



Bayesian hypothesis testing and estimation under the marginalized random-effects meta-analysis model

Robbie C. M. van Aert¹ · Joris Mulder¹

Accepted: 11 March 2021 / Published online: 22 June 2021
© The Author(s) 2021

Abstract

Meta-analysis methods are used to synthesize results of multiple studies on the same topic. The most frequently used statistical model in meta-analysis is the random-effects model containing parameters for the overall effect, between-study variance in primary study's true effect size, and random effects for the study-specific effects. We propose Bayesian hypothesis testing and estimation methods using the marginalized random-effects meta-analysis (MAREMA) model where the study-specific true effects are regarded as nuisance parameters which are integrated out of the model. We propose using a flat prior distribution on the overall effect size in case of estimation and a proper unit information prior for the overall effect size in case of hypothesis testing. For the between-study variance (which can attain negative values under the MAREMA model), a proper uniform prior is placed on the proportion of total variance that can be attributed to between-study variability. Bayes factors are used for hypothesis testing that allow testing point and one-sided hypotheses. The proposed methodology has several attractive properties. First, the proposed MAREMA model encompasses models with a zero, negative, and positive between-study variance, which enables testing a zero between-study variance as it is not a boundary problem. Second, the methodology is suitable for default Bayesian meta-analyses as it requires no prior information about the unknown parameters. Third, the proposed Bayes factors can even be used in the extreme case when only two studies are available because Bayes factors are not based on large sample theory. We illustrate the developed methods by applying it to two meta-analyses and introduce easy-to-use software in the R package `BFpack` to compute the proposed Bayes factors.

Keywords Bayes factor · Meta-analysis · Random-effects model · Heterogeneity

Introduction

The rapidly expanding scientific literature calls for methods to synthesize research such as a systematic review. Meta-analysis is an important part of a systematic review and refers to applying statistical methods to combine findings of different studies on the same topic. The first step of a meta-analysis is to obtain a standardized effect size (e.g., standardized mean difference or correlation coefficient) for each included study and these effect sizes are subsequently combined by means of meta-analysis methods to summarize the included studies. Meta-analysis is nowadays seen as the gold standard for research synthesis (Aguinis et al., 2011; Head et al., 2015) and is often used for policy making as its results

are seen as best available evidence (Thompson & Sharp, 1999; Cordray & Morphy, 2009; Hedges & Olkin, 1985).

The random-effects model is the most commonly used statistical model in meta-analysis (Borenstein et al., 2010; 2009). In the random-effects model, each study is assumed to have an unknown true underlying effect size. The main parameter of interest in this model is the overall effect size of the studies included in the meta-analysis. However, estimation and statistical inferences for the between-study variance in the random-effects model is just as important (Higgins et al., 2009) because both parameters focus on distinct and relevant aspects of a meta-analysis. For example, the overall effect size in a meta-analysis on the efficacy of a psychological treatment refers to the average efficacy of the treatment across studies and the between-study variance quantifies the heterogeneity across studies.

The vast majority of meta-analyses uses frequentist analysis techniques for estimation and drawing statistical inferences. However, Bayesian meta-analysis methods have been proposed as well and have gained in popularity (Xu

✉ Robbie C. M. van Aert
r.c.m.vanaert@tilburguniversity.edu

¹ Department of Methodology and Statistics, Tilburg University, Tilburg, the Netherlands

et al., 2008). Bayesian methods are especially well suited for analyzing meta-analytic data (Smith et al., 1995; Sutton & Abrams, 2001; Lunn et al., 2013; Turner & Higgins, 2019) because the multilevel structure of a random-effects meta-analysis can be straightforwardly taken into account. Moreover, estimation of the between-study variance is imprecise using frequentist meta-analysis methods in the common situation of a meta-analysis containing a small number of studies (Chung, Rabe-Hesketh, & Choi, 2013; Sidik & Jonkman, 2007; Kontopantelis, Springate, & Reeves, 2013). Bayesian meta-analysis methods are advantageous in this situation because (i) externally available information about the between-study variance can be incorporated in the prior distribution if available, and (ii) the methodology does not directly rely on large sample theory.

Meta-analysts generally want to estimate and conduct hypothesis tests for the parameters in the random-effects model. The vast majority of the literature on Bayesian meta-analysis methods has been focused on parameter estimation using either empirical Bayes or fully Bayesian estimation (e.g., (Turner et al., 2015; Rhodes et al., 2015; Normand, 1999; Lambert et al., 2002)). However, hypothesis testing using Bayes factors has also been proposed, which quantifies evidence of one model relative to another model (Kass & Raftery, 1995). (Berry, 1998) proposed a Bayes factor for testing the null hypothesis of no between-study variance in a meta-analysis of studies using 2x2 contingency tables. Rouder and Morey (2011) developed a Bayes factor to test the null hypothesis of no effect in a meta-analysis of studies using a two-independent groups design. (Scheibehenne et al., 2017) and (Gronau et al., 2017) proposed a Bayesian model averaging approach to compute an average Bayes factor for testing the null hypothesis of no effect. In this approach, an average Bayes factor is computed by averaging over posterior model probabilities obtained with the random-effects model and the equal-effect model (a.k.a. fixed-effect or common-effect model) where the study's true effect sizes are assumed to be homogeneous.

The first contribution of our paper is that we propose, in contrast to existing Bayesian meta-analysis methods, to use a *marginalized* random-effects meta-analysis (MAREMA) model rather than the random-effects model. In this MAREMA model, the study-specific true effects are regarded as nuisance parameters and integrated out of the probability density function. The elimination of nuisance parameters via integration is common in integrated likelihood methods (e.g., Berger et al. (1999)), and has recently been extended to marginal random intercept models (Mulder & Fox, 2013; 2019), and marginal item response theory models (Fox et al., 2017).

The proposed MAREMA model encompasses three important meta-analysis models. First, the equal-effect model is included if the between-study variance is equal to

zero. Second, the random-effects model is included when the between-study variance is positive, which implies that the differences between studies' effect sizes cannot be fully explained by sampling error. Third, the random-effects model in case of a negative between-study variance is also included. A negative between-study variance indicates that the differences between studies' effect sizes are smaller than expected based on sampling error. A negative between-study variance may yield relevant insights for meta-analysts because it may indicate that the assumption of independence of primary studies in a meta-analysis is violated or that the computation of the effect sizes of the primary studies is incorrect (Ioannidis et al., 2006).

Our proposed methodology is also distinctive from other Bayesian meta-analysis methods because we place a prior distribution on the proportion of total variance that can be attributed to between-study variance rather than directly on the between-study variance parameter. This proportion is known as I^2 (Higgins & Thompson, 2002; Higgins et al., 2003) in the meta-analysis literature and it is frequently used to quantify the relative heterogeneity around the true effect size (Borenstein et al., 2017). An advantage of placing a prior on I^2 is that it is a bounded parameter which enables us to place a proper (noninformative) uniform prior to compute Bayes factors (note that Bayes factors generally cannot be computed using improper priors while at the same time arbitrarily vague proper priors should also be avoided due to Bartlett's phenomenon (Jeffreys, 1961; Lindley, 1957; Bartlett, 1957)). Due to the uniformity of the prior, the proposed Bayes factors can be used for a default Bayesian meta-analysis without requiring external prior knowledge about the model parameters.

The proposed Bayes factors enable testing point and one-sided hypotheses. Examples of a point and one-sided hypothesis are testing whether the overall effect size in a meta-analysis equals zero (i.e., $H : \mu = 0$) or is larger than zero (i.e., $H : \mu > 0$). Moreover, the proposed Bayes factors also enable testing multiple hypotheses simultaneously (e.g., $H : \mu = 0$ vs. $H : \mu > 0$ vs. $H : \mu < 0$). Another attractive property is that the quantification of relative evidence between hypotheses is exact even in the extreme case of only two studies in a meta-analysis as Bayes factors do not rely on large sample theory.

The outline of this paper is as follows. We continue by further introducing the MAREMA model and illustrate Bayesian estimation under the MAREMA model. Subsequently, we introduce Bayes factor testing under the MAREMA model and elaborate on the specification of the prior distributions. In Section “[Bayesian hypothesis testing in examples](#)”, the MAREMA model is used to compute Bayes factors for a meta-analysis on the efficacy of two treatments for post-traumatic stress disorder (PTSD) and a meta-analysis on the effect of a smartphone application and

cessation support on long-term smoking abstinence. We conclude the paper with a discussion section.

Marginalized random-effects meta-analysis model

The conventional random-effects meta-analysis model (Borenstein et al., 2009; Konstantopoulos & Hedges, 2019) assumes that $i = 1, 2, \dots, k$ independent effect sizes are extracted from studies on the same topic. The statistical model of the random-effects model can be written as

$$\begin{aligned} y_i &\sim N(\theta_i, \sigma_i^2) \\ \theta_i &\sim N(\mu, \tau^2) \end{aligned} \tag{1}$$

where y_i is the observed effect size of the i th study, θ_i is the study-specific (unknown) true effect size, μ is the overall true effect size, τ^2 is the between-study variance in true effect size, and σ_i^2 is the known sampling variance of the i th study, which is conventionally estimated in practice and then assumed to be known in the analysis. The random-effects model simplifies to an equal-effect model if $\tau^2 = 0$ because all studies then share a common true effect size. Note that other distributions for the random effects than the normal distribution have been proposed (e.g., Baker & Jackson, 2008; 2016; Lee & Thompson, 2008), but a normal distribution for the random effects remains to be used in almost any random-effects meta-analysis.

The study-specific true effects θ_i are generally treated as nuisance parameters in the random-effects model. We integrate these out of the random-effects model in Eq. 1 to obtain the MAREMA model, which is given by (see also Raudenbush & Bryk, 1985 for instance)

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

Multiple estimators have been proposed for estimating the between-study variance τ^2 (Veroniki et al., 2016; Langan et al., 2016). Estimates of τ^2 cannot be compared across meta-analyses if these meta-analyses used different effect size measures. That is, a τ^2 estimate in a meta-analysis of standardized mean differences cannot be compared to one of correlation coefficients, and this was one of the reasons to develop I^2 that will be described next (Higgins & Thompson, 2002).

Quantifying heterogeneity using I^2

A commonly used way to quantify the relative heterogeneity in a meta-analysis is using I^2 (Higgins & Thompson, 2002; Higgins et al., 2003),

$$I^2 = \frac{\tau^2}{\tau^2 + \tilde{\sigma}^2} \tag{3}$$

where $\tilde{\sigma}^2$ is the typical within-study sampling variance that is computed with

$$\tilde{\sigma}^2 = \frac{(k - 1) \sum 1/\sigma_i^2}{(\sum 1/\sigma_i^2)^2 - \sum 1/\sigma_i^4}. \tag{4}$$

I^2 has an intuitive interpretation because it is the proportion of total variance that can be attributed to between-study variance in true effect size. Note that I^2 resembles the intraclass correlation coefficient (ICC) that is routinely reported in multilevel analysis. This ICC indicates the proportion of total variance that can be attributed to taking into account the dependence of observations within the level 2 units (e.g., (Hox et al., 2018)). However, a major difference between I^2 and the ICC is that the total variance in the computation of I^2 (i.e., $\tau^2 + \tilde{\sigma}^2$) is a function of the studies' sample size whereas the ICC is not a function of the sample size of the level 2 units. Hence, I^2 artificially increases if the sample size of the primary studies increases while τ^2 remains constant (Rücker et al., 2008).

Next, we reparameterize the MAREMA model in Eq. 2 using I^2 . We replace I^2 with the Greek letter ρ to make explicit that it is an unknown parameter and can attain negative values. The MAREMA model can then equivalently be written as

$$y_i \sim N(\mu, \sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)). \tag{5}$$

where $\tilde{\sigma}^2 \rho / (1 - \rho)$ transforms ρ to its corresponding τ^2 . In Eq. 5, ρ must lie in the interval $(\rho_{min}, 1)$ with

$$\rho_{min} = \frac{-\sigma_{min}^2}{-\sigma_{min}^2 + \tilde{\sigma}^2} \tag{6}$$

where σ_{min}^2 is the smallest sampling variance of the studies included in the meta-analysis. ρ_{min} is the smallest possible value of the parameter ρ given the observed data and is always negative. Note that the special case $\rho = 0$ (i.e., the equal-effect model) does not lie on the boundary of the parameter space, which enables testing the hypothesis of homogeneous true effect size.

Bayesian estimation under the MAREMA model

The MAREMA model can be estimated using flat priors if prior information is absent,

$$\begin{aligned} \pi(\mu, \rho) &= \pi(\mu)\pi(\rho), \text{ with} \\ \pi(\mu) &\propto 1 \\ \pi(\rho) &= U(\rho_{min}, 1) \end{aligned} \tag{7}$$

where U refers to the uniform distribution. Flat priors are used for estimation to minimize the impact of the priors on the results. We illustrate Bayesian estimation under the

MAREMA model by analyzing data of two meta-analyses. The first meta-analysis is by Ho & Lee, (2012) on the efficacy of eye movement desensitization and reprocessing (EMDR) therapy versus exposure based cognitive behavior therapy (CBT) to treat PTSD. This meta-analysis consists of ten standardized mean differences (i.e., Hedges' g) and a positive effect size indicates that EMDR therapy is more efficacious than CBT. The second meta-analysis is by Whittaker et al., (2019) on the difference between using a smartphone app for smoking cessation support and lower intensity support on long-term abstinence. Three studies are included in this meta-analysis and log risk ratio is the effect size measure of interest. A positive log risk ratio indicates that using a smartphone app for smoking cessation support yields more long-term abstinence than lower intensity support. These examples were selected to illustrate that the proposed methodology can be used for different effect size measures and meta-analyses with only a small number of primary studies, which is especially common in medical research (Rhodes et al., 2015; Turner et al., 2015).

A Gibbs sampler is proposed for Bayesian estimation under the MAREMA model (see Appendix A). We also analyzed the data of these meta-analyses with frequentist random-effects meta-analysis using the restricted maximum likelihood estimator (Raudenbush, 2009) for estimating τ^2 as implemented in the R package `metafor` (Viechtbauer, 2010). R code of these analyses is available at <https://osf.io/jcge7/>.

The posterior distributions of μ and ρ when fitting the MAREMA model to the meta-analyses by Ho & Lee, (2012) (solid lines) and Whittaker et al., (2019) (dashed lines) are presented in Fig. 1. Remarkably, the posterior distributions

of ρ (right panel of Fig. 1) are very wide for both meta-analyses. There is also considerable posterior support for negative ρ values. This could suggest that the random-effects model, which is employed for the frequentist analysis, may not be appropriate for these data. We will reflect on causes and the implications of a negative value for ρ in the discussion section.

Parameter estimates obtained with the MAREMA model and also the frequentist random-effects model are presented in Table 1. The first row of Table 1 shows the results of the MAREMA model and the second row those of the frequentist random-effects model. The first three columns show the results of estimating μ and the last three columns those of estimating ρ . The `metafor` package was used for the frequentist meta-analysis, which does not report a standard error for $\hat{\rho}$.

Parameter estimates of the MAREMA model and frequentist meta-analysis of the meta-analysis by Ho & Lee, (2012) were comparable. However, the estimates for μ were slightly larger under the MAREMA model relative to the frequentist meta-analysis estimate. Furthermore, as expected, the 95% credibility interval for ρ under the MAREMA model was considerably wider than the 95% confidence interval of the frequentist meta-analysis due to the fact that the random-effects model does not allow negative values for ρ and therefore there is less “room” for ρ to vary. To conclude, EMDR therapy was more efficacious than CBT therapy for treating PTSD, and heterogeneity was imprecisely estimated close to zero (indicating homogeneity).

Parameter estimates for μ under the MAREMA model were approximately zero whereas the frequentist meta-analytic estimate for μ was slightly larger for the

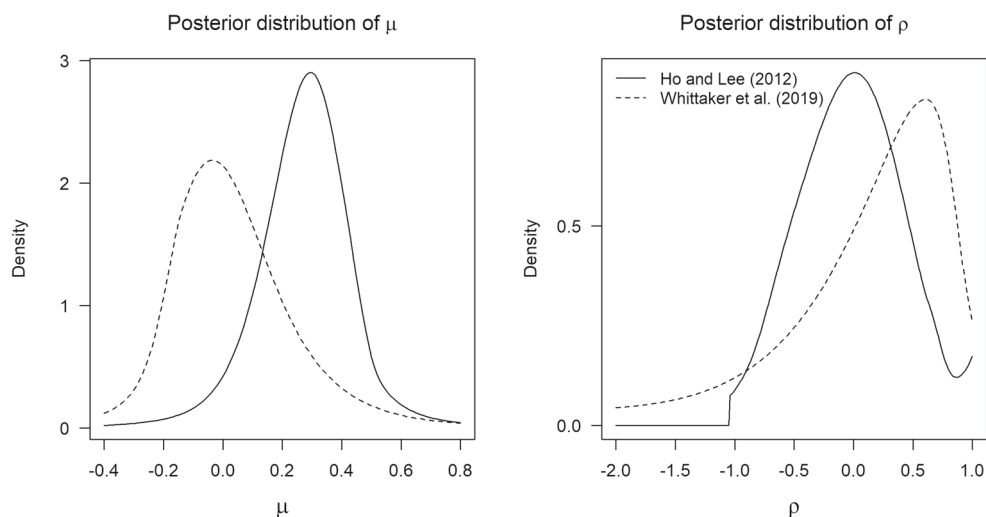


Fig. 1 Posterior distributions of μ (left panel) and ρ (right panel) for the meta-analyses by Ho & Lee, (2012) (solid lines) and Whittaker et al., (2019) (dashed lines). The posterior distributions in the figure are smoothed using a logspline as implemented in the R package `logspline` (Kooperberg, 2020). $\rho_{min} = -1.045$ for the meta-analysis by Ho & Lee, (2012) and $\rho_{min} = -2.326$ for the meta-analysis by Whittaker et al., (2019)

Table 1 Results of Bayesian estimation under the marginalized random-effects meta-analysis (MAREMA) model (using posterior means) and frequentist random-effects meta-analysis when estimating the parameters in the meta-analysis by Ho & Lee, (2012) (*first panel*) and Whittaker et al., (2019) (*second panel*)

	$\hat{\mu}$			$\hat{\rho}$		
	Estimate	SD/SE	95% CI	Estimate	SD/SE	95% CI
Ho & Lee, (2012)						
MAREMA	0.274 (0.327)	0.29	(-0.109;0.638)	-0.026 (-0.016)	0.425	(-0.837;0.812)
Frequentist	0.249	0.129	(-0.003;0.502)	0.022	–	(0;0.747)
Whittaker et al., (2019)						
MAREMA	0.033 (0.043)	0.381	(-0.413;0.625)	0.089 (0.597)	0.68	(-1.752;0.922)
Frequentist	0.114	0.326	(-0.525;0.753)	0.696	–	(0;0.993)

Both the posterior mean and mode (in brackets) based on the draws from the posterior distribution are presented as parameter estimates of the MAREMA model; SD refers to the standard deviation of posterior draws; SE refers to the standard error; and 95% CI refers to the credibility interval in case of Bayesian estimation under the MAREMA model and confidence interval in case of the frequentist meta-analysis; no standard error is reported in the output of the `metafor` package for $\hat{\rho}$; the 95% CI of $\hat{\rho}$ in the frequentist meta-analysis was computed using the Q -profile method (Viechtbauer, 2007)

meta-analysis by Whittaker et al., (2019). Furthermore, due to the skewness of the posterior of ρ under the MAREMA model there is a considerable difference between the posterior mean and posterior mode, where the latter is close to the estimate under the frequentist random-effects model. To conclude, estimates based on the MAREMA model and frequentist meta-analysis differed slightly for the meta-analysis of Whittaker et al., (2019). These difference were probably caused by the meta-analysis only containing three studies.

Given the uncertainty in the unconstrained estimates it is particularly useful to test precise null hypotheses on the overall effect μ and the relative heterogeneity ρ . Hence, Bayesian hypothesis testing under the MAREMA model is discussed next.

Bayesian hypothesis testing under the MAREMA model

Testing hypotheses plays a fundamental role in scientific research in general and psychological science in particular. In this section, we propose multiple hypothesis tests for the mean μ and the relative between-study heterogeneity ρ separately.

We propose testing hypotheses for both μ and ρ under the MAREMA model. The following hypotheses are being tested for μ :

$$\begin{aligned}
 H_0 : \mu &= 0 \\
 H_1 : \mu &< 0 \\
 H_2 : \mu &> 0,
 \end{aligned}
 \tag{8}$$

where support for H_0 , H_1 , or H_2 indicates that the overall effect μ is equal to zero, is negative, or is positive, respectively. For ρ we test the following hypotheses:

$$\begin{aligned}
 H_0 : \rho &= 0 \\
 H_1 : \rho &< 0 \\
 H_2 : \rho &> 0,
 \end{aligned}
 \tag{9}$$

where support for H_0 , H_2 , or H_1 indicates a good fit of an equal-effect model, a random-effects model, or a model which assumes less variance due to sampling error (and thus a misfit of the equal-effect or random-effects model), respectively.

Bayes factors are used for testing the proposed hypotheses (Jeffreys, 1961; Kass & Raftery, 1995). A Bayes factor quantifies the evidence in the data for one hypothesis relative to a contrasting hypothesis via the ratio of the marginal likelihood,

$$B_{12} = \frac{m_1(\mathbf{y})}{m_2(\mathbf{y})}
 \tag{10}$$

where \mathbf{y} is a vector of the observed effect size y_i and m_1 and m_2 are the marginal likelihoods under H_1 and H_2 , respectively. For example, $B_{12} = 1$ indicates that both hypotheses are equally supported by the data whereas $B_{12} = 20$ indicates that there is 20 times more support in the data for H_1 relative to H_2 .

Prior specification

In Bayesian hypothesis testing for the overall effect size, it is not recommended to use an arbitrarily vague prior to avoid Bartlett’s paradox (Jeffreys, 1961; Lindley, 1957;

Bartlett, 1957). Instead, following (Mulder & Fox, 2019), we propose a unit-information prior for μ conditional on ρ in combination with a proper uniform prior for ρ . Under the unconstrained MAREMA model this boils down to

$$\begin{aligned}\pi_u(\mu, \rho) &= \pi_u(\mu|\rho)\pi_u(\rho), \text{ with} \\ \pi_u(\mu|\rho) &= N\left(\mu, k\left(\mathbf{1}'\Sigma_\rho^{-1}\mathbf{1}\right)^{-1}\right) \\ \pi(\rho) &= U(\rho_{min}, 1).\end{aligned}\quad (11)$$

The unit-information prior $\pi(\mu|\rho)$ contains the amount of information of a single study (Zellner, 1986). This is visible in the variance of $\pi(\mu|\rho)$ because the number of studies in a meta-analysis k is multiplied by the variance of $\hat{\mu}$, which is $(\mathbf{1}'\Sigma_\rho^{-1}\mathbf{1})^{-1}$ where $\mathbf{1}$ is a column vector of ones and $\Sigma_\rho = \text{diag}(\sigma_1^2 + \tau^2, \dots, \sigma_k^2 + \tau^2)$. Unit-information priors are commonly used for computing Bayes factors in model selection and hypothesis testing problems. For example, the well-known Bayesian information criterion (BIC, (Schwarz, 1978)) is based on an approximation of the marginal likelihood using a unit-information prior (Raftery, 1995; Kass & Wasserman, 1995). This class of priors is also employed in many other Bayesian testing scenarios (e.g., (Liang et al., 2008; Rouder & D Morey, 2012; Mulder et al., in press)). The usefulness of unit-information priors lies in the fact that they cover a reasonable range of possible values for the model parameters. As the prior contains the information of a single study, these priors are neither too informative nor too vague (to avoid Bartlett's paradox). The prior $\pi(\mu|\rho)$ depends on ρ because τ^2 is included in the variance-covariance matrix Σ_ρ . Note that the prior $\pi(\mu|\rho)$ is different from the prior used for Bayesian estimation under the MAREMA model in Section "Bayesian estimation under the MAREMA model" as we used an improper prior for estimation whereas the prior $\pi(\mu|\rho)$ used for testing is a proper prior. The proper prior $\pi(\rho)$ is the same prior distribution as we proposed for estimation under the MAREMA model.

The unconstrained prior in Eq. 11 is used as building block for the priors under the hypotheses of interest. In particular, for the one-sided hypotheses, truncated priors are considered while the precise null hypotheses receive a point mass at the null value. Figure 2 illustrates the proposed prior distributions for testing unconstrained, one-sided, and point hypotheses under the MAREMA model. The left panel shows prior distributions for testing hypotheses regarding μ where ρ is left unconstrained. The dashed line refers to the prior distribution of the unconstrained hypothesis in Eq. 11. The asterisk at the top of the figure illustrates the point mass for testing the hypothesis $H_0 : \mu = 0$. The solid and dotted lines refer to the prior distributions for the one-sided hypotheses $H_1 : \mu < 0$ and $H_2 : \mu > 0$, respectively. Their heights are twice as large as the unconstrained prior to ensure that the distributions integrate to 1. The right panel

shows prior distributions for testing hypotheses regarding ρ . The dashed, solid, and dotted lines refer to the prior distributions of the unconstrained (H_u), left-sided ($H_1 : \rho < 0$), and right-sided ($H_2 : \rho > 0$) hypotheses whereas the asterisk refers to the point mass for the hypothesis of no between-study variance ($H_0 : \rho = 0$).

Marginal likelihood

The marginal likelihoods of the different hypotheses differ with respect to the prior distributions. Hence, the marginal likelihoods of the different hypotheses can be computed by using different prior distributions in combination with adjusting the limits of integration.

For example, the marginal likelihood for the one-sided hypothesis $H_1 : \mu < 0$ with ρ unconstrained can be written as a function of the marginal likelihood under the unconstrained model,

$$m_1(\mathbf{y}) = \int_{\mu < 0} f(\mathbf{y}|\mu, \rho)\pi_1(\mu, \rho)d\mu d\rho = m_u(\mathbf{y}) \frac{P(\mu < 0|\mathbf{y}, H_u)}{P(\mu < 0|H_u)}, \quad (12)$$

where $f(\mathbf{y}|\mu, \rho)$ is the likelihood function of the MAREMA model in Eq. 5, $m_u(\mathbf{y})$ is the marginal likelihood under the unconstrained model, $P(\mu < 0|\mathbf{y}, H_u)$ and $P(\mu < 0|H_u)$ are the posterior and prior model probabilities for $\mu < 0$ under the hypothesis H_u , and the prior under H_1 can be written as a truncation of the unconstrained prior

$$\pi_1(\mu, \rho) = \pi_u(\mu, \rho)I(\mu < 0)/P(\mu < 0|H_u),$$

where $I(\cdot)$ is the indicator function. Note here that $P(\mu < 0|H_u) = \frac{1}{2}$ because the unconstrained prior for μ is centered around 0. The posterior probability in Eq. 12 can be computed as the proportion of unconstrained posterior draws satisfying $\mu < 0$. Interested readers are referred to Appendix B for further computational details.

Software: BFpack

The R package `BFpack` (Mulder et al., in press) contains functions for computing Bayes factors for a large set of statistical models (e.g., multivariate regression, generalized linear models, and correlation analysis). As an argument the main function "BF" needs a fitted modeling object which defines under which model Bayes factors need to be computed. To execute the Bayes factor test under the MAREMA model using the "BF" function, the function needs as argument an object returned by fitting a random-effects meta-analysis model with the `metafor` package. The `metafor` package is popular software for conducting meta-analysis that can be used for any effect size measure

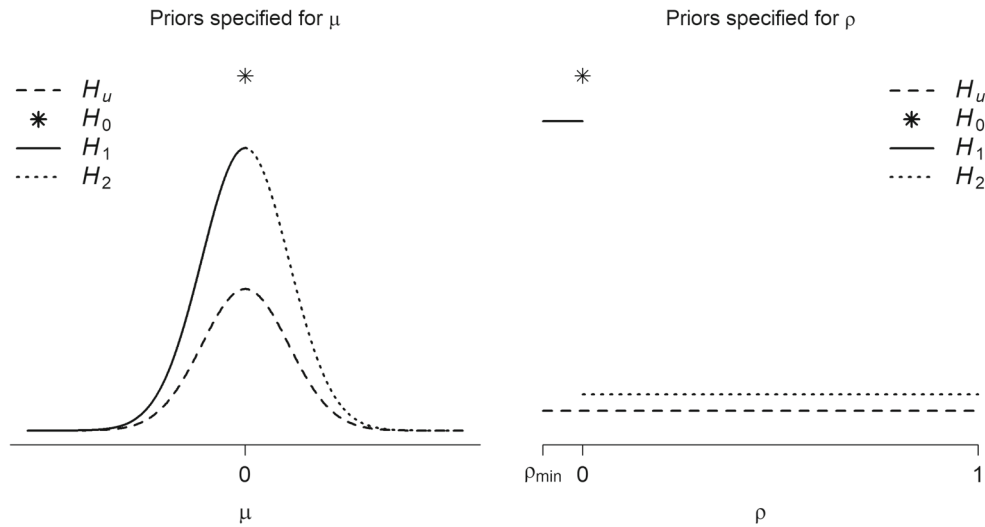


Fig. 2 Prior distributions of μ (left panel) and ρ (right panel) for Bayesian hypothesis testing under the MAREMA model. The different prior distributions refer to the hypotheses listed in Eqs. 8 and 9

and requires the meta-analyst to supply the observed effect sizes of the primary studies and the corresponding sampling variances (or standard errors). Hence, researchers familiar with the `metafor` package can readily compute the Bayes factors that we propose using the function `BF`. The “BF” function also returns unconstrained estimates of μ and ρ based on a Gibbs sampler.

Bayesian hypothesis testing in examples

We compute Bayes factors using the MAREMA model for the two examples introduced in Section “Bayesian estimation under the MAREMA model”. The MAREMA model was fitted to the two meta-analyses using the proposed unit-information prior on μ and uniform prior on ρ .¹ We tested the three hypotheses for μ and ρ listed in Eqs. 8 and 9, respectively. We analyzed the data using R (R Core Team, 2020) and the R packages `metafor` (Viechtbauer, 2010) and `BFpack` (Mulder et al., in press) in particular. R code illustrating how to compute Bayes factors and posterior probabilities for the hypotheses for the two examples is available at <https://osf.io/ejfsv/>.

The Bayes factors and posterior model probabilities for hypotheses on μ are presented in Table 2. For the meta-analysis by Ho & Lee, (2012), the Bayes factor comparing H_2 with H_1 is the largest of the tested hypotheses, which

implies that the hypothesis $H_2 : \mu > 0$ is 15.810 times more likely than the hypothesis $H_1 : \mu < 0$. Moreover, the Bayes factor comparing H_2 with $H_0 : \mu = 0$ equaled 3.779 implying that μ in this meta-analysis is likely larger than zero. Also the posterior probabilities (last row in Table 2; assuming equal prior probabilities) suggested that a positive effect (H_2) was most likely after observing the data with a posterior probability of 0.753. The two-tailed frequentist hypothesis test of $H_0 : \mu = 0$ was not statistically significant using a significance level of 0.05 ($z = 1.936$, two-tailed p value = 0.053).

The Bayes factors and posterior model probabilities for hypotheses on ρ are shown in the first columns of Table 3 for the study by Ho & Lee, (2012). The Bayes factors comparing the hypotheses $H_0 : \rho = 0$ with $H_1 : \rho < 0$ and $H_2 : \rho > 0$ indicated that H_0 is approximately four and five times more likely, respectively. These results were also corroborated by the posterior probability that was the largest for H_0 (with 0.689). This indicates that there was most evidence for an equal-effect model. The commonly used Q -test for testing whether the studies are homogeneous in a frequentist meta-analysis was not statistically significant ($Q(9) = 9.417$, $p = 0.4$). Note here that a nonsignificant result does not imply evidence for the null, as p values cannot be used for quantifying evidence for the null (because p values are by definition uniformly distributed if H_0 would be true). To conclude, the EMDR treatment was observed to be on average more efficacious than the CBT treatment and effects are homogeneous across studies. However, due to the uncertainty in the posterior probabilities and the parameter estimates (see Section “Bayesian estimation under the MAREMA model”), more studies are required in order to draw more definite conclusions.

¹The random walk procedures were based on 100,000 iterations (burn-in period of 5000 iterations) and starting values ρ_0 equal to the estimate of a frequentist random-effects meta-analysis with the restricted maximum likelihood as estimator for the between-study variance and $s^2 = \sqrt{0.1}$

Table 2 Bayes factors and posterior model probabilities ($P(H_q|\mathbf{y})$) for hypotheses on μ . The results based on the meta-analysis by Ho & Lee, (2012) are shown in the first columns and by Whittaker et al., (2019) in the last columns of the table

		Ho & Lee, (2012)			Whittaker et al., (2019)		
		H_0	H_1	H_2	H_0	H_1	H_2
Bayes factors	H_0	1.000	4.183	0.265	1.000	2.558	2.115
	H_1	0.239	1.000	0.063	0.391	1.000	0.827
	H_2	3.779	15.810	1.000	0.473	1.209	1.000
$P(H_q \mathbf{y})$		0.199	0.048	0.753	0.537	0.210	0.254

$H_0 : \mu = 0$, $H_1 : \mu < 0$, and $H_2 : \mu > 0$

Bayes factors for testing hypotheses on μ for the meta-analysis by Whittaker et al., (2019) are shown in the last columns of Table 2. Hypothesis $H_0 : \mu = 0$ received more support than hypotheses $H_1 : \mu < 0$ and $H_2 : \mu > 0$, but there was no strong evidence for any of the hypotheses (largest Bayes factor equaled 2.558). This absence of strong evidence was likely also caused by this meta-analysis only consisting of three studies. The posterior model probabilities (last row Table 2) also showed that H_0 was most likely (probability is 0.537). Application of a two-tailed frequentist hypothesis test resulted in a nonsignificant result, and thus H_0 could not be rejected ($z = 0.349$, two-tailed p value = 0.727).

The hypothesis $H_0 : \rho = 0$ received 10.958 more evidence than $H_1 : \rho < 0$ and 2.901 more evidence than $H_2 : \rho > 0$. Moreover, the hypothesis $H_2 : \rho > 0$ was 3.778 more likely than $H_1 : \rho < 0$. This corroborated the posterior model probabilities indicating that the effect sizes were most likely to be either homogeneous (with a probability 0.696) or heterogeneous (with a probability 0.240). Interestingly, the frequentist Q -test was statistically significant ($Q(2) = 6.24$, $p = 0.044$). This may be due to the fact that the significance tests rely on large sample theory, which may not be realistic here given there are only three studies. Using the posterior probabilities to get conditional error probabilities, there is a probability of $0.210 + 0.254 = 0.464$ that we would be wrong when concluding that the hypothesis of no effect is true, and a

probability of $0.064 + 0.240 = 0.304$ that the hypothesis of homogeneous effects is true given the observed data. Thus (as expected) more data would be needed in order to receive more pronounced evidence about which hypothesis is likely to be true.

Discussion

The main goals of meta-analyses are estimating the overall effect and the heterogeneity in effect size as well as drawing inferences for these parameters. This paper proposes novel Bayesian estimation and hypothesis testing methods to achieve these goals. Our approach is novel compared to alternative Bayesian meta-analysis methods because the framework builds on the MAREMA model where the (nuisance) study-specific effects are integrated out. This MAREMA model encompasses both an equal-effect and a random-effects meta-analytic model, and also encompasses a model which assumes that there is less variance than under the equal-effect and random-effects models.

Another major contribution is that we place a uniform prior distribution on the proportion of the variance that can be explained by between-study heterogeneity in true effect size rather than on the between-study variance or standard deviation in true effect size directly, as is usually done in Bayesian meta-analyses (Berry, 1998; Scheibehenne et al., 2017; Gronau et al., 2017). This relative variance has a clear

Table 3 Bayes factors and posterior model probabilities ($P(H_q|\mathbf{y})$) for hypotheses on ρ . The results based on the meta-analysis by Ho & Lee, (2012) are shown in the first columns and by Whittaker et al., (2019) in the last columns of the table

		Ho & Lee, (2012)			Whittaker et al., (2019)		
		H_0	H_1	H_2	H_0	H_1	H_2
Bayes factors	H_0	1.000	3.977	4.979	1.000	10.958	2.901
	H_1	0.251	1.000	1.252	0.091	1.000	0.265
	H_2	0.201	0.799	1.000	0.345	3.778	1.000
$P(H_q \mathbf{y})$		0.689	0.173	0.138	0.696	0.064	0.240

$H_0 : \rho = 0$, $H_1 : \rho < 0$, and $H_2 : \rho > 0$

standardized scale (Higgins et al., 2003) which facilitates its interpretation. Furthermore, the bounded parameter allows specifying a proper noninformative uniform prior in case prior information is absent or when a default (reference) test is preferred. Other advantages are that this prior does not depend on the effect size measure used in the meta-analysis, and avoids the need of eliciting a prior scale for which the Bayes factor can be highly sensitive.

We illustrated the proposed Bayesian hypothesis testing and estimation in two illustrative examples. Both the estimation and hypothesis testing results revealed large posterior uncertainty regarding the heterogeneity in true effect size. The uncertainty can be explained by the relatively small number of studies which is common in meta-analyses (Rhodes et al., 2015; Turner et al., 2015), and therefore the obtained quantifications for posterior uncertainty seemed reasonable. More convincing evidence for one of the tested hypotheses will be obtained if more studies become available and can be included in the meta-analysis. Note here that because Bayes factors are consistent the evidence in favor of the true hypothesis would go to infinity if the number of studies in a meta-analysis tends to infinity. The frequentist test for between-study heterogeneity on the other hand was statistically significant even though the example only included three studies. This may be another example that p values tend to overestimate the evidence against a null hypothesis (e.g., (Berger & Delampady, 1987; Sellke et al., 2001; Benjamin & Berger, 2019), and that Bayes factors and posterior probabilities may better capture the evidence in favor of or against statistical hypotheses in case of small samples.

In both applications, evidence was found that the heterogeneity (ρ) may be negative. This implies that an equal-effect meta-analysis model (where $\rho = 0$) or a random effects model (where $\rho > 0$) may not be appropriate, and might cause bias in the results. A negative ρ may also indicate a violation of the assumptions of the meta-analysis model, which is “corrected” by allowing ρ to be negative. For example, the reported within-study sampling variances σ_i^2 may be overestimated, the assumption of independent primary study’s effect sizes in a meta-analysis is violated, or the effect sizes of the primary studies were incorrectly computed (Ioannidis et al., 2006). More research is needed to explore the possible causes of a negative ρ , and how it affects the results. A good starting point would be (Nielsen et al., 2021) who explored the effect of negative intraclass correlations in a multilevel model.

We have only tested point hypotheses where the parameter was constrained to zero or one-sided hypotheses comparing hypotheses where the parameter was smaller or larger than zero. Other hypotheses that can be tested are combined hypotheses (e.g., $H : \mu > 0$ & $\rho > 0$) or hypotheses testing a parameter to another hypothesized value than zero. For example, meta-analysts may want to test

point hypotheses using the rules of thumb for a small, medium, and large effect size as defined for many common effect size measures (for rules of thumb for multiple effect size measures see (Cohen, 1988)). Point hypotheses on the heterogeneity may be constrained to $\rho = 0.25, 0.5,$ and 0.75 to resemble low, moderate, and high heterogeneity according to the thresholds proposed by (Higgins et al., 2003). It is important to note that any hypothesis on the proportion of total variance that can be explained by heterogeneity in true effect size can directly be transformed to a hypothesis on the between-study variance in true effect size. For example, testing the hypothesis $H : \rho = 0$ is equivalent to testing $H : \tau^2 = 0$. If a hypothesis on ρ is tested against another value than 0, the equivalent hypothesis on the between-study variance in true effect size can be obtained by transforming ρ to the between-study variance in true effect size.

We proposed a framework for Bayesian hypothesis testing and estimation using minimally informative default prior distributions, but our methodology can readily be extended to other prior distributions. Informative priors can be specified to incorporate external information about the heterogeneity in true effect size or overall effect from research in comparable fields. This may provide better estimates and statistical inferences especially if the number of primary studies in the meta-analysis is small. For example, if prior information about the heterogeneity parameter ρ is available this could be translated to an informative stretched beta prior distribution (following, (Mulder & Fox, 2019) for example) while prior information about the effect size could be translated to an informative normal prior. By considering different prior distributions one can assess the robustness of the quantification of the relative evidence in the data between statistical hypotheses under the MAREMA model.

Meta-analysts are generally not only interested in estimating and drawing inferences about the overall effect size and the between-study variance, but they are also interested in studying whether systematic differences between primary studies can explain the between-study variance. In case of large between-study variance, examining whether systematic differences between studies exist might even be more insightful than focusing on the overall effect size (Rouder et al., 2019). The between-study variance in a meta-analysis can be explained by including moderator variables in a meta-analysis model in a so-called meta-regression (Thompson & Sharp, 1999; Van Houwelingen et al., 2002). Future research may focus on extending our MAREMA model such that Bayesian estimation and Bayes factor testing can be conducted when moderators are included.

Future research may also focus on extending our methods to more advanced meta-analysis models such as multilevel meta-analysis (van den Noortgate & Onghena, 2003;

Konstantopoulos, 2011) and multivariate meta-analysis (Jackson et al., 2011; Hedges, 2019). These models relax the strong assumption of the conventional random-effects model that effect sizes in a meta-analysis have to be independent. Another avenue for future research is studying whether relaxing the assumptions of the random-effects model benefit Bayesian meta-analysis under the MAREMA model in particular and Bayesian meta-analysis in general. For example, the random-effects model assumes that the within-study sampling variances are known (van Aert & Jackson, 2019; Jackson, 2009; Konstantopoulos & Hedges, 2019), and we adopted this assumption in the MAREMA model. This assumption is not tenable in practice which can be problematic if the sample size of studies is small. Estimates of the within-study sampling variance are then imprecise, which potentially lead to biased parameter estimates and inaccurate statistical inferences. This strong assumption can be avoided in a Bayesian meta-analysis by taking the uncertainty in these variances into account by means of a prior distribution instead of using a plug-in estimate. A logical choice for a prior distribution on the within-study sampling variances is the inverse-gamma distribution.

Another topic for future research is studying to what extent the proposed estimation and Bayes factor test under the MAREMA model are affected by publication bias. Publication bias implies that especially studies with statistically significant findings are more likely to be published than studies with non-significant findings. Consequently, studies with non-significant findings are more difficult to locate and are less likely to be included in a meta-analysis. Due to publication bias, effect sizes of primary studies and, in turn, also the overall effect size of a meta-analysis are most likely positively biased (Kraemer et al., 1998; van Assen et al., 2015; Lane & Dunlap, 1978). We expect estimation and inferences based on the MAREMA model to be inaccurate if publication bias is severe, and recommend researchers to also apply and report methods that correct for publication bias in this case.

To conclude, we have proposed Bayesian hypothesis testing and estimation using the MAREMA model. The proposed Bayes factors allow testing point and one-sided hypotheses for both the overall effect and heterogeneity in true effect size. We hope that our methods together with the easy-to-use software included in the R package `BFPack` enables researchers to frequently use Bayesian methods in their meta-analyses.

Appendix A: Gibbs sampler for Bayesian estimation under the MAREMA model

We show in this appendix how the MAREMA model can be estimated by means of a Gibbs sampler. The likelihood

of the MAREMA model given the prior distributions in Eq. 7 is

$$h(\mathbf{y}|\mu, \rho) = (2\pi)^{-\frac{k}{2}} \prod \left[\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho) \right]^{-\frac{1}{2}} \times \exp \left[-\frac{1}{2} \sum \frac{(y_i - \mu)^2}{\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)} \right] (1 - \rho_{min})^{-1}. \quad (13)$$

We first derive the conditional distribution of μ , i.e.,

$$\mu|\mathbf{y}, \rho \propto \exp \left[-\frac{1}{2} \sum \frac{(y_i - \mu)^2}{\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)} \right] \times \exp \left[-\frac{1}{2} \left(\mu^2 \sum \frac{1}{\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)} - \mu \sum \frac{2y_i}{\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)} \right) \right]. \quad (14)$$

If we let $c = \sum 1 / (\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho))$ and $v = \sum 2y_i / (\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho))$ and we complete the square for μ , we get

$$\mu|\mathbf{y}, \rho \sim N \left(\frac{\sum y_i / (\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho))}{\sum 1 / (\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho))}, \frac{1}{\sum 1 / (\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho))} \right). \quad (15)$$

Using the conditional distribution of μ and a random walk procedure to approximate the conditional distribution of ρ , the Gibbs sampler can be written as

1. Set the initial value ρ_0 and the variance of the proposal distribution s^2 .
2. Repeat the following steps to obtain $j = 1, 2, \dots, J$ draws and verify after every 100 draws whether the acceptance rate a in the random walk procedure is between 0.15 and 0.5 and adjust s^2 to $\left[\sqrt{s^2} \left(\frac{a-0.5}{1-0.15} + 1 \right) \right]^2$ if $a > 0.5$ and to $\left[\frac{\sqrt{s^2}}{2-a/0.15} \right]^2$ if $a < 0.15$.
 - (a) Draw a candidate sample ρ^* from the proposal distribution $q(\rho^*|\rho_{j-1}, s^2)$, which is a truncated normal distribution with mean ρ_{j-1} , variance s^2 , lower bound of truncation $-\sigma_{min}^2$, and upper bound of truncation 1.
 - (b) Draw a μ_j from the conditional distribution of μ in Eq. 15 with $\rho = \rho^*$.
 - (c) Compute the ratio

$$R = \frac{h(\mathbf{y}|\mu_j, \rho^*)/q(\rho^*|\rho_{j-1}, s^2)}{h(\mathbf{y}|\mu_j, \rho^*)/q(\rho_{j-1}|\rho^*, s^2)}.$$
 - (d) Draw a number from a uniform distribution ranging from 0 to 1 and denote this by b .
 - (e) If $R < b$, $\rho_j = \rho^*$ otherwise $\rho_j = \rho_{j-1}$.
3. Exclude a burn-in period from the draws.

Appendix B: Computation of marginal likelihoods

In this appendix, we describe how the marginal likelihoods of the hypotheses can be computed. Note that point and one-sided hypotheses are always specified for one parameter while the other parameter not included in hypothesis is left unconstrained.

B.1: Computing the marginal likelihood of hypotheses with μ and ρ unconstrained (i.e., H_u)

The marginal likelihood of the hypothesis where both μ and ρ are unconstrained is computed by first integrating out μ in $f(\mathbf{y}|\mu, \rho)\pi(\mu, \rho)$ (see Appendix C for the derivation),

$$\begin{aligned}
 m_q(\mathbf{y}) &= \int g(\mathbf{y}|\rho)d\rho = \int (2\pi)^{-\frac{k+1}{2}} \prod [\sigma_i^2 + \tilde{\sigma}^2\rho/(1-\rho)]^{-\frac{1}{2}} \\
 &\times \exp\left[-\frac{1}{2}(\mathbf{y} - \hat{\mu}\mathbf{1})'\Sigma_\rho^{-1}(\mathbf{y} - \hat{\mu}\mathbf{1})\right] \left[k(\mathbf{1}'\Sigma_\rho^{-1}\mathbf{1})^{-1}\right]^{-\frac{1}{2}} (1 - \rho_{min})^{-1} \\
 &\times \exp\left[-\frac{1}{2}(\mathbf{1}'\Sigma_\rho^{-1}\mathbf{1})(\hat{\mu}^2 - \frac{\hat{\mu}^2}{1+k^{-1}})\right] (2\pi)^{\frac{1}{2}} (\mathbf{1}'\Sigma_\rho^{-1}\mathbf{1})^{-\frac{1}{2}} (1+k^{-1})^{-\frac{1}{2}} d\rho \quad (16)
 \end{aligned}$$

where $g(\mathbf{y}|\rho)$ is the likelihood where μ is integrated out and q refers to the particular hypothesis that is being tested. This marginal likelihood in Eq. 16 can then be approximated by using importance sampling,

$$m_u(\mathbf{y})d\rho = \int g(\mathbf{y}|\rho) \frac{\pi(\rho)}{w(\rho)} d\rho \approx J^{-1} \sum_{j=1}^J g(\mathbf{y}|\rho_j) \frac{\pi(\rho_j)}{w(\rho_j)} \quad (17)$$

where $w(\rho)$ is a proposal density. This proposal density is a stretched beta distribution and was proposed by (Mulder & Fox, 2019) as a prior distribution for the ICC in hierarchical models. This stretched beta distribution ranges from ρ_{min} till 1. The shape parameters of this stretched beta distribution are computed with $\alpha = (\rho_{min} - \bar{\rho})[(\bar{\rho} - 1)\bar{\rho} - \bar{\rho}\rho_{min} + \rho_{min} + \lambda]/[(\rho_{min} - 1)\lambda]$ and $\beta = \alpha(\bar{\rho} - 1)/(\bar{\rho} - \rho_{min})$ where $\bar{\rho}$ and λ are the mean and variance of draws from the posterior distribution of ρ . Both shape parameters are multiplied by 0.6 to make sure that the proposal density has heavier tails than the posterior distribution. These heavier tails yield a more accurate approximation of the marginal likelihood.

B.2: Computation of the marginal likelihood of $H_0 : \mu = 0$

The marginal likelihood of the point hypothesis $H_0 : \mu = 0$ has to be approximated because ρ could not be integrated

out of $f(\mathbf{y}|\mu, \rho)\pi(\mu, \rho)$. Importance sampling is used for approximating this marginal likelihood,

$$\begin{aligned}
 m_0(\mathbf{y}) &= \int I(\mu = 0) f(\mathbf{y}|\mu, \rho) d\rho \\
 &= \int I(\mu = 0) f(\mathbf{y}|\mu, \rho) \frac{\pi(\rho)}{w(\rho)} d\rho \quad (18) \\
 &\approx J^{-1} \sum_{j=1}^J I(\mu = 0) f(\mathbf{y}|\mu, \rho_j) \frac{\pi(\rho_j)}{w(\rho_j)}
 \end{aligned}$$

The proposal density $w(\rho)$ is the stretched beta distribution that is also used for computing the marginal likelihood of hypothesis H_u (see Appendix B.1).

B.3: Computation of the marginal likelihood of $H_1 : \mu < 0$

The computation of the marginal likelihood for the hypothesis $H_1 : \mu < 0$ is described in Section “Marginal likelihood”. The posterior model probability $P(\mu < 0|\mathbf{y}, H_u)$ is approximated using a random walk procedure that closely resembles the one in the Gibbs sampler in Appendix A and can be described as

1. Set the initial value ρ_0 and the variance of the proposal distribution s^2 .
2. Repeat the following steps to obtain $j = 1, 2, \dots, J$ draws and verify after every 100 draws whether the acceptance rate a is between 0.15 and 0.5 and adjust s^2 to $\left[\sqrt{s^2} \left(\frac{a-0.5}{1-0.15} + 1\right)\right]^2$ if $a > 0.5$ and to $\left[\frac{\sqrt{s^2}}{2-a/0.15}\right]^2$ if $a < 0.15$.
 - (a) Draw a candidate sample ρ^* from the proposal distribution $q(\rho^*|\rho_{j-1}, s^2)$, which is a truncated normal distribution with mean ρ_{j-1} , variance s^2 , lower bound of truncation $-\sigma_{min}^2$, and upper bound of truncation 1.
 - (b) Compute the ratio

$$R = \frac{g(\mathbf{y}|\rho^*)/q(\rho^*|\rho_{j-1}, s^2)}{g(\mathbf{y}|\rho_j)/q(\rho_j|\rho_{j-1}, s^2)}$$
 - (c) Draw a number from a uniform distribution ranging from 0 to 1 and denote this by b .
 - (d) If $R < b$, $\rho_j = \rho^*$ otherwise $\rho_j = \rho_{j-1}$.
3. Draw samples from the conditional posterior of μ (denoted by μ_j) given the sampled ρ_j . The conditional posterior of μ is

$$N\left(\frac{\sum y_i/(\sigma_i^2 + \tilde{\sigma}^2\rho_j/(1-\rho_j)) + \mu_0/\sigma_j^2}{\sum 1/(\sigma_i^2 + \tilde{\sigma}^2\rho_j/(1-\rho_j)) + 1/\sigma_j^2}, \frac{1}{\sum 1/(\sigma_i^2 + \tilde{\sigma}^2\rho_j/(1-\rho_j)) + 1/\sigma_j^2}\right) \quad (19)$$

where $\mu_0 = 0$ and $\sigma_j^2 = \frac{k}{\sum 1/(\sigma_i^2 + \tilde{\sigma}^2\rho_j/(1-\rho_j))}$.

4. Exclude a burn-in period from the draws.
5. Compute how many draws of μ_j are smaller than 0 to approximate $P(\mu < 0|\mathbf{y}, H_u)$.

B.4: Computation of the marginal likelihood of $H_2 : \mu > 0$

Computation of the marginal likelihood of the hypothesis $H_2 : \mu > 0$ is highly similar to that of $H_1 : \mu < 0$. That is, the marginal likelihood can be approximated with

$$m_1(\mathbf{y}) = \iint_{\mu>0} f(\mathbf{y}|\mu, \rho)\pi_2(\mu, \rho)d\mu d\rho = m_u(\mathbf{y}) \frac{P(\mu > 0|\mathbf{y}, H_u)}{P(\mu > 0|H_u)} \tag{20}$$

where $P(\mu > 0|\mathbf{y}, H_u)$ and $P(\mu > 0|H_u)$ are the posterior and prior model probabilities for $\mu > 0$ under the hypothesis H_u where both μ and ρ are unconstrained. The probability $P(\mu > 0|H_u)$ is computed using the prior distribution $\pi(\mu|\rho)$. The posterior model probability $P(\mu > 0|\mathbf{y}, H_u)$ is approximated by means of the random walk procedure described in Appendix B.3 where is computed how many draws are larger instead of smaller than 0 in the final step.

B.5: Computation of the marginal likelihood of $H_0 : \rho = 0$

The marginal likelihood of the hypothesis $H_0 : \rho = 0$ can be computed with the likelihood function $g(\mathbf{y}|\rho)$ in Eq. 16 where μ is integrated out,

$$m_0(\mathbf{y}) = I(\rho = 0)g(\mathbf{y}|\rho). \tag{21}$$

B.6: Computation of the marginal likelihood of $H_1 : \rho < 0$

The marginal likelihood of the hypothesis $H_1 : \rho < 0$ is computed in a similar way as the marginal likelihood for one-sided hypotheses on μ . That is,

$$m_1(\mathbf{y}) = \int_{\rho<0} g(\mathbf{y}|\rho)d\rho = m_0(\mathbf{y}) \frac{P(\rho < 0|\mathbf{y}, H_u)}{P(\rho < 0|H_u)} \tag{22}$$

where $P(\rho < 0|\mathbf{y}, H_u)$ and $P(\rho < 0|H_u)$ are the posterior and prior model probabilities for $\rho < 0$ under the hypothesis H_u where μ and ρ are unconstrained. The probability $P(\rho < 0|H_u)$ is computed using the prior $\pi(\rho)$. The probability $P(\rho < 0|\mathbf{y}, H_u)$ is approximated by means of the random walk procedure described in Appendix B.3 by computing how many draws of ρ_j are smaller than 0.

B.7: Computation of the marginal likelihood of $H_2 : \rho > 0$

Computation of the marginal likelihood of the hypothesis $H_2 : \rho > 0$ is highly equivalent to that of $H_1 : \rho < 0$,

$$m_2(\mathbf{y}) = \int_{\rho>0} g(\mathbf{y}|\rho)d\rho = m_0(\mathbf{y}) \frac{P(\rho > 0|\mathbf{y}, H_u)}{P(\rho > 0|H_u)} \tag{23}$$

with the exception that $P(\rho > 0|\mathbf{y}, H_u)$ and $P(\rho > 0|H_u)$ are now the posterior and prior model probability of $\rho > 0$. The posterior model probability $P(\rho > 0|\mathbf{y}, H_u)$ is approximated by computing how many draws of ρ_j are larger than 0 in the random walk procedure described in Appendix B.3.

Appendix C: Derivation of marginal likelihood of MAREMA model where μ is integrated out

We show in this appendix how the marginal likelihood in Eq. 16 can be obtained by integrating out μ in $f(\mathbf{y}|\mu, \rho)\pi(\mu, \rho)$. The marginal likelihood of the MAREMA model in Eq. 16 equals

$$\begin{aligned} \iint f(\mathbf{y}|\mu, \rho)d\mu d\rho &= \int \int (2\pi)^{-\frac{k+1}{2}} \prod [\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)]^{-\frac{1}{2}} \\ &\times \exp \left[-\frac{1}{2} \sum \frac{(y_i - \mu)^2}{\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)} \right] (2\pi)^{-\frac{1}{2}} \left[\frac{k}{\sum (\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho))} \right]^{-\frac{1}{2}} \\ &\times \exp \left[-\frac{1}{2} \frac{\mu^2}{k / \sum (\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho))} \right] (1 - \rho_{min})^{-1} d\mu d\rho. \end{aligned} \tag{24}$$

Combining powers and replacing some expressions with their equivalents in matrix algebra yields

$$\begin{aligned} \iint f(\mathbf{y}|\mu, \rho)d\mu d\rho &= \int \int (2\pi)^{-\frac{k+1}{2}} \prod [\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)]^{-\frac{1}{2}} \\ &\times \exp \left[-\frac{1}{2} (\mathbf{y} - \hat{\mu}\mathbf{1})' \Sigma_\rho^{-1} (\mathbf{y} - \hat{\mu}\mathbf{1}) \right] \left[k(\mathbf{1}' \Sigma_\rho^{-1} \mathbf{1})^{-1} \right]^{-\frac{1}{2}} (1 - \rho_{min})^{-1} \\ &\times \exp \left[-\frac{1}{2} \left(\frac{(\mu - \hat{\mu})^2}{(\mathbf{1}' \Sigma_\rho^{-1} \mathbf{1})^{-1}} + \frac{k^{-1} \mu^2}{(\mathbf{1}' \Sigma_\rho^{-1} \mathbf{1})^{-1}} \right) \right] d\mu d\rho. \end{aligned} \tag{25}$$

Rewriting Eq. 25 to facilitate integrating out μ gives

$$\begin{aligned} \iint f(\mathbf{y}|\mu, \rho)d\mu d\rho &= \int \int (2\pi)^{-\frac{k+1}{2}} \prod [\sigma_i^2 + \tilde{\sigma}^2 \rho / (1 - \rho)]^{-\frac{1}{2}} \\ &\times \exp \left[-\frac{1}{2} (\mathbf{y} - \hat{\mu}\mathbf{1})' \Sigma_\rho^{-1} (\mathbf{y} - \hat{\mu}\mathbf{1}) \right] \left[k(\mathbf{1}' \Sigma_\rho^{-1} \mathbf{1})^{-1} \right]^{-\frac{1}{2}} (1 - \rho_{min})^{-1} \\ &\times \exp \left[-\frac{1}{2} \frac{(\mu - \hat{\mu} / (1 + k^{-1}))^2}{(\mathbf{1}' \Sigma_\rho^{-1} \mathbf{1})^{-1} (1 + k^{-1})^{-1}} \right] \exp \left[-\frac{1}{2} (\mathbf{1}' \Sigma_\rho^{-1} \mathbf{1}) (\hat{\mu}^2 - \hat{\mu}^2 / (1 + k^{-1})) \right] d\mu d\rho. \end{aligned} \tag{26}$$

Integrating out μ of Eq. 26 yields the marginal likelihood in Eq. 16.

Acknowledgements We would like to thank Jean-Paul Fox for insightful discussions about marginalized random effects models.

Funding Robbie C.M. van Aert is supported by the European Research Council. Grant Number: 726361(IMPROVE).

Availability of data and materials The data used in the analyses is available at <https://osf.io/jcge7/> and <https://osf.io/ejfsv/>.

Code Availability All R code is available at <https://osf.io/jcge7/> and <https://osf.io/ejfsv/>.

Declarations

Conflict of Interests There are no conflicts of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

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. <https://doi.org/10.1002/job.719>
- Baker, R., & Jackson, D. (2008). A new approach to outliers in meta-analysis. *Health Care Management Science*, 11(2), 121–131. <https://doi.org/10.1007/s10729-007-9041-8>
- Baker, R., & Jackson, D. (2016). New models for describing outliers in meta-analysis. *Research Synthesis Methods*, 7(3), 314–328. <https://onlinelibrary.wiley.com/doi/abs/10.1002/jrsm.1191>
- Bartlett, M. S. (1957). A comment on D. V. Lindley's statistical paradox. *Biometrika*, 44(3-4), 533–534.
- Benjamin, D. J., & Berger, J. O. (2019). Three recommendations for improving the use of *p* values. *The American Statistician*, 73(sup1), 186–191. <https://doi.org/10.1080/00031305.2018.1543135>
- Berger, J. O., & Delampady, M. (1987). Testing precise hypotheses. *Statistical Science*, 2(3), 317–335. <http://www.jstor.org/stable/2245772>
- Berger, J. O., Liseo, B., & Wolpert, R. L. (1999). Integrated likelihood methods for eliminating nuisance parameters. *Statistical Science*, 14(1), 1–28. <https://doi.org/10.1214/ss/1009211804>
- Berry, S. M. (1998). Understanding and testing for heterogeneity across 2×2 tables: Application to meta-analysis. *Statistics in Medicine*, 17(20), 2353–2369. <https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291097-0258%2819981030%2917%3A20%3C2353%3A%3AAID-SIM923%3E3.0.CO%3B2-Y>
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2009). *Introduction to meta-analysis*. Chichester: John Wiley & Sons, Ltd. ISBN 9780470057247 0470057246
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2010). A basic introduction to fixed-effect and random-effects models for meta-analysis. *Research Synthesis Methods*, 1(2), 97–111. <https://doi.org/10.1002/jrsm.12>
- Borenstein, M., Higgins, J. P. T., Hedges, L. V., & Rothstein, H. R. (2017). Basics of meta-analysis: I2 is not an absolute measure of heterogeneity. *Research Synthesis Methods*, 8(1), 5–18. <https://onlinelibrary.wiley.com/doi/abs/10.1002/jrsm.1230>
- Kooperberg, C. (2020). logspline: Routines for logspline density estimation. <https://cran.r-project.org/package=logspline>
- Chung, Y., Rabe-Hesketh, S., & Choi, I. H. (2013). Avoiding zero between-study variance estimates in random-effects meta-analysis. *Statistics in Medicine*, 32(23), 4071–4089.
- Cohen, J. (1988). *Statistical power analysis for the behavioral sciences*, (2nd edition). Hillsdale: Lawrence Erlbaum Associates. ISBN 0805802835 9780805802832
- Cordray, D. S., & Morphy, P. (2009). Research synthesis and public policy. In Cooper, H., Hedges, L. V., & Valentine, J. C. (Eds.) *The Handbook of Research Synthesis and Meta-Analysis*, (pp. 473–493). New York: Russell Sage Foundation.
- Liang, F., Paulo, R., Molina, G., A Clyde, M., & O Berger, J. (2008). Mixtures of *g* priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103(481), 410–423. <https://doi.org/10.1198/016214507000001337>
- Fox, J.-P., Mulder, J., & Sinharay, S. (2017). Bayes factor covariance testing in item response models. *Psychometrika*, 82(4), 979–1006. <https://doi.org/10.1007/s11336-017-9577-6>
- 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, (Vol. 2, pp. 123–138), <https://doi.org/10.1080/23743603.2017.1326760>
- Head, M. L., Holman, L., Lanfear, R., Kahn, A. T., & Jennions, M. D. (2015). The extent and consequences of *p*-hacking in science. *PLoS Biology*, 13(3), e1002106. <https://doi.org/10.1371/journal.pbio.1002106>
- Hedges, L. V. (2019). Stochastically dependent effect sizes. In Cooper, H., Hedges, L. V., & Valentine, J. C. (Eds.) *The handbook of research synthesis and meta-analysis*. (3rd edition, pp. 281–297). New York: Russell Sage Foundation.
- Hedges, L. V., & Olkin, I. (1985). *Statistical methods for meta-analysis*. Orlando: Academic Press. ISBN 0123363802 9780123363800
- Higgins, J. P. T., & Thompson, S. G. (2002). Quantifying heterogeneity in a meta-analysis. *Statistics in Medicine*, 21(11), 1539–1558. <https://doi.org/10.1002/sim.1186>
- Higgins, J. P. T., Thompson, S. G., Deeks, J. J., & Altman, D. G. (2003). Measuring inconsistency in meta-analyses. *British Medical Journal*, 327(7414), 557–560. <https://doi.org/10.1136/bmj.327.7414.557>
- Higgins, J. P. T., Thompson, S. G., & Spiegelhalter, D. J. (2009). A re-evaluation of random-effects meta-analysis. *Journal of the Royal Statistical Society*, 172(1), 137–159.
- Ho, M. S. K., & Lee, C. W. (2012). Cognitive behaviour therapy versus eye movement desensitization and reprocessing for post-traumatic disorder - is it all in the homework then?. *Revue Européenne de Psychologie Appliquée/European Review of Applied Psychology*, 62(4), 253–260.
- Hox, J. J., Moerbeek, M., & Van de Schoot, R. (2018). *Multilevel analysis: Techniques and applications*. Routledge, New York. ISBN 9781315650982 1315650983 9781138121409 1138121401 9781138121362 1138121363 9781317308683 1317308689.
- Ioannidis, J. P., Trikalinos, T. A., & Zintzaras, E. (2006). Extreme between-study homogeneity in meta-analyses could offer useful insights. *Journal of Clinical Epidemiology*, 59(10), 1023–1032. <https://doi.org/10.1016/j.jclinepi.2006.02.013>

- Jackson, D. (2009). The significance level of the standard test for a treatment effect in meta-analysis. *Statistics in Biopharmaceutical Research*, 1(1), 92–100. <https://doi.org/10.1198/sbr.2009.0009>
- Jackson, D., Riley, R., & White, I. R. (2011). Multivariate meta-analysis: Potential and promise. *Statistics in Medicine*, 30(20), 2481–2498. <https://doi.org/10.1002/sim.417>
- Jeffreys, H. (1961). *Theory of probability*, (3rd edition). Oxford: Clarendon Press.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430), 791.
- Kass, R. E., & Wasserman, L. (1995). A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *Journal of the American Statistical Association*, 90(431), 928–934. <https://doi.org/10.1080/01621459.1995.10476592>
- Konstantopoulos, S. (2011). Fixed effects and variance components estimation in three-level meta-analysis. *Research Synthesis Methods*, 2(1), 61–76. <https://onlinelibrary.wiley.com/doi/abs/10.1002/jrsm.35>
- Konstantopoulos, S., & Hedges, L. V. (2019). Statistically analyzing effect sizes Fixed- and random-effects models. In Cooper, H., Hedges, L. V., & Valentine, J. C. (Eds.) *The handbook of research synthesis and meta-analysis*. (3rd edition, pp. 245–280). New York: Russell Sage Foundation.
- Kontopantelis, E., Springate, D. A., & Reeves, D. (2013). A re-analysis of the Cochrane Library data. The dangers of unobserved heterogeneity in meta-analyses. *PLOS ONE* 8 (7).
- Kraemer, H. C., Gardner, C., Brooks, J., & Yesavage, J. A. (1998). Advantages of excluding underpowered studies in meta-analysis. Inclusionist versus exclusionist viewpoints. *Psychological Methods*, 3(1), 23–31. <https://doi.org/10.1037/1082-989X.3.1.23>
- Lambert, P. C., Sutton, A. J., Abrams, K. R., & Jones, D. R. (2002). A comparison of summary patient-level covariates in meta-regression with individual patient data meta-analysis. *Journal of Clinical Epidemiology*, 55(1), 86–94. <http://www.sciencedirect.com/science/article/pii/S0895435601004140>
- Lane, D. M., & Dunlap, W. P. (1978). Estimating effect size Bias resulting from the significance criterion in editorial decisions. *British Journal of Mathematical & Statistical Psychology*, 31, 107–112.
- Langan, D., Higgins, J. P. T., & Simmonds, M. (2016). Comparative performance of heterogeneity variance estimators in meta-analysis: A review of simulation studies. *Research Synthesis Methods*, 8(2), 181–198. <https://doi.org/10.1002/jrsm.1198>
- Lee, K. J., & Thompson, S. G. (2008). Flexible parametric models for random-effects distributions. *Statistics in Medicine*, 27(3), 418–434. <https://doi.org/10.1002/sim.2897>
- Lindley, D. V. (1957). A statistical paradox. *Biometrika*, 44(1-2), 187–192. <https://doi.org/10.1093/biomet/44.1-2.187>
- Lunn, D., Barrett, J., Sweeting, M., & Thompson, S. G. (2013). Fully Bayesian hierarchical modelling in two stages, with application to meta-analysis. *Journal of the Royal Statistical Society: Series C*, 62(4), 551–572. <https://rss.onlinelibrary.wiley.com/doi/abs/10.1111/rssc.12007>
- Mulder, J., & Fox, J.-P. (2013). Bayesian tests on components of the compound symmetry covariance matrix. *Statistics and Computing*, 23(1), 109–122. <https://doi.org/10.1007/s11222-011-9295-3>
- Mulder, J., & Fox, J.-P. (2019). Bayes factor testing of multiple intraclass correlations. *Bayesian Analysis*, 14(2), 521–552. <https://projecteuclid.org:443/euclid.ba/1533866668>
- Mulder, J., Williams, D. R., Gu, X., Tomarken, A., Boeing-Messing, F., Olsson-Collentine, A., . . . , van Lissa, C. (in press). BFPack: Flexible Bayes factor testing of scientific theories in R. *Journal of Statistical Software*.
- Nielsen, N. M., A C Smink, W., & Fox, J.-P. (2021). Small and negative correlations among clustered observations: Limitations of the linear mixed effects model. *Behaviormetrika*, 48(1), 51–77. <https://doi.org/10.1007/s41237-020-00130-8>
- Normand, S.-L. T. (1999). Meta-analysis: Formulating, evaluating, combining, and reporting. *Statistics in Medicine*, 18(3), 321–359.
- R Core Team, R. (2020). A language and environment for statistical computing. <http://www.r-project.org/>
- Raftery, A. E. (1995). Bayesian model selection in social research. *Sociological Methodology*, 25, 111–163. <https://doi.org/10.2307/271063>
- Raudenbush, S. W. (2009). Analyzing effect sizes: Random-effects models. In Cooper, H., Hedges, L. V., & Valentine, J. C. (Eds.) *The Handbook of Research Synthesis and Meta-Analysis*, (pp. 295–315). New York: Russell Sage Foundation.
- Raudenbush, S. W., & Bryk, A. S. (1985). Empirical Bayes meta-analysis. *Journal of Educational Statistics*, 10(2), 75–98.
- Rhodes, K. M., Turner, R. M., & Higgins, J. P. (2015). Predictive distributions were developed for the extent of heterogeneity in meta-analyses of continuous outcome data. *Journal of Clinical Epidemiology*, 68(1), 52–60.
- Rouder, J. N., & D Morey, R. (2012). Default Bayes factors for model selection in regression. *Multivariate Behavioral Research*, 47(6), 877–903. <https://doi.org/10.1080/00273171.2012.734737>
- Rouder, J. N., Haaf, J. M., Davis-Stober, C. P., & Hilgard, J. (2019). Beyond overall effects: A Bayesian approach to finding constraints in meta-analysis. *Psychological Methods*, 24(5), 606–621. <https://doi.org/10.1037/met0000216>
- Rouder, J. N., & Morey, R. D. (2011). A Bayes factor meta-analysis of Bems ESP claim. *Psychonomic Bulletin & Review*, 18(4), 682–689.
- Rücker, G., Schwarzer, G., Carpenter, J. R., & Schumacher, M. (2008). Undue reliance on I² in assessing heterogeneity may mislead. *BMC Medical Research Methodology*, 8(1), 79. <https://doi.org/10.1186/1471-2288-8-79>
- Scheibehenne, B., Gronau, Q. F., Jamil, T., & Wagenmakers, E.-J. (2017). Fixed or random? A resolution through model averaging: Reply to Carlsson, Schimmack, Williams, and Bürkner (2017). <https://journals.sagepub.com/doi/abs/10.1177/0956797617724426> (Vol. 28, pp. 1698–1701).
- Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2), 461–464.
- Sellke, T., Bayarri, J. M., & Berger, O. J. (2001). Calibration of *p* values for testing precise null hypotheses. *The American Statistician*, 55(1), 62–71. <http://www.jstor.org/stable/2685531>
- Sidik, K., & Jonkman, J. N. (2007). A comparison of heterogeneity variance estimators in combining results of studies. *Statistics in Medicine*, 26(9), 1964–1981.
- Smith, T. C., Spiegelhalter, D. J., & Thomas, A. (1995). Bayesian approaches to random-effects meta-analysis: A comparative study. *Statistics in Medicine*, 14(24), 2685–2699.
- Sutton, A. J., & Abrams, K. R. (2001). Bayesian methods in meta-analysis and evidence synthesis. *Statistical Methods in Medical Research*, 10(4), 277–303. <https://journals.sagepub.com/doi/abs/10.1177/09622820101000404>
- Thompson, S. G., & Sharp, S. J. (1999). Explaining heterogeneity in meta-analysis a comparison of methods. *Statistics in Medicine*, 18(20), 2693–2708. [https://doi.org/10.1002/\(sici\)1097-0258\(19991030\)18:20%3C2693::aid-sim235%3E3.0.co;2-v](https://doi.org/10.1002/(sici)1097-0258(19991030)18:20%3C2693::aid-sim235%3E3.0.co;2-v)
- Turner, R. M., & Higgins, J. P. T. (2019). Bayesian meta-analysis. In Cooper, H., Hedges, L. V., & Valentine, J. C. (Eds.) *The handbook of research synthesis and meta-analysis*. (3rd edition, pp. 299–314). New York: Russell Sage Foundation.
- Turner, R. M., Jackson, D., Wei, Y., Thompson, S. G., & Higgins, J. P. T. (2015). Predictive distributions for between-study heterogeneity and simple methods for their application in Bayesian meta-analysis. *Statistics in Medicine*, 34(6), 984–998.
- van Aert, R. C. M., & Jackson, D. (2019). A new justification of the Hartung-Knapp method for random-effects meta-analysis based

- on weighted least squares regression. *Research Synthesis Methods*, 10(4), 515–527. <https://onlinelibrary.wiley.com/doi/abs/10.1002/jrsm.1356>
- van Assen, M. A. L. M., van Aert, R. C. M., & Wicherts, J. M. (2015). Meta-analysis using effect size distributions of only statistically significant studies. *Psychological Methods*, 20(3), 293–309. <https://doi.org/10.1037/met0000025>
- van den Noortgate, W., & Onghena, P. (2003). Multilevel meta-analysis a comparison with traditional meta-analytical procedures. *Educational and Psychological Measurement*, 63(5), 765–790.
- Van Houwelingen, H. C., Arends, L. R., & Stijnen, T. (2002). Advanced methods in meta-analysis: Multivariate approach and meta-regression. *Statistics in Medicine*, 21(4), 589–624. [1002/sim.1040](https://doi.org/10.1002/sim.1040).
- 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. <https://doi.org/10.1002/jrsm.1164>
- Viechtbauer, W. (2007). Confidence intervals for the amount of heterogeneity in meta-analysis. *Statistics in Medicine*, 26(1), 37–52. <https://doi.org/10.1002/sim.2514>
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3), 1–48. <https://doi.org/10.18637/jss.v036.i03>
- Whittaker, R., McRobbie, H., Bullen, C., Rodgers, A., Gu Y, & Dobson, R. (2019). Mobile phone text messaging and app-based interventions for smoking cessation. *Cochrane Database of Systematic Reviews*,(10). <https://doi.org/10.1002/14651858.CD006611.pub5>
- Xu, H., Platt, R. W., Luo, Z.-C., Wei, S., & Fraser, W. D. (2008). Exploring heterogeneity in meta-analyses: Needs, resources and challenges. *Paediatric and Perinatal Epidemiology*, 22(Suppl 1), 18–28. <http://europepmc.org/abstract/MED/18237348> <https://doi.org/10.1111/j.1365-3016.2007.00908.x>
- Zellner, A. (1986). On assessing prior distributions and Bayesian regression analysis with g-prior distributions. In Goel, P. K., & Zellner, A. (Eds.) *Bayesian inference and decision techniques: Essays in honor of Bruno de Finetti*. Amsterdam: Elsevier.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.