). A Bayesian model-averaged meta-analysis of the power pose effect with informed and default priors: the case of felt power. *Comprehensive Results in Social Psychology*, *2*(1), 123–138. doi: 10.1080/

23743603.2017.1326760

Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Batailler, C., Birt, A. R., . . . others (2016). A multi-lab preregistered replication of the ego-depletion effect. *Perspectives on Psychological Science*, *11*(4), 546–573.

Hedeker, D., Mermelstein, R. J., & Demirtas, H. (2008). An application of a mixed-effects location scale model for analysis of ecological momentary assessment (EMA) data. *Biometrics*, *64*(2), 627–634. doi: 10.1111/j.1541-0420.2007.00924.x

Hedeker, D., Mermelstein, R. J., & Demirtas, H. (2012). Modeling between-subject and within-subject variances in ecological momentary assessment data using mixed-effects location scale models. *Statistics in Medicine*, *31*(27), 3328–3336. doi: 10.1002/sim.5338

Hedges, L. V. (1992). Meta-analysis. *Journal of Educational Statistics*, *17*(4), 279–296.

Hedges, L. V., & Pigott, T. D. (2001). The power of statistical tests in meta-analysis. *Psychological Methods*, *6*(4), 203–217. doi: 10.1037/1082-989x.6.3.203

Hedges, L. V., & Vevea, J. L. (1998). Fixed- and random-effects models in meta-analysis. *Psychological Methods*, *3*(4), 486–504. doi: 10.1037/1082-989X.3.4.486

Higgins, J., & Thompson, S. (2002). Quantifying heterogeneity in meta - analysis. *Statistics in Medicine*, *21*(11), 1539–1558.

Higgins, J., Thompson, S. G., & Spiegelhalter, D. J. (2009). A re-evaluation of random-effects meta-analysis. *Journal of the Royal Statistical Society. Series A: Statistics in Society*, *172*(1), 137–159. doi: 10.1111/j.1467-985X.2008.00552.x

Hoff, P. D. (2009). *A first course in Bayesian statistical methods* (Vol. 580). Springer.

IntHout, J., Ioannidis, J. P., Rovers, M. M., & Goeman, J. J. (2016). Plea for routinely presenting prediction intervals in meta-analysis. *BMJ open*, *6*(7), e010247. doi: 10.1136/bmjopen-2015-010247

Jeffreys, H. (1961). *The theory of probability*. Oxford: Oxford University Press.

Kalin, A. M., Pfluger, M., Gietl, A. F., Riese, F., Jäncke, L., Nitsch, R. M., & Hock, C. (2014). Intraindividual variability across cognitive tasks as a potential marker for prodromal Alzheimers disease. *Frontiers in Aging Neuroscience*, *6*, 147. doi: 10.3389/fnagi.2014.00147

Kelley, T. L. (1927). *The Interpretation of Educational Measurements*. New York: World Book.

Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Adams Jr, R. B., Alper, S., . . . others (2018). Many Labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, *1*(4), 443–490.

Kruschke, J. K., & Liddell, T. M. (2017). The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*. doi: 10.3758/s13423-016-1221-4

Langan, D., Higgins, J. P., & Simmonds, M. (2017). Comparative performance of heterogeneity variance estimators in meta-analysis: a review of simulation studies. *Research Synthesis Methods*, *8*(2), 181–198. doi: 10.1002/jrsm.1198

Larsson, E. E. C., & Von Stumm, S. (2015). Seeing red? The effect of colour on intelligence test performance. *Intelligence*, *48*, 133–136.

Lee, Y., & Nelder, J. A. (2006). Double hierarchical generalized linear models (with discussion). *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, *55*(2), 139–185.

Lin, L. (2020). Comparison of four heterogeneity measures for meta-analysis. *Journal of Evaluation in Clinical Practice*, *26*(1), 376–384. doi: 10.1111/jep.13159

Lindley, D. V. (1957). A STATISTICAL PARADOX. *Biometrika*, *44*(1-2), 187–192. doi: 10.1093/biomet/44.1-2.187

MacDonald, S. W. S., Hultsch, D. F., & Dixon, R. A. (2003). Performance variability is related to change in cognition: evidence from the Victoria Longitudinal Study. *Psychology and aging*, *18*(3), 510.

Meng, X.-L. (1994). Posterior Predictive p-Values. *The Annals of Statistics*, *22*(3), 1142–1160. doi: 10.1214/aos/1176325622

Nagashima, K., Noma, H., & Furukawa, T. A. (2019). Prediction intervals for random-effects meta-analysis: A confidence distribution approach. *Statistical methods in medical research*, *28*(6), 1689–1702.

Plummer, M. (2013). *JAGS: Just Another Gibs Sampler*. Retrieved from http://mcmc-jags.sourceforge.net/

Quillian, L., Heath, A., Pager, D., Midtbøen, A. H., Fleischmann, F., & Hexel, O. (2019). Do some countries discriminate more than others? Evidence from 97 field experiments of racial discrimination in hiring. *Sociological Science*, *6*, 467–496.

Ram, N., Stollery, B., Rabbitt, P., & Nesselroade, J. R. (2005). Cognitive performance inconsistency: Intraindividual change and variability. *Psychology and Aging*, *20*(4), 623–633. doi: 10.1037/0882-7974.20.4.623

Rast, P. (2011). Verbal Knowledge, Working Memory, and Processing Speed as Predictors of Verbal Learning in Older Adults. *Developmental Psychology*, *47*(5), 1490–1498. doi: 10.1037/a0023422

Raudenbush, S. W., & Bryk, A. S. (1985). Empirical Bayes meta-analysis. *Journal of educational statistics*, *10*(2), 75–98.

Riley, R. D., Higgins, J. P., & Deeks, J. J. (2011). Interpretation of random effects meta-analyses. *Bmj*, *342*(7804), 964–967. doi: 10.1136/bmj.d549

Rouder, J. N., Haaf, J. M., & Vandekerckhove, J. (2018). Bayesian inference for psychology, part IV: parameter estimation and Bayes factors. *Psychonomic Bulletin and Review*, *25*(1), 102–113. doi: 10.3758/s13423-017-1420-7

Röver, C. (2020). Bayesian random-effects meta-analysis using the bayesmeta R package. *Journal of Statistical Software*, *93*(6), 1–51.

Rubio-Aparicio, M., López-López, J. A., Viechtbauer, W., Marin-Martinez, F., Botella, J., & Sánchez-Meca, J. (2020). Testing categorical moderators in mixed-effects meta-analysis in the presence of heteroscedasticity. *The Journal of Experimental Education*, *88*(2), 288–310.

Ruscio, J., & Roche, B. (2012). Variance heterogeneity in published psychological research: A review and a new index. *Methodology*, *8*(1), 1–11. doi: 10.1027/1614-2241/a000034

Shieh, G. (2016). Choosing the best index for the average score intraclass correlation coefficient. *Behavior Research Methods*, *48*(3), 994–1003. doi: 10.3758/s13428-015-0623-y

Stein, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In *Proceedings*

of the third berkeley symposium on mathematical statistics and probability, 1954–1955, vol. i (pp. 197–206). University of California Press, Berkeley and Los Angeles.

Thompson, C. G., & Becker, B. J. (2020). A group-specific prior distribution for effect-size heterogeneity in meta-analysis. *Behavior research methods*, 1–11.

Thompson, S. G., & Sharp, S. J. (1999). Explaining heterogeneity in meta-analysis: A comparison of methods. *Statistics in Medicine*, *18*(20), 2693–2708. doi: 10.1002/(SICI)1097 -0258(19991030)18:20<2693::AID-SIM235>3.0.CO;2-V

van Agteren, J., Iasiello, M., Lo, L., Bartholomaeus, J., Kopsaftis, Z., Carey, M., & Kyrios, M. (2021). A systematic review and meta-analysis of psychological interventions to improve mental wellbeing. *Nature Human Behaviour*, 1–22.

van Erp, S., Verhagen, J., Grasman, R., & Wagenmakers, E. (2013). Estimates of Between-Study Heterogeneity for 705 Meta-Analyses Reported in Psychological Bulletin From 1990-2013. (1), 1–18.

Vehtari, A., Gelman, A., & Gabry, J. (2016). *Efficient leave-one-out cross-validation and WAIC for Bayesian models.* Retrieved from https://github.com/jgabry/loo

Veroniki, A. A., Jackson, D., Viechtbauer, W., Bender, R., Bowden, J., Knapp, G., . . . Salanti, G. (2016). Methods to estimate the between-study variance and its uncertainty in meta-analysis. *Research Synthesis Methods*, *7*(1), 55–79. doi: 10.1002/jrsm.1164

Viechtbauer, W. (2010). Conducting Meta-Analyses in R with the metafor Package. *Journal of Statistical Software*, *36*(3), 1–48. doi: 10.18637/jss.v036.i03

Von Hippel, P. T. (2015). The heterogeneity statistic I2 can be biased in small meta-analyses. *BMC Medical Research Methodology*, *15*(1), 1–8. doi: 10.1186/s12874-015-0024-z

Williams, D. R., Rast, P., & Bürkner, P.-C. (2018). Bayesian Meta-Analysis with Weakly Informative Prior Distributions. doi: https://doi.org/10.31234/osf.io/7tbrm