

This is a repository copy of *On weakly informative prior distributions for the heterogeneity parameter in Bayesian random-effects meta-analysis*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/170174/>

Version: Published Version

Article:

Röver, Christian, Bender, Ralf, Dias, Sofia orcid.org/0000-0002-2172-0221 et al. (5 more authors) (2021) On weakly informative prior distributions for the heterogeneity parameter in Bayesian random-effects meta-analysis. *Research Synthesis Methods*. pp. 448-474. ISSN: 1759-2887

<https://doi.org/10.1002/jrsm.1475>

Reuse






This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

On weakly informative prior distributions for the heterogeneity parameter in Bayesian random-effects meta-analysis

Christian Röver¹  | Ralf Bender²  | Sofia Dias³  |
Christopher H. Schmid⁴  | Heinz Schmidli⁵ | Sibylle Sturtz² |
Sebastian Weber⁶ | Tim Friede¹ 

¹Department of Medical Statistics,
University Medical Center Göttingen,
Göttingen, Germany

²Department of Medical Biometry,
Institute for Quality and Efficiency in
Health Care (IQWiG), Köln, Germany

³Centre for Reviews and Dissemination,
University of York, York, UK

⁴Department of Biostatistics and Center
for Evidence Synthesis in Health, Brown
University School of Public Health,
Providence, Rhode Island, USA

⁵Statistical Methodology, Development,
Novartis Pharma AG, Basel, Switzerland

⁶Advanced Exploratory Analytics,
Novartis Pharma AG, Basel, Switzerland

Correspondence

Christian Röver, Department of Medical
Statistics, University Medical Center
Göttingen, Humboldtallee 32, Göttingen
37075, Germany.

Email: christian.roever@med.uni-
goettingen.de

Funding information

Deutsche Forschungsgemeinschaft, Grant/
Award Number: FR 3070/3-1

Abstract

The normal-normal hierarchical model (NNHM) constitutes a simple and widely used framework for meta-analysis. In the common case of only few studies contributing to the meta-analysis, standard approaches to inference tend to perform poorly, and Bayesian meta-analysis has been suggested as a potential solution. The Bayesian approach, however, requires the sensible specification of prior distributions. While noninformative priors are commonly used for the overall mean effect, the use of *weakly informative* priors has been suggested for the heterogeneity parameter, in particular in the setting of (very) few studies. To date, however, a consensus on how to generally specify a weakly informative heterogeneity prior is lacking. Here we investigate the problem more closely and provide some guidance on prior specification.

KEYWORDS

Bayes factor, GLMM, hierarchical model, marginal likelihood, variance component

1 | INTRODUCTION

In meta-analysis, researchers commonly encounter a certain amount of variability between experiments, to a degree going beyond what could be attributed to measurement error alone. Hierarchical models are commonly

used in order to account for such (“between-study”) heterogeneity.^{1,2} In the present paper, we focus on the special simple case of meta-analysis within the framework of the *normal-normal hierarchical model (NNHM)*. The NNHM approximates estimates from separate sources and their standard errors via normal distributions, and

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. *Research Synthesis Methods* published by John Wiley & Sons Ltd.

implements heterogeneity at a second level using another normal variance component. In meta-analysis applications, the NNHM provides a good approximation for many types of endpoints or effect measures.^{3,4} The normal approximation has its limitations,⁵ some of which are less of a problem in a Bayesian context.⁶ A small number of studies tends to pose a problem especially for frequentist methods, in particular regarding the construction of confidence intervals (CIs) with good coverage properties.⁷⁻¹⁰ A common convention is to exercise extra caution when the number of studies is small.⁹

Bayesian approaches to meta-analysis have been advocated for quite a while,¹¹⁻¹⁷ and analyses may technically be performed using MCMC methods¹ or semi-analytical integration.¹⁸ Within the R software, for example, the *bayesmeta*^{19,20} or *bmeta*²¹ packages are available. Performing a Bayesian analysis is not technically challenging; computations are straightforward and valid for any number of studies, although less data will mean that results are more sensitive to prior specifications (especially when it comes to variance parameters). A crucial condition is that the explicitly implemented normal approximation needs to hold, which may break down, for example, for meta-analyses of *small* studies.^{5,6} While for large numbers of studies, the choice of prior distributions usually has little impact, for few studies the exact form of the prior distributions chosen may become crucial, as one cannot rely on the prior information being overruled by the data in that case. At least part of this problem may be considered “shared” for frequentist and Bayesian methods as long as one tries to get by without using a proper, informative prior.²² Some supposedly *non-informative* prior distributions can probably be argued to be less influential than others, but ultimately these are unlikely to be the best choice in few-study problems. Beyond meta-analysis, the use of informative priors for regularization in the estimation of certain parameters is also common.²³ Especially for few studies, this may be a promising approach.²⁴ The case of “few” studies is hard to define; there is no obvious threshold, and in fact there may actually be no need to distinguish: use of an informative prior will not be harmful for analyses of “many” studies. Indeed, a proper prior is necessary irrespective of the number of studies in case the analysis requires the calculation of marginal likelihoods. In the present manuscript, we will investigate examples ranging in size between two and five studies. These are the cases where the use of an informative prior will make the greatest difference, and such situations have been discussed in the context of up to 4,⁹ 3–10,⁷ or only 2 studies.⁸

Heterogeneity priors have been investigated previously from different angles; some discussed general considerations for variance parameters^{15,25,26} while others motivated particular settings for specific example

cases^{27,28} or investigated commonly used settings in a systematic literature review.²⁹ The aim of the present investigation is to provide general guidance for judging and deriving *weakly informative* heterogeneity priors, and to suggest consensus examples for some common types of effect measures. This may also aid in the design and justification of prior settings, or the prospective pre-specification of Bayesian meta-analyses³⁰ and it may help avoid (suspicion of) post-hoc tweaking of prior assumptions.

The remainder of this article is structured as follows. In the next section, the normal-normal hierarchical model (NNHM) along with its parameters and prior distributions are formally introduced. Section 3 discusses prior distributions for the heterogeneity parameter and some general motivating considerations and implications. Section 4 motivates heterogeneity priors for a selection of common types of endpoints and effect measures based on the previously discussed ideas. In Section 5, examples of meta-analyses with different endpoints are introduced, and analyses are performed using the suggested prior settings. Section 6 closes with conclusions and recommendations.

2 | THE STATISTICAL MODEL

2.1 | The normal-normal hierarchical model

The *normal-normal hierarchical model* (NNHM) represents measurements y_i from k different sources using two hierarchy levels. Along with the estimates, their associated standard errors σ_i need to be available. The σ_i are assumed to be fixed and known (which commonly is only an approximation).^{5,31} Each estimate y_i is assumed to measure an underlying true value θ_i , which is not necessarily identical across all k measurements; (“between-study”) variability among the θ_i is accounted for by an additional variance component whose magnitude is given by the heterogeneity $\tau \geq 0$:

$$y_i | \theta_i \tilde{N}(\theta_i, \sigma_i^2), \quad (1)$$

$$\theta_i | \mu, \tau \tilde{N}(\mu, \tau^2) \quad \text{for } i = 1, \dots, k, \quad (2)$$

where the estimates y_i (as well as the θ_i) are modeled as exchangeable. The overall mean effect μ is often the figure of primary interest. By marginalizing over the θ_i values, the model may be written in simplified form:

$$y_i | \mu, \tau \tilde{N}(\mu, \sigma_i^2 + \tau^2). \quad (3)$$

This is a random-effects model, which in the special case of $\tau = 0$ simplifies to the common-effect model (also

known as the fixed-effect model).^{3,4,20,32} The NNHM provides a good approximation for many types of effect measures where the estimates as well as between-study variability may be assumed to be (approximately) normally distributed.⁵

While often the aim of a meta-analysis is estimation of the overall mean μ , it is sometimes useful to also infer the study-specific means θ_i or a prediction θ_{k+1} . The amount of information gained on θ_i or θ_{k+1} through the joint meta-analysis depends very much on the amount of heterogeneity τ . If there was no heterogeneity ($\tau = 0$), then we would have $\theta_1 = \theta_2 = \dots = \theta_{k+1} = \mu$, and all data would essentially contribute to the estimation of a single common parameter. If, on the other hand, τ was very large, then different parameters θ_i would only be very loosely connected (2), and consideration of additional data would only add very little to the estimation of any particular θ_i or to a prediction θ_{k+1} . In between, for moderate τ values, estimates of θ_i are somewhat “shrunk” towards the overall mean μ , and the prediction θ_{k+1} is also more tightly constrained. Estimation of the heterogeneity τ hence also has distinct effects on the so-called “shrinkage estimates” θ_i as well as predictions θ_{k+1} .^{20,33}

2.2 | Prior distributions

2.2.1 | Effect and heterogeneity priors

In the NNHM, there are two unknowns requiring prior specification, namely the overall mean effect μ and the heterogeneity τ . In the following, we will assume that the prior may be factored into $p(\mu, \tau) = p(\mu) \times p(\tau)$, implying prior independence of μ and τ ; note though that one may also argue in favor of a dependent prior.^{22,34} In a sense, dependence is often implicitly implemented, for example, in the case of log-transformed effect scales: on the back-transformed (exponentiated) scale, the amount of heterogeneity then scales with the value of the effect.

The effect prior $p(\mu)$ may often, also for technical convenience, be taken to be (improper) uniform or normal.²⁰ In case a proper, informative effect prior is used, this may also have implications for the heterogeneity prior; in particular the prior variance of μ may be relevant when considering reasonable τ values (see also Section 3.4.2 below).

Here we are first of all concerned with the prior distribution for the heterogeneity, $p(\tau)$. A number of priors have been proposed that may be considered “non-informative” in particular senses (e.g., improper uniform or Jeffreys priors, which may be motivated using invariance or information-theoretic arguments),²⁰ but these usually cause problems especially when the number of studies (k) is sufficiently small, or when the computation

of marginal likelihoods (or Bayes factors) is desired. In the following, we will hence be concerned with proper, (weakly) informative priors.

2.2.2 | Different views of prior specification

There may be different perspectives on the role or purpose of prior specification within a Bayesian analysis; we sketch three aspects here:

1. **Epistemic point-of-view:** The posterior distribution depends on the prior via Bayes' theorem; the prior inevitably needs to enter inference, reflecting the state of information beyond the data at hand.^{1,35} Prior assumptions simply add to the line of other assumptions being made, like a normal likelihood, independence, known standard errors, etc.
2. **Regularization point-of-view:** The aim is to introduce “weakly informative priors, which attempt to let the data speak while being strong enough to exclude various ‘unphysical’ possibilities which, if not blocked, can take over a posterior distribution in settings with sparse data” (Gelman, 2009).³⁶ This perspective is closely connected to regularization or penalization approaches in general.³⁷ While in the likelihood framework it may sometimes be perceived as a rather *ad hoc* fix, it constitutes a transparent, readily interpretable model component in the Bayesian case.
3. **Pragmatic point-of-view:** The resulting estimates may be judged solely based on their operating characteristics (which may be frequentist or Bayesian¹), without worrying about their exact theoretic underpinning.

The first viewpoint is probably the most “constructive” one here, in the sense of providing guidance on sensible prior choices. An example of a regularization approach in the NNHM context is given by the procedure proposed by Chung et al.,³⁸ where regularization is used to implement preference for positive τ values. Alternatively, one may also give preference to small τ values, as these imply a less complex model, which is the idea behind penalized complexity priors³⁹ (and which here would lead to an exponential prior). Comparisons of operating characteristics (also including frequentist approaches) were done, for example, by Friede et al.⁷ There are probably more perspectives beyond or between these three (e.g., Bayarri and Berger⁴⁰ and Kass⁴¹). For example, meta-analyses may be thought of as constituting draws from a “population” whose associated heterogeneities are reflected in the prior distribution—an “aleatory” interpretation of (prior) probability, which may lend a

somewhat frequentist flavor to the analysis. An important point to stress is that there is not necessarily a single “correct” prior: the use of different priors may be seen as basing inferences on different preconditions, and the choice of prior depends on which information one is willing to incorporate into the analysis; different analysts may hence draw different conclusions from the same data, when these are founded on differing prior beliefs.⁴² In a sense, the posterior inherits its meaning from the prior to some extent.⁴³ Other common shortcuts taken or approximations and asymptotics relied upon may in fact often be potentially more influential and relevant than the choice among the (usually limited) set of reasonable prior distributions (see, e.g., Jackson and White⁵).

2.2.3 | Implications for interval estimation

While (frequentist) confidence intervals aim to provide coverage of the true parameter *uniformly*, independent of the actual current parameter value, this is generally not the case for (Bayesian) credible intervals. In some cases, it is possible to specify (often improper) priors leading to posterior distributions that also provide proper frequentist coverage, but usually such a prior is not available.⁴⁴ Credible intervals are calibrated and yield proper coverage *on average* across the prior distribution; for the point-wise coverage this means that there may be over-coverage in certain regions of parameter space and undercoverage in others.^{20,45–47} For example, in the present case this may mean that long-run coverage may be above the nominal level if data were repeatedly generated based on heterogeneity values from the lower end of the prior range, and below the nominal level otherwise.

3 | HETEROGENEITY PRIORS

3.1 | Aim

For meta-analyses involving many studies (large k), the choice of prior distribution often has little impact, and an (improper) uniform prior for τ may be a good choice, not least due to its invariance property.^{20,25} Here we are concerned first of all with the case of few studies (small k); a uniform prior may not actually be an option here, as it requires $k \geq 3$ studies in order to yield a proper, integrable posterior,²⁵ and it may otherwise generally be considered overly conservative.^{7,8,25} Similar problems arise also with the Jeffreys prior for the NNHM model;²⁰ this kind of issue is common in Bayesian analysis.²³ Another case where a proper, weakly informative prior may be

required (not only for few studies) is when marginal likelihoods or Bayes factors are of interest.

While the availability of a “noninformative” prior comes with a certain convenience (one less issue to worry about), in the present case its failure to provide reasonable estimates in certain instances will often appear somewhat contradictory to common sense. The introduction of an informative prior then may entail a trade-off of the introduced regularization versus simplicity and robustness. On the other hand, the explicit consideration of relevant prior information may also be seen as an advantage.

From a merely “technical” perspective, a heterogeneity prior must (in order to ensure integrability of the posterior) have a shorter-than-uniform upper tail (an eventually decreasing, integrable density function) and also an integrable density towards zero. In that spirit, it may also make sense to consider near-origin- and upper-tail-behaviors separately. While an (improper) uniform prior may be considered noninformative for several reasons (e.g., due to its scale-invariance property²⁰), its overly heavy upper tail may also be considered “anti-conservative”.⁴⁸ On the other hand, it may be possible to “rescue” some of the desirable behavior and robustness, for example, by the use of heavy-tailed priors.⁴⁹ Besides upper-tail considerations, priors may also behave quite differently near zero; for example, depending on whether the prior density approaches zero, a finite value, or infinity. A finite prior density may ensure a near-zero behavior roughly like a uniform prior, while a zero density may be useful, for example, in bounding maximum-a-posteriori (MAP) point estimates away from zero;³⁸ in particular from the regularization perspective, the prior density’s derivative near zero may also be of interest (as it determines how small τ values may be pushed towards or away from zero).

While the concept of “weak informativeness” remains somewhat elusive (just like that of a “noninformative” prior), the information content (or “vagueness”) of a prior is commonly related to its variance, its entropy,⁵⁰ or its associated effective sample size (ESS).^{51,52} In many cases it is also helpful to consider the informativeness of a prior relative to a reference,⁵³ for example, a unit information prior.^{26,54} Since the posterior draws its interpretation in part from the prior, it is important to make the prior specification plausible and transparent. Following the parsimony principle (*Ockham’s razor*), it may be constructive to seek the (in some sense) *simplest* prior distribution within any relevant constraints.⁵⁵ Possible approaches to implement such a notion in practice may work, for example, via maximization of the entropy,⁵⁰ pre-specification of an effective sample size,^{51,52} or matching of moments.

Despite the aim of a weakly informative formulation, one should also anticipate the case where the data have little information to add, so that the posterior closely resembles the prior and hence the analysis results are largely determined by the prior settings. This may happen especially in cases of few studies and is also suggested in some of the examples that will be discussed below (see Figure 8); such cases highlight the importance of a transparent and convincing prior specification.

In the remainder of this section, we aim to facilitate a structured approach to interpreting heterogeneity and specifying heterogeneity prior distributions by pointing out relevant perspectives and highlighting consequences of certain heterogeneity settings. Similar ideas are to some degree also utilized in prior elicitation in general.^{56,57} A set of guiding questions is eventually suggested in Table 6.

3.2 | General properties of the NNHM

When considering prior distributions for the heterogeneity τ , it is useful to recall that $\tau \geq 0$ is a scale parameter, and that its square τ^2 denotes a variance component within the NNHM. Immediate associations of variance priors useful in a simple normal model however may be misleading: inverse-gamma (or inverse- χ^2) distributions are usually not recommended, as these arise as conjugate distributions only in related, yet distinctly different circumstances. An inverse-gamma distribution is conjugate in the simple case of estimating the variance of a normal distribution with known mean.¹ In such a case, an unequal pair of two data points for example implies that the variance must be positive (a zero variance would have a zero likelihood); in the present NNHM context, however, unequal y_i values may be consistent with zero heterogeneity ($\tau = 0$), so that such priors are not a natural choice here, and their use is generally discouraged.^{2,25,58,59} Supposedly noninformative settings based on inverse-gamma distributions commonly tend to result in sensitivity to specification details,²⁵ and often too much probability is allocated to very large heterogeneity values.⁶⁰

For uniform or normal effect prior distributions, the resulting *conditional* effect posterior $p(\mu|\tau, y)$ again is normal. While for increasing τ the (conditional) posterior mean of μ shifts from the inverse-variance weighted mean towards the unweighted average of the estimates y_i , the (conditional) posterior variance of μ is proportional to τ .²⁰ At the same time, larger heterogeneity values also imply wider prediction intervals and less shrinkage^{16,20,61–63} (see also Section 2.1). Varying τ between zero and infinity essentially also means varying between the

extremes of *pooled* and *separate* analyses of individual studies. In a sense, overestimation of τ may hence often be considered a “conservative” or “less harmful” form of bias. In that spirit, one might argue that—within reasonable limits—a prior that is *stochastically larger* than another is also *more conservative*.⁶⁴ A simple way to implement stochastically ordered distribution families is by using parameterizations that include a scale parameter.⁶⁵ Use of a scale parameter does not actually impose a restriction; if not already included in the parameterization, it may easily be introduced. Note that simple re-scaling of a prior distribution $p(\tau)$ then also implies a (re) scaling of the corresponding marginal prior predictive distributions $p(\theta_i|\mu)$ by the same factor. In general, stochastically ordered priors also imply the same ordering for the resulting posteriors.^{63,66,67} Consideration of stochastically ordered alternative priors may hence also offer a framework for sensitivity analyses (see also Supporting Information Appendix D.4).

3.3 | Reasonable (proper) distributional families

A simple way to implement the “technical” requirements (as suggested in Section 3.1) may be to require roughly uniform behavior near zero (implying indifference among small heterogeneity values on the τ scale and ensuring integrability in the lower tail), and a monotonically decaying tail with increasing heterogeneity values (implying decreasing probability for increasing τ values and ensuring integrability in the upper tail). This may be achieved, for example, by using half-normal, half-Student- t , half-Cauchy, half-logistic, exponential or Lomax distributions. A sample of such distributions is sketched in Figure 1. Note that for comparability, the distributions in the figure are all scaled such that they have a common median of 1; their corresponding parameters are also listed in Table 4 below. In particular, half-normal, half-Student- t , or half-Cauchy distributions have been recommended as appropriate families within the NNHM, also due to favorable frequentist properties.^{2,25,58} The half-Student- t distribution (including the half-Cauchy as a special case, and the half-normal as a limiting case) may be derived as conditionally conjugate distributions in an extended parameterization of the NNHM.² The exponential distribution might be motivated as the *maximum entropy* distribution for a pre-specified prior expectation,⁵⁰ or as the *penalized complexity* prior.³⁹ The half-logistic distribution combines a zero derivative (implying near-uniform behavior) at the origin with an upper tail behavior close to that of an exponential distribution.




FIGURE 1 A selection of potential probability densities for the heterogeneity. All distributions are scaled so that their prior median is at unity ($\tau = 1$, dashed line; see also Table 4) [Colour figure can be viewed at wileyonlinelibrary.com]

Half-Student- t and Lomax distributions here may be considered as heavy-tailed variants of the half-normal and exponential distributions, respectively. In the spirit of a *contaminated* prior, encompassing priors “close to an elicited one”,^{68,69} these may also be motivated as scale mixtures, where the (exponential or half-normal) scale parameter is associated with some variability or uncertainty. The scale mixture connection is also derived in detail in Supporting Information Appendix C. The special case of a Lomax ($\alpha = 1$) distribution also coincides with the form of prior distribution suggested by DuMouchel (a log-logistic prior for τ).^{70,71} Similarly, the exponential distribution may also be motivated as a scale mixture of a half-normal distribution with Rayleigh-distributed scale. The use of heavy-tailed prior distributions has the advantage of ensuring some degree of robustness against prior misspecification (or prior/data conflict)⁴⁹ at the cost of sacrificing some of its “regularization” power. Another simple way of implementing some degree of robustness is by combining “informative” and “heavy-tailed” elements in a two-component mixture distribution.^{72,73}

Another simple and common prior distribution is the (proper) bounded uniform distribution defined on an interval $[0, a]$. It inherits certain qualities from the (improper) uniform distribution, but it introduces a sharp cutoff at the upper bound a , which may be hard to motivate or justify. Although, if the bound is large enough, then it may be very reasonable (e.g., for log-ORs).

Among the above examples, the Student- t and Lomax distributions possess “shape” parameters in addition to scale parameters, which here essentially regulate the degree of heavy-tailedness. If considered desirable, more complex prior assumptions may be implemented using more complex distributions, for example, using folded

non-central Student- t distributions with a non-zero mode,^{2,25,58} however, additional degrees of complexity would probably require solid justification to be convincing. In the context of a *penalization* interpretation of the prior, a mode at zero also implies a corresponding “penalty term” that is monotonically increasing in τ ; this applies, for example, for a penalized-complexity prior³⁹ that aims to give preference to sparse models. In empirical investigations based on meta-analyses archived in the *Cochrane Database of Systematic Reviews*, log-Normal and log-Student- t_5 distributions have been fitted to empirical data.^{74,75} The log-normal and log- t distributions here were found to fit the predictive distributions best, however, only few alternatives (log-normal, log- t_5 and inverse-gamma,⁷⁴ or log-normal, inverse-gamma and gamma distributions⁷⁵ for τ^2) were considered as candidates in these comparisons. Some properties of the distributions discussed here are also listed in Supporting Information Appendix B.

In practice, the half-normal distribution is quite commonly used; the reasons for its popularity are probably its simple and familiar form, its near-uniform behavior at the origin along with a reasonably quickly decaying upper tail, as well as considerations of numerical stability. In the following, we will focus mostly on half-normal distributions. In our experience, minor differences between similar prior densities are of rather minor practical relevance, while it is most important what heterogeneity ranges the bulk of prior probability is assigned to.

When eventually formulating prior assumptions in terms of a parametric prior probability distribution, it is first of all necessary to be able to judge the meaning and implications of certain heterogeneity settings; these issues will be discussed in the following section.

3.4 | Interpreting heterogeneity values

3.4.1 | Units of τ

Informative priors naturally *always* need to be considered in the context of the endpoint under consideration. In order to specify a sensible prior for τ , it is important to recapitulate its role in the NNHM (see Section 2.1). The heterogeneity τ is a *scale parameter* that relates to the probable size of *differences* (between-study differences) in effects (θ_i and μ ; see Equation 2). With that, the units of measurements (y_i), effects (θ_i , μ) and heterogeneity (τ) are the same; if the effect is measured, say, in meters, then so is the heterogeneity. Or both may be dimensionless, as, for example, in the case of log-transformed ratios (like log-odds-ratios (log-ORs), log-incidence-rate-ratios (log-IRRs), log-hazard-ratios (log-HRs), and so on) or

standardized mean differences (SMDs). One may in fact argue that the nature of the effect scale is the most important aspect to consider for prior specification.²⁴ In case the effects y_i have been transformed prior to analysis, then it is often useful to consider implications on the back-transformed scale. Transformations are usually introduced to achieve a better fit to the normality assumptions within the NNHM; for example, using logarithmic or arcsine transforms.^{3,4,76} In such cases, also considering the back-transformed (exponential or sine) effect scales is often instructive.

In case the effect scale has definite upper and lower bounds (which is often the case, for example, for endpoints measured as scores), this also provides information on the plausible (and possible) between-study variability. In case of bounded scales, it may for example be useful to consider the extreme cases of a continuous uniform distribution across the considered range (which would have standard deviation $\frac{b-a}{\sqrt{12}} = \frac{b-a}{3.46}$, where a and b are the lower and upper bounds, respectively), or a discrete distribution with probabilities of $\frac{1}{2}$ concentrated at both margins a and b (which would have standard deviation $\frac{b-a}{2}$). Such considerations may define absolute “worst-case” settings for the heterogeneity. Any normal approximation employed on a bounded parameter space with a standard deviation of, say, $> \frac{b-a}{4}$ would inevitably have substantial overlap with out-of-domain values; any heterogeneity value that is not $\ll \frac{b-a}{4}$ should raise suspicion and might actually call for a different approach (e.g., transformation to a different parameter space).

3.4.2 | Magnitudes of other effects

Relevant hints may originate from considering the magnitude of other (known or plausible) effects of interventions or covariates. The reasonable range for the overall mean effect μ may also have implications for the expected range of study-specific means θ_i ; in case an informative prior for μ is used (or is at least plausible), its variance may help constraining also the between-trial variability. Heterogeneity may often be attributed to differences in the composition of the populations underlying each estimate, and the distribution of relevant covariates within (which may be observed or unobserved). If the observed heterogeneity is assumed to be due to different constitutions of populations, then the heterogeneity relates to accumulated effects of associated covariates. With that, within- and between-study variability in effects are related to within- and between-study differences among subjects and the plausible magnitude of covariates' effects. For example, if a treatment effect is known to differ between males and females by a certain amount, this

difference between genders may help judging or motivating plausible magnitudes of effect differences between studies. In case the variability between centers within the same study has been investigated, this may also provide a hint on between-study variability (which will then most likely be larger).

3.4.3 | Implications of a fixed heterogeneity value

Specific values of the heterogeneity τ may be judged and compared based on the implied distribution of true effects θ_i , which is given by the (conditional) prior predictive distribution $p(\theta_i|\mu, \tau)$ (see Equation 2), where τ defines the distribution's standard deviation. The effects θ_i (conditional on μ) then vary within a range of $\mu \pm 1.96\tau$ with 95% probability. For a randomly picked pair of effects (θ_i and θ_j), their difference ($\theta_i - \theta_j$) follows a $N(0, 2\tau^2)$ -distribution (2), and their absolute difference $|\theta_i - \theta_j|$ then has a median of 0.95τ . Quite commonly, the effects θ_i are transformed prior to analysis, so that it may be helpful to consider the implications on the back-transformed scale. A very common example is the logarithmic transformation, which is often used for analyses involving, for example, odds ratios (ORs), relative risks (RRs) or hazard ratios (HRs), and where the inverse transform is the exponential function. Ninety-five percentage predictive intervals and median differences are shown for a range of τ values in Table 1 along with the corresponding exponentiated figures.

An extensive discussion of these conditional distributions is given in Spiegelhalter et al. (Sec. 5.7).¹⁵ By working out what range of θ_i values is expected, or what difference between a randomly picked pair of θ_i values is expected, corresponding plausible ranges of τ values may be determined. Based on such considerations, Spiegelhalter et al.¹⁵ categorized ranges of τ values in the

TABLE 1 Implications of certain fixed heterogeneity values τ on the probable ranges of true effects θ_i (conditional prior predictive distributions) and the corresponding exponentiated ranges (the latter are relevant for log-transformed effect scales)

τ	95% Predictive interval		Random pair $ \theta_i - \theta_j $	
	$\theta_i - \mu$	$\exp(\theta_i - \mu)$	Median	$\exp(\text{median})$
0.1	[−0.20, 0.20]	[0.82, 1.22]	0.10	1.10
0.2	[−0.39, 0.39]	[0.68, 1.48]	0.19	1.21
0.5	[−0.98, 0.98]	[0.38, 2.66]	0.48	1.61
1.0	[−1.96, 1.96]	[0.14, 7.10]	0.95	2.60
2.0	[−3.92, 3.92]	[0.020, 50.4]	1.91	6.74

context of log-ORs as “reasonable,” “fairly high” or “fairly extreme” as shown in Table 2. Such investigations may help judging what τ values are reasonable or unrealistic and with that may help specifying e.g. the heterogeneity prior's tail quantiles.

For example, Prevost et al. (Sec. 4)²⁷ aimed to constrain the predictive interval ($\exp(\theta_i - \mu)$) to a range of [0.5, 2.0], which is achieved for $\tau = 0.35$. Considering this range as extreme and unlikely, a half-Normal prior with scale 0.18 (implying $P[\tau \leq 0.35] = 0.95$) was eventually suggested for a log-RR. R code to illustrate these arguments using Monte Carlo sampling and exact calculations is provided in Supporting Information Appendix D.1.

3.4.4 | Implications of a heterogeneity distribution

Besides considering the *conditional* distribution for fixed τ values ($p(\theta_i|\mu, \tau)$, see previous subsection), one may also investigate the *marginal* prior predictive distribution $p(\theta_i|\mu)$, marginalized over a particular heterogeneity prior, which technically results as the integral $p(\theta_i|\mu) = \int_0^\infty p(\theta_i|\mu, \tau)p(\tau)d\tau$. Since $p(\theta_i|\mu, \tau)$ is normal (2), the marginal $p(\theta_i|\mu)$ is a *normal (scale) mixture* distribution. Its form may usually either be derived numerically,^{18–20} or it may easily be explored using *collapsed Gibbs sampling*, that is, generating a Monte Carlo sample by repeatedly sampling from the heterogeneity prior ($p(\tau)$), and subsequently from the conditional predictive distribution ($p(\theta_i|\tau)$). Investigating the marginal prior predictive distribution may help judging the prior scale or distributional family.

Table 3 illustrates a range of prior predictive distributions for a set of half-normal priors that differ in their scale. The implied probabilities for the (log-OR) categories shown in Table 2 are also given. Note that a simple re-scaling of the heterogeneity prior implies proportional scaling of mean and quantiles for τ as well as θ_i (as can be seen in Table 3). In this spirit, Dias et al.²⁸ for example proposed a half-normal(0.32)-prior for a log-OR based on the implied prediction interval for $\exp(\theta_i - \mu)$ of [0.5, 2.0]. R code to illustrate these arguments using Monte

Carlo sampling and exact calculations is provided in Supporting Information Appendix D.2.

Similarly, Table 4 illustrates a range of prior predictive distributions for a set of heterogeneity priors from different distributional families; what they have in common is the prior median of 1.0 for τ . Quantiles or mean of τ or θ_i for other scalings of $p(\tau)$ may be derived by proportional re-scaling (as in Table 3). For example, a half-Cauchy distribution that has its median heterogeneity matched to that of a half-normal distribution requires a scale parameter that is smaller by a factor of $\approx 2/3$. From the table, one can also read off the ratio of 95% quantile over the median, which may be a useful indicator of the heavy-tailedness of the different distribution families. The distributions from Table 4 are also illustrated in Figure 1. Some additional properties of these distributions are provided in Supporting Information Appendix B.

Different distributional families for the prior $p(\tau)$ imply differing marginal prior predictive distributions $p(\theta_i|\mu, \tau)$. Concrete prior information on $p(\theta_i|\mu, \tau)$ then may help constraining the shape of $p(\tau)$, however, the prior family may also be selected based on considerations of heavy-tailedness, near-zero behavior, or simplicity.

3.4.5 | The role of the unit information standard deviation

Consider the simple case of an effect measure that for each study is determined as an average of independent identically distributed observations. In such a case, the associated standard error is simply of the form

$$\sigma_i = \frac{\sigma_1}{\sqrt{n_i}}, \quad (4)$$

where n_i is the sample size, and σ_1 is the common “population” standard deviation of each single observation that was averaged over. This figure describes the *population*-, or *within-study-standard deviation*,⁵⁴ which for the moment we take to be constant across studies. This figure is also called the *unit information standard deviation (UISD)*, as it relates to an observational unit's contribution to a study's likelihood. One may now relate the heterogeneity τ to σ_1 and ask whether the between-study variability (τ) is likely to exceed the within-study variability (σ_1), or what ratios of these two are plausible. Figure 2 illustrates the relationship of within-study and between-study standard deviations σ_1 and τ . Usually, one would expect $\tau \ll \sigma_1$, implying that while study means (θ_i) may differ to some degree, the distributions of subjects within studies will still be largely overlapping (see Figure 2, left panel). In that sense, the UISD σ_1 may constitute an

TABLE 2 Categories of heterogeneity and corresponding τ ranges in the context of log-ORs, according to Spiegelhalter et al.,¹⁵ Sec. 5.7

Category	Range
“Reasonable”	$0.1 < \tau < 0.5$
“Fairly high”	$0.5 < \tau < 1.0$
“Fairly extreme”	$\tau > 1.0$

TABLE 3 Implications of a range of half-normal heterogeneity priors $p(\tau)$ on probable values of heterogeneity τ and predicted effects θ_i (*marginal prior predictive distributions*)

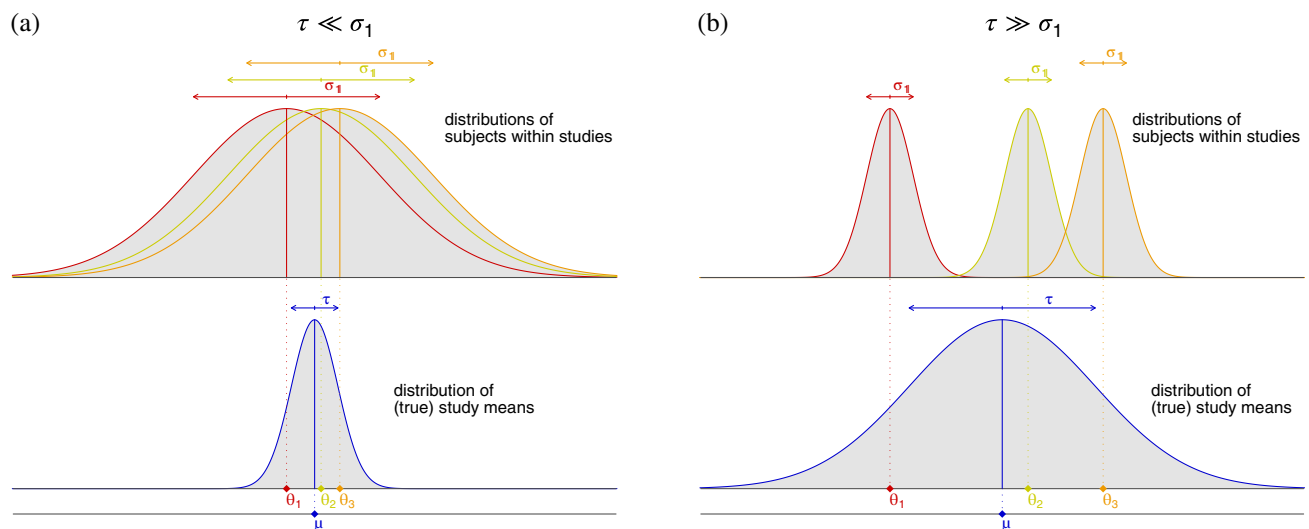
$p(\tau)$	Heterogeneity τ			95% Predictive interval		Category probability (%)		
	Median	Mean	95% quant.	$\theta_i - \mu$	$\exp(\theta_i - \mu)$	Reasonable	Fairly high	Fairly extreme
Half-normal(0.1)	0.07	0.08	0.20	[-0.22, 0.22]	[0.80, 1.24]	32	0	0
Half-normal(0.2)	0.13	0.16	0.39	[-0.44, 0.44]	[0.65, 1.55]	60	1	0
Half-normal(0.5)	0.34	0.40	0.98	[-1.09, 1.09]	[0.34, 2.98]	52	27	5
Half-normal(1.0)	0.67	0.80	1.96	[-2.18, 2.18]	[0.11, 8.89]	30	30	32
Half-normal(2.0)	1.35	1.60	3.92	[-4.37, 4.37]	[0.013, 79.0]	16	19	62

Note: The three rightmost columns show the corresponding probabilities for the three categories from Table 2.

TABLE 4 Implications of a range of heterogeneity priors $p(\tau)$ from different families on probable values of heterogeneity τ and predicted effects θ_i (*marginal prior predictive distributions*)

$p(\tau)$	Scale	Heterogeneity τ			95% Predictive interval	
		Median	Mean	95% quant.	$\theta_i - \mu$	$\exp(\theta_i - \mu)$
half-normal(1.48)	1.48	1.00	1.18	2.91	[-3.24, 3.24]	[0.039, 25.5]
half-Student- $t_{\nu=4}$ (1.35)	1.35	1.00	1.28	3.75	[-3.85, 3.85]	[0.021, 46.8]
half-Cauchy(1.00)	1.00	1.00		12.7	[-10.10, 10.10]	[0.000041, 24,371]
half-logistic(0.91)	0.91	1.00	1.26	3.33	[-3.55, 3.55]	[0.029, 34.7]
exponential(0.69)	1.44	1.00	1.44	4.32	[-4.33, 4.33]	[0.013, 75.9]
Lomax $_{\alpha=6}$ (8.17)	8.17	1.00	1.63	5.29	[-5.04, 5.04]	[0.0065, 155]
Lomax $_{\alpha=1}$ (1.00)	1.00	1.00		19.0	[-14.74, 14.74]	[0.00000040, 2,520,157]

Note: For comparability, the different priors are all scaled to a common median of 1.0. Except for the exponential distribution, which is commonly parameterized by its *rate* (or inverse scale), all distributions have a scale parameter.

**FIGURE 2** Illustration of the relationship of between-study heterogeneity τ and unit information standard deviation (UISD) σ_1 . The left panel (a) shows the commonly expected setup, in which the heterogeneity τ is relatively small compared to the within-study standard deviation ($\tau \ll \sigma_1$). The right panel (b) shows that a larger τ would imply that the distributions of subjects from different studies were eventually barely overlapping. Note that the eventual *estimates* (y_i) resulting from the different studies then may have different standard errors $\sigma_i = \frac{\sigma_1}{\sqrt{n_i}} < \sigma_1$ associated, depending on the studies' sample sizes n_i [Colour figure can be viewed at wileyonlinelibrary.com]

important “landmark” on the heterogeneity continuum and thus may help constraining the range of plausible heterogeneity values.²⁶

This concept of within-study standard deviation may be extended to other types of effect scales—for example, the standard error of a log-OR derived from a 2×2 -table is approximately given by $\sigma_i = \frac{4}{\sqrt{n_i}}$, so that, *heuristically*, the UISD here equals $\sigma_1 = 4$ per subject (at least),²⁰ Supporting Information Appendix A.1. Sometimes it may also make more sense to define UISDs not *per subject* but rather *per event* (see also Supporting Information Appendix A.3 for an example), but care also needs to be taken in order not to confuse these two figures. For a given set of log-OR estimates, the UISD may alternatively also be investigated by inverting Equation (4) (see also Equation⁶ and the examples in Section 5.3 below).

Another link may be drawn between σ_1 and τ via shrinkage estimation (see Section 2.1) and the consideration of *prior effective sample sizes*.^{52,77} Consider the case where a meta-analysis of k studies is available, and a new ($k + 1$)th study is conducted. The previous meta-analysis of course provides (prior) information on the new study's estimate θ_{k+1} , the exact amount of which is determined by the number of studies k , their sample sizes n_i , the UISD σ_1 , but also by the amount of heterogeneity.^{33,72} If τ is large, then separate studies are only loosely related and the previous data add little information. If on the other hand τ is very small (i.e., studies are almost homogeneous), then they may contribute a lot of information. With that, the amount of heterogeneity is related to whether studies should rather be pooled or viewed as essentially independent pieces of information. One may then consider the idealized limiting case of infinitely many ($k \rightarrow \infty$) infinitely large ($n_i \rightarrow \infty$) studies as the previous data source, so that the amount of contributed information solely depends on τ . In that case, the historical data may be thought of as effectively contributing a number of n_∞^* additional subjects to the $k + 1$ th study. This *prior maximum sample size* then relates to σ_1 and τ as.⁷⁷

$$\frac{\tau}{\sigma_1} = \frac{1}{\sqrt{n_\infty^*}}. \quad (5)$$

Table 5 illustrates this relationship. For example, if in the ideal case (i.e., $k = \infty$, $n_i = \infty$) the additional data should add information equivalent to *at most* 16 subjects, then this would correspond to τ amounting to *at most* a quarter of σ_1 . If one has an idea of how much information a meta-analysis may (or should) contribute to a single study's shrinkage estimate (in the idealized case of very many very large studies), then such considerations may help

TABLE 5 Correspondence between prior maximum sample sizes (n_∞^*) and the magnitude of the heterogeneity (τ) relative to the unit information standard deviation (UISD) (σ_1) (see (5))⁷⁷

τ/σ_1	0	1/16	1/8	1/4	1/2	1	∞
n_∞^*	∞	256	64	16	4	1	0

constraining probable magnitudes of τ , or associating probabilities with ranges of τ values.

Note that a number of priors have been proposed which are defined relative to the magnitude of the σ_i values (or their harmonic mean), for example, the *Jeffreys*, *DuMouchel* or *uniform shrinkage* priors.²⁰ In view of the above arguments, it might also make sense to define priors relative to the UISD, or its estimated value. Inverting (4) yields $\sigma_1 = \sqrt{n_i \sigma_i^2}$ for a single study, and based on a given data set we suggest the more general empirical estimate

$$s_1 = \sqrt{\bar{n} \bar{s}_h^2} = \sqrt{\frac{\sum_{i=1}^k n_i}{\sum_{i=1}^k \sigma_i^{-2}}} \quad (6)$$

where $\bar{n} = \frac{1}{k} \sum_{i=1}^k n_i$ is the average (arithmetic mean) sample size, and $\bar{s}_h^2 = \left(\frac{1}{k} \sum_{i=1}^k \sigma_i^{-2} \right)^{-1}$ is the harmonic mean of the squared standard errors (variances). This estimator is defined so that in the special case of a common-effect analysis (i.e., assuming $\tau = 0$), the overall mean estimate's variance (which then is given by $\left(\sum_{i=1}^k \sigma_i^{-2} \right)^{-1}$) consistently also equals $\frac{s_1^2}{\sum_i n_i}$.

3.4.6 | Empirical information on τ

Empirical data, for example, from earlier investigations in a related area,⁷⁸ may also contribute to a-priori information. Informative priors based on empirical information have been derived for standardized mean differences (SMDs) and log-ORs in medical applications by investigating large numbers of meta-analyses published in the *Cochrane Database of Systematic Reviews* by Rhodes et al.⁷⁴ and Turner et al.⁷⁵ Additional evidence for certain types of effect scales may be found, for example, in the works by Pullenayegum,³⁴ Turner et al.,⁷⁹ Kontopantelis et al.,⁸⁰ Steel et al.,⁸¹ van Erp et al.,⁸² Seide et al.,^{83,84} and Günhan et al.⁸⁵ Note that some references provide information directly on the heterogeneity *parameter*, while others summarize *estimates* of heterogeneity.

Empirical information often entails the question of how representative the external information is for the

study at hand, or what may be the relevant data subset, or what to do if no such sample may be available. In terms of the *epistemic* view discussed in Section 2.2.2, the inclusion of empirical evidence in the prior specification affects the interpretation of the prior, and with that, of the posterior. Empirical data may then often be seen as a somewhat complementary source of evidence. When there is doubt about the immediate applicability of empirical information for the problem at hand, this may also be reflected, for example, in a robustified two-component mixture prior.^{72,73}

3.5 | Guiding questions

In order to summarize the above arguments, Table 6 lists some guiding questions that may aid in structuring the specification of a prior for the heterogeneity. These are mostly based on the arguments laid out in Sections 3.3 and 3.4. Firstly, plausible heterogeneity magnitudes (in terms of τ or θ_i ranges) need to be determined. These reflections may then also help choosing a parametric family for the prior, or the distributional family may also be selected based on considerations of near-zero behavior, heavy-tailedness or simplicity. Beyond the mere type of endpoint or effect measure, the context also may determine whether smaller or larger amounts of heterogeneity

TABLE 6 Some guiding questions for judging reasonable prior distributions for the heterogeneity parameter τ

Prior information:	
(i)	What is the effect scale, what (<i>between-study</i>) differences are expected or plausible?
(ii)	What is the magnitude of other known (or plausible) effects? Do these provide guidance? Is an informative effect prior used? If so, what is its variance? Does it provide guidance?
(iii)	Is a plausible “unit information standard deviation (UISD)” available? Does it provide guidance?
(iv)	Is relevant external empirical information on heterogeneity available? Should it be considered in the analysis?
Translation into a prior probability distribution:	
(v)	Does the prior information help pinpointing prior quantiles (of τ)?
(vi)	Does the prior information help pinpointing prior predictive quantiles (of θ_i)?
(vii)	Does the prior information suggest particular properties for the prior (-density)? (Monotonicity? A non-zero mode? A heavy tail? Certain near-zero behavior? ...)

are to be expected, for example, depending on whether studies' designs and populations were similar. Special considerations in the context of specific common types of effect scales are discussed in detail in Section 4. These are then illustrated using actual data examples in Section 5.

4 | MOTIVATING HETEROGENEITY PRIORS IN VARIOUS SETTINGS

4.1 | Means and mean differences

This general case covers endpoints measured on *absolute* scales, hence it is not possible to give universally applicable advice on a plausible prior scale. For example, the same analysis may require different scalings of the prior depending on whether an endpoint is expressed, say, in terms of hours or minutes. In particular, in case of effects that are defined as averages, the UISD (see also Section 3.4.5) may provide some guidance; if standard errors σ_i scale with sample size ($\sigma_i \approx \frac{\sigma_1}{\sqrt{n_i}}$, see also Equation (4)), then σ_1 (or an estimate s_1 , (6)) may provide some orientation based on the considered (or other related) data. Relating effects to “within-population standard deviations” is actually an approach that is also formalized in the case of *standardized* mean differences (SMDs); see the following section.

Mean differences are another very common special case. These are often used in order to “normalize” outcomes; for example, in controlled clinical trials, each study's *treatment* group is usually related to a *control* group in order to express the treatment effect *relative to the unexposed group*. In the simplest case, the study's outcome then is defined as $y_i = \bar{x}_{2,i} - \bar{x}_{1,i}$, where $\bar{x}_{1,i}$ and $\bar{x}_{2,i}$ are the i th study's averages from control and treatment group, respectively. When considering UISDs, the relevant sample size will then result as the sum of the two treatment groups' sizes ($n_i = n_{1,i} + n_{2,i}$). In the simple case of two equally-sized groups ($n_{1,i} = n_{2,i} = \frac{n_i}{2}$) and equal variances within groups (so that $\text{Var}(\bar{x}_{1,i}) = \text{Var}(\bar{x}_{2,i}) = \frac{\sigma_w^2}{n_i/2}$) the UISD simply results as $\sigma_1 = \sqrt{2\sigma_w^2}$, where σ_w^2 is the within-group variance.

Again a special case arises when considering *paired differences*.⁸⁶ In general, analogous considerations apply for un-paired as well as for paired differences; only for the latter case the UISD σ_1 may be expressed as $\sigma_1^2 = \text{Var}(x_{1;j}) + \text{Var}(x_{2;j}) - 2\text{Cov}(x_{1;j}, x_{2;j})$ where j is the index identifying the j th pair of observations in the i th study. We can see how the individual (paired) observation's variance contribution results as a sum of the two observations' marginal variances and their covariance. Now, since any pair of observations ($y_{1;j}$ and $y_{2;j}$) is

usually positively correlated ($\text{Cov}(y_{1,ij}, y_{2,ij}) > 0$), the sum of individual variances ($\text{Var}(x_{1,ij}) + \text{Var}(x_{2,ij})$), if known, may provide an upper bound on σ_1 .

Finally, there are generic cases of parameter estimates that are reported along with a standard error, but which do not necessarily have a “sample size” (n_i) associated, as is sometimes the case, for example, for laboratory experiments.⁸⁷

4.2 | Standardized mean differences

Standardized mean differences (SMDs) aim to compare mean differences measured on different scales by normalizing them through their population standard deviation. Effectively, these measure *by how many standard deviations* the two study groups differ; SMDs are always dimensionless. Their aim is to estimate $\delta_i = \frac{\mu_{2,i} - \mu_{1,i}}{\varsigma_i}$, where $\mu_{2,i}$ and $\mu_{1,i}$ are the two groups' true means and ς_i is the within-group standard deviation (which may be defined with respect to one or the other or both treatment groups, or which may also be externally informed). Note that ς_i here bears some similarity to the UISD σ_1 (when considering the latter with respect to the *unstandardized* differences). Slightly differing, but essentially similar approaches are given, for example, by the “Cohen's d ,” “Hedges' g ” or “Glass' Δ ” estimators, which differ in details like bias correction or standardization terms.^{3,4} Essentially, these aim to estimate the mean difference ($\mu_{2,i} - \mu_{1,i}$) by the difference of averages ($\bar{x}_{2,i} - \bar{x}_{1,i}$), and also the standard deviation by an empirical one. SMDs (along with the correlations treated below) are somewhat different here from the “general” mean differences, in that they are explicitly designed and utilized in order to compare endpoints measured on different scales, which are not *directly* comparable. A heterogeneity of $\tau = 0$ may hence be considered particularly unlikely. A value of $\tau = 1$ would mean that the between-study heterogeneity (among δ_i values) was equal to the within-group variability ς_i . Closely related to SMDs are *standardized regression coefficients*, which are re-scaled as if both the regressor's as well as the response's variance were normalized to unity.⁸⁸ Similar arguments would apply for analyses involving standardized regression coefficients, and arguments applicable to correlation coefficients (see Section 4.5 below) may also be relevant.

Effects on the SMD scale have been categorized as 0.2 = “small,” 0.5 = “medium,” 0.8 = “large,”⁸⁹ Sec. 2.2.3, where an extension has recently been proposed to include the grades of 0.1 = “very small,” 1.2 = “very large,” and 2.0 = “huge.”⁹⁰ Consequently, such a ranking might be utilized in order to bound between-study effects to mostly non-extreme values, for example, by

anticipating mostly up to “large” heterogeneity and hence formulating a bound on $P(\tau \leq 1)$. Neglecting estimation uncertainty for the denominator, and for simplicity assuming equal sample sizes for each of the i th study's groups, leads to a UISD of $\sigma_1 = 2$ (see Supporting Information Appendix A.1).

Empirical evidence on heterogeneities between SMDs based on an analysis of studies archived in the *Cochrane Database of Systematic Reviews* is given by Rhodes et al.⁷⁴; for a general healthcare setting (not restricted to a particular outcome type), a log-Student- t distribution with parameters $\mu = -1.72$, $\sigma = 1.295$, and 5 degrees of freedom was derived (implying a median and 95% quantile of 0.18 and 2.43, respectively). Heterogeneity *estimates* reported in studies published in the *Psychological Bulletin* are provided by van Erp et al.⁸²; the 189 τ -estimates for SMDs that were quoted in 32 publications had a median and 95% quantile of 0.20 and 0.66, respectively.

4.3 | Log-transformed odds, rates and effect scales

Many outcomes are commonly analyzed on a logarithmic scale, which may be advantageous for several reasons; firstly, the domain of positive numbers is mapped to the complete real line, which makes strictly positive scales tractable for normal models like the NNHM, which is often convenient. Secondly, additive effects on the log-scale translate to multiplicative effects on the original scale. Symmetry of the normal distribution (2) on the log-scale then implies a “symmetric” treatment of multiplicative factors and their inverses (since $\exp(\mu + x) = \exp(\mu) \times \exp(x)$ while $\exp(\mu - x) = \exp(\mu) \times \frac{1}{\exp(x)}$). This is useful, for example, when dealing with outcomes like rates, odds, rate ratios, odds ratios, relative risks, hazard ratios or concentration measurements. An offset of, say, 0.1 on the log-scale translates (approximately) to a change of 10% on the back-transformed (exponentiated) scale, regardless of the original value. Thirdly, the normal approximation to the likelihood that is used in the NNHM (1) may provide a better fit on the logarithmic scale.

When considering heterogeneity values on the logarithmic scale, a more intuitive approach is usually to examine the corresponding implications on the back-transformed scale. Note that a normal model on the log-scale actually corresponds to a log-normal model on the original scale. In a sense, an analysis on the logarithmic scale may also be viewed as an implementation of a dependent joint prior for effect and heterogeneity^{22,34} on the original (exponentiated) scale. The consequences of certain heterogeneity values or heterogeneity distributions were already investigated in some detail in

Sections 3.4.3 and 3.4.4; the important issue to judge is what *relative* (multiplicative) difference between studies is deemed plausible; see also the extensive discussion by Spiegelhalter et al. (Sec. 5.7).¹⁵

A common type of effect are log-transformed odds (or *logits*).^{91,92} For example, in epidemiology or at the design stage of a clinical trial it may be of interest to infer the magnitude and variability of the prevalence of a certain condition, or historical information may be utilized to support the control group in a clinical trial.⁷² The prevalence may be expressed in terms of the probability $p \in [0, 1]$ or the odds $\frac{p}{1-p} \in [0, \infty]$, while for meta-analysis purposes it then makes sense to move to the log-odds scale $\log\left(\frac{p}{1-p}\right) \in \mathbb{R}$. Rather than viewing this as a case of a logarithmic transformation of the odds, one might as well consider this as a *logit* transformation of probabilities, mapping the interval $[0, 1]$ to the real line via the *logit* function $f(p) = \log\left(\frac{p}{1-p}\right)$. Besides considerations of what ratios the odds may plausibly be spanning, here it may be helpful to consider a uniform distribution in proportions as an extreme case; for the log-odds, this implies a logistic distribution that has a standard deviation of $\frac{\pi}{\sqrt{3}} = 1.81$. The USD in this case amounts to (at least) $\sigma_1 = 2$ (see Supporting Information Appendix A.2). Similarly, event rates (based on a Poisson model) are commonly combined in meta-analyses based on a log-transformation.

Similarly to the cases of means and mean differences discussed earlier, a log-transform is also commonly applied in the context of two-group comparisons, for example, for log-OR, log-IRR, log-RR or log-HR effect measures. Logarithmic ORs are a natural extension of the log-odds case above, since the logarithmic *ratio* of odds is simply a *difference* of log-odds; other pairwise group comparisons generalize similarly from single-group estimates. USDs for log-ORs and log-RRs are derived by Röver,²⁰ and for log-IRRs in Supporting Information Appendix A.3; the corresponding figures for log-HRs are discussed by Spiegelhalter et al. (Sec. 2.4.2).¹⁵ When discussing USDs for count outcomes, it is important to clearly indicate whether these relate to *subjects* or *events* (e.g., for ORs the numbers are four per subject²⁰ and two per event¹⁵).

Empirical evidence on the magnitude of heterogeneities within meta-analyses published in the *Cochrane Database of Systematic Reviews* is given by Turner et al.^{75,79} For example, for a log-OR effect in a general healthcare setting (without restricting to a specific type of outcome), a log-normal distribution with $\mu = -1.28$ and $\sigma = 0.87$ was derived, implying a median and 95% quantile of 0.28 and 1.16, respectively (see also Table 3). Similarly, Günhan et al.⁸⁵ in a re-analysis of data from the Cochrane Database of Systematic Reviews determined a 95% quantile of heterogeneity *estimates* of 1.05 for analyses based on binary data and log-ORs.

Consider for example the common case of a meta-analysis of log-OR estimates. If we want to restrict prior probabilities mostly to “reasonable” to “fairly high” heterogeneity levels (according to Table 2 in Section 3.4.3), one could use a half-normal prior with scale 0.5, implying $P(\tau > 1.0) = 4.6\%$ and assigning 52% and 27% probability to the “reasonable” and “fairly high” categories, respectively. Figure 3 illustrates the half-normal(0.5) prior along a half-normal(1.0) prior, and the prior proposed by Turner et al.⁷⁵ (log-normal with $\mu = -1.28$ and $\sigma = 0.87$). The heterogeneity categories from Table 2 are marked, and at the bottom, the probabilities for the categories are shown. The probabilities assigned by the half-normal(0.5) prior and the “empirical” prior are roughly in agreement, while the half-normal(1.0) prior would assign more or less equal probabilities to the “reasonable,” “fairly high” and “fairly extreme” categories, and leave only 8% probability for smaller values. Similar arguments hold also for other log-transformed effect scales.

4.4 | Regression slopes

Very closely related to mean differences is the more general case of meta-analysis of regression parameters (slopes or interactions) and their standard errors.⁹³ In the special case of a single binary covariate, the regression effectively reduces to a two-group comparison, and consideration of additional covariates then may allow for some “adjustment.” When the covariate is continuous, however, extra care needs to be taken, since not only the endpoint’s scaling is relevant (the regression’s “y variable”), but also the regressor’s scaling (the regression’s “x variable”). Whether the regressor is expressed in, say, days or weeks, affects the resulting slope parameter (and its standard error) by a corresponding re-scaling by a factor of seven. The regressor’s scaling will then similarly also affect the scale of the anticipated heterogeneity: when combining estimated (linear) regression coefficients, which are to be interpreted as “the expected change in y for a one-unit change in x,” the heterogeneity between estimates depends on the units of x. For example, the variability expected among temporal changes that are expressed on a *per-week* scale rather than a *per-day* scale should be seven times as large.

The immediate question then is what increment in the regressor to base heterogeneity considerations on; what is eventually needed is a statement of the form “for a change in the regressor by a difference of Δ_x , the associated effects are anticipated to vary by a magnitude of τ ,” and that difference Δ_x needs to be specified. Sometimes there may be obvious “natural” units to be used, for example in the common case of a binary (zero/one) coded covariate (e.g., for treatment vs. control or males vs. females); the




FIGURE 3 Comparison of the heterogeneity prior proposed by Turner et al.⁷⁵ for log-ORs in a general setting (a log-normal distribution with $\mu = -1.28$ and $\sigma = 0.87$, shown in blue) with half-normal priors (with scales 0.5 and 1.0). The bottom plots especially contrast the implied prior probabilities for the heterogeneity categories proposed by Spiegelhalter et al. (Sec. 5.7)¹⁵ (see also Tables 2 and 3) [Colour figure can be viewed at wileyonlinelibrary.com]

obvious difference to consider here is an increment of $\Delta_x = 1$. Otherwise the width of the regressor's distribution may be relevant.⁹⁴ Consider again the case of a binary covariate and a balanced setup; the standard deviation of the binary variable will then be $\frac{1}{2}$, so that *twice the standard deviation* might generally be a sensible scale to consider. Note though that this is by no means universally applicable, as such scales may be affected by many factors (e.g., inclusion criteria in clinical trials) and might also be very different between studies. Note that the Δ_x value needs to be the same across the considered studies.

Once the "reference" increment Δ_x has been determined, a prior for the associated heterogeneity may be formulated. In case the actual analysis then is done with respect to a differing scaling, the prior needs to be re-scaled accordingly. For example, if a prior with scale s was determined for a *per-week* increment, but the actual analysis is based on the *per-day* regression coefficients, then their prior should have scale $\frac{s}{7}$. The UISD σ_1 then also scales proportionally.

Note that the above arguments extend beyond simple linear regressions with continuous outcomes, for example,

logistic regressions, Poisson regressions or survival analyses, in which regression parameters then relate to log-ORs, log-IRRs or log-HRs. Once a reference increment Δ_x has been determined, the arguments regarding log-transformed endpoints discussed earlier in Section 4.3 apply, and potential re-scaling issues still need to be considered. A way to circumvent considerations of regressor's or response's scales may be to move to *standardized regression coefficients* instead, which are unitless and are somewhat similar to SMDs (see also Section 4.2) or correlations (see Section 4.5).⁸⁸ Depending on the exact type of regression analysis and the standardization technique (e.g., in case of a logistic regression, and when standardization is done based only on the regressor's scale),⁹⁵⁻⁹⁷ arguments relevant for log-transformed endpoints might also apply.

4.5 | Correlation coefficients

Estimated correlation coefficients (Pearson's r) are commonly quoted and summarized for studies dealing with paired observations.^{3,4,98} Correlation coefficients are

restricted to the domain $[-1, 1]$, with values of $|r| = 1$ indicating perfectly linear (positive or negative) correlation, and $r = 0$ indicating uncorrelatedness.⁹⁹ Due to the problems with bounded parameter spaces, correlation coefficients are commonly analyzed after an appropriate transformation using *Fisher's z transform*, which is defined as $z_i = \frac{1}{2} \log\left(\frac{1+r_i}{1-r_i}\right) = \text{arctanh}(r_i)$. This transformation maps the original domain to the real line, and in particular, it is also a *variance stabilizing transformation*; the (approximate) standard error of the transformed z_i value only depends on the i th study's sample size n_i and is given by $\frac{1}{\sqrt{n_i-3}}$. Correlation values within the range $-0.5 < r_i < 0.5$ are little affected by the transformation, which makes more of a difference for more extreme values.

An upper limit to the expected heterogeneity may be specified by considering a uniform distribution of θ_i values across the range of correlation coefficients as a “worst case.” For plain (correlation r) values, this would imply a variance of $\frac{1}{3} = 0.58^2$. On the scale of z -transformed values, this implies a distribution with probability density function $p(z) = \frac{2}{(\exp(-z) + \exp(z))^2}$, that has a zero mean and a variance of $\frac{\pi^2}{12} \approx 0.91^2$ (these moments might actually motivate a prior for the overall effect μ , too). The standard error of z_i values after transformation (see above) implies a UISD of approximately $\sigma_1 = 1.0$. With that, it should usually be safe to expect heterogeneity values well below $\tau = 1.0$.

If τ values near unity (or 0.91) already imply rather extreme heterogeneity, the question remains what constitutes “large,” yet reasonable heterogeneity. For that, we may consider the somewhat more moderate cases of $r \sim \text{Uniform}(-0.5, 0.5)$ or $r \sim \text{Uniform}(0.0, 0.8)$. Both these cases happen to lead to similar variances of $\text{Var}(z) = 0.30^2$ on the transformed scale, so that $\tau = 0.30$ may already be considered “large” heterogeneity.

While the use of “plain,” un-transformed correlation values within the NNHM framework is a bit problematic due to the bounded parameter space that is not reflected in the model, it is not uncommon. We have already seen some hints of what amounts of between-study variance for plain correlations may be possible or plausible in the considerations above; a value of $\tau = \frac{1}{\sqrt{3}} = 0.58$ (corresponding to a uniform distribution in r) would already be extreme; one would most likely expect values way less than even half as much.

Van Erp et al.⁸² collected heterogeneity estimates reported in studies that were published in the *Psychological Bulletin*. Although the figures were not identified as being based on Fisher- z transformation or not (apparently a mix of both was encountered), these numbers may provide some empirical motivation. Among the

observed heterogeneity estimates for correlation endpoints in 539 analyses from 25 studies, a median and 95% quantile of 0.12 and 0.29, respectively, were found. Similarly, Steel et al.⁸¹ quote heterogeneity estimates from 292 management-related meta-analyses in the range of 0.0–0.4, with a median of 0.16.

5 | EXAMPLE APPLICATIONS

5.1 | Mean differences

Grande et al.¹⁰⁰ investigated the effect of physical exercise (vs. no exercise as control) on the duration of acute respiratory infections (ARIs). Four studies were jointly considered in a meta-analysis, the endpoint of interest was the mean difference in the *number of symptom days per episode*. The relevant data are shown in Table 7.

The outcome here is measured in units of *days* (change in symptom duration for treated patients relative to the control group). For the purpose of the present analysis, ARIs were defined as “infections of the respiratory tract that last for less than 30 days,”¹⁰⁰ while ARI durations generally are substantially shorter, lasting of the order of a week.^{101,102} With that, the reduction in symptom days cannot be more than (roughly) a week. ARIs may be caused by bacterial or viral pathogens; the effect of antibiotic treatment is in a shortening of the order of 1 day.¹⁰³ From the data (Table 7), we can derive estimates of the UISD, which here is at an average of $s_1 = 3.9$.

The treatment effect may be expected to be of the order of days (anything below 1 day would probably not be considered clinically meaningful), and a similar magnitude may be expected for the heterogeneity. Values $\tau > 1$ would make the between-study heterogeneity larger than the effect of antibiotics, which seems implausible. Variations in treatment effects of the order of several days would probably imply that the effect was several times larger in some studies than in others.

A τ value of 1.0 would imply a median difference in true effects of ≈ 1 day for a random pair of studies (see Table 1), which might be at the upper end of the plausible range. A half-normal(0.5) prior would imply $P(\tau \leq 1) \approx 95\%$, and considering the corresponding prior predictive distribution (see Table 3), we can see that this implies a 95% prior predictive interval of roughly ± 1 day around the overall mean effect.

For the present example, we would hence suggest a half-normal(0.5) prior. Note that this is a common, well-researched condition. For more uncertain cases, one might want to go for a heavier-tailed prior. A meta-analysis based on the half-normal(0.5) prior is illustrated in Figure 4. Among the four studies considered, one

suggests a stronger effect than the others, however, due to its relatively small size and correspondingly large associated standard error, it is still consistent with the remaining three. The estimated heterogeneity (the median and 95% credible interval [CI] are shown in the bottom left of the forest plot) here has barely changed from the a priori anticipated amount (see Table 3). The heterogeneity's posterior is also illustrated in Figure 8; prior and posterior are very similar in this case. The resulting combined estimate then also suggests a more moderate effect, namely, a reduction of the order of one symptom day, with an uncertainty of about a factor of two. The estimated heterogeneity is relatively low compared to the width of the overall mean's CI, and so the prediction interval is only slightly longer, and the shrinkage intervals show substantially greater precision than the original estimates. Sensitivity to other prior choices is also investigated for this example in Supporting Information Appendix D.4.

5.2 | Standardized mean differences

Aalbers et al.¹⁰⁴ investigated the short-term effect of music therapy on depression symptoms; four studies

comparing music therapy plus treatment-as-usual (TAU) versus TAU alone were found. Within these four studies, differing clinician-rated symptom scores were utilized in order to quantify depression severity: the Hamilton rating scale for depression (HAM-D), considering potentially differing numbers of items between studies, as well as the Montgomery-Åsberg depression rating scale (MADRS). In order to facilitate a joint analysis, the meta-analysis was based on SMDs (here: Hedges' g); the relevant data are shown in Table 8.

The outcome measured on the SMD scale means that a unit change in y_i corresponds to a *one standard deviation change* in the symptom severity score. Considering, for example, the Alborno (1992) study,¹⁰⁵ which was measuring change in symptom severity using the 17-item HAM-D scale with a within-group standard deviation of about 5 (see Table 8), a difference of 1 on the SMD scale here would roughly correspond to a five-point change in HAM-D score.^{106–109} In terms of SMD, this would already be considered a “large” effect.^{89,90} The UISD for SMDs is predicted at $\sigma_1 = 2$, while from the present data here we get a very similar empirical average of $s_1 = 2.2$.

For the between-study differences, we would assume that they would be mostly in the “small” to “medium” range ($\ll 1$)—otherwise effects would be differing by a

<i>i</i>	Study	Treatment group			Control group			MD	
		$\bar{x}_{1,i}$	$s_{1,i}$	$n_{1,i}$	$\bar{x}_{2,i}$	$s_{2,i}$	$n_{2,i}$	y_i	σ_i
1	Nieman (1990)	3.60	2.97	18	7.00	5.94	18	−3.40	1.57
2	Çiloğlu (2005)	5.15	1.56	60	6.10	1.00	30	−0.95	0.27
3	Barrett (2012)	9.30	5.13	47	11.40	5.75	51	−2.10	1.10
4	Sloan (2013)	5.30	1.50	16	6.30	2.20	16	−1.00	0.67

TABLE 7 Mean difference (MD) example data due to Grande et al.¹⁰⁰ \bar{x} , s and n denote the treatment and control groups' empirical means, standard deviations and sample sizes

Note: The y_i are the derived MDs and σ_i the associated standard errors that eventually go into the analysis (see Section 2.1). Here, mean differences are on the scale of days (change in disease duration). Negative estimates y_i indicate a beneficial effect.




FIGURE 4 Forest plots for the two examples discussed in Sections 5.1 and 5.2. In both cases, a half-normal(0.5) prior for the heterogeneity τ was used. Besides the intervals based on the quoted estimates, the shrinkage intervals are shown in gray. At the bottom, the credible interval for the overall mean (μ) is shown along with the prediction interval for a “new” additional study effect θ_{k+1} . The estimated heterogeneity (τ) is quoted in terms of the posterior median and shortest 95% credible interval

TABLE 8 Standardized mean difference (SMD) example data due to Aalbers et al.¹⁰⁴ \bar{x} , s and n denote the treatment and control groups' empirical means, standard deviations and sample sizes

<i>i</i>	Study	Treatment group			Control group			SMD	
		$\bar{x}_{1;i}$	$s_{1;i}$	$n_{1;i}$	$\bar{x}_{2;i}$	$s_{2;i}$	$n_{2;i}$	y_i	σ_i
1	Chen (1992)	−98.23	15.19	34	−67.06	15.19	34	−2.03	0.30
2	Radulovic (1996)	−16.50	10.00	30	−10.60	10.00	30	−0.58	0.26
3	Albornoz (2011)	−8.17	5.89	12	−3.83	5.31	12	−0.75	0.42
4	Erkkilä (2011)	−10.70	8.40	30	−6.05	8.06	37	−0.56	0.25

Note: The y_i are the derived SMDs and σ_i the associated standard errors that eventually go into the analysis (see Section 2.1). The original data are based on different depression symptom scores that are measured on different scales. Negative estimates y_i indicate a reduction in symptom severity.

standard deviation or more between studies, and also the studies' confidence intervals (which are roughly of the size $\sigma_i \approx \frac{\sigma_1}{\sqrt{n_i}} = \frac{2}{\sqrt{n_i}}$) would be unlikely to have any overlap. Rhodes et al.⁷⁴ in their empirical investigation based on the *Cochrane Database of Systematic Reviews* predicted a median and 95% quantile of 0.18 and 2.43 for the heterogeneity τ (where the large upper quantile appears rather extreme, based on the above arguments). Similarly, van Erp et al.⁸² inferred a median and 95% quantile of 0.20 and 0.66, respectively, based on heterogeneity *estimates* within a smaller data base.

A value of $\tau = 1.0$ would imply a median difference of ≈ 0.95 ("large") for a random pair of true study means θ_i (see Table 1), which already appears like a rather extreme amount; values of $\tau = 0.5$ (implying mostly "medium" sized between-study differences) or below seem to be more plausible. A half-normal(0.5) prior would cover this range and would imply a prior median (for τ) slightly above the magnitude suggested the empirical investigations (see also Table 3).

For the present example, we would then suggest a half-normal(0.5) prior as a slightly conservative choice, in order to reflect the potential heavy-tailedness suggested by Rhodes et al.,⁷⁴ and to account for the fact that the empirical data might be of limited relevance for the present example data. A meta-analysis based on the half-normal(0.5) prior is illustrated in Figure 4. Among the four studies, three consistently indicate estimates in the range 0.5–0.8, while the first one shows a huge effect estimate of the order of 2.0; a positive amount of heterogeneity appears to be present (the CI for τ is in a strictly positive range; see also Figure 8), and the eventual combined estimate indicates a "small" to "very large" average effect. Given the pronounced heterogeneity one might discuss whether the estimation of a pooled effect is meaningful. Nevertheless, we use this example to illustrate the use of Bayesian methods in heterogeneous situations, where heterogeneity cannot be explained and good reasons are available to perform a quantitative meta-analysis despite of large heterogeneity. The large estimated heterogeneity here results in a wide CI for the overall effect,

a very wide prediction interval, and also very little shrinkage for the estimated study-specific effects θ_i .

5.3 | Log-transformed effect scales

5.3.1 | Log odds ratio

A systematic review was performed by Crins et al.¹¹⁰ to investigate the effect of Interleukin-2 receptor antagonists (IL2-RA) on recovery of pediatric patients following liver transplantation. One aspect of interest was the occurrence of *acute rejection* (AR) reactions as a common adverse event. Two randomized controlled trials reporting such data were found, the event counts along with the corresponding (logarithmic) odds ratios and standard errors are shown in Table 9. Both studies indicated a reduction in the chances of an AR event for the treatment group.

The treatment effect is expressed and analyzed on a logarithmic scale here. A heterogeneity magnitude of $\tau = 1.0$ would imply that any random pair of studies would be expected to exhibit effects differing by a factor of 2.6 (see Table 1), which seems quite extreme already; values like $\tau = 0.5$ or below seem more plausible. In a similar investigation involving 14 studies and based on adult patients (Goralczyk et al.¹¹¹), a mean treatment effect (log-OR) of −0.26, corresponding to an OR of 0.77, was found. The UISD for a log-OR is at $\sigma_1 \approx 4$ per subject, while for the present data here we get an estimate of $s_1 = 5.4$. An empirical study based on a large number of meta-analyses predicts a median (95% quantile) of 0.28 (1.16) for the heterogeneity (Turner et al.⁷⁵), and an investigation of heterogeneity *estimates* found a median (95% quantile) of 0.00 (1.05) (Günhan et al.,⁸⁵). In the data from the closely related meta-analysis by Goralczyk et al.¹¹¹ the heterogeneity is estimated at 0.12 (0.38).

A half-normal(0.5) prior would mostly cover values $\tau < 1.0$ (up to "fairly high" heterogeneity according to Table 2) with an expectation and median below 0.5 (see also Table 3). The resulting 95% prior predictive interval

would still include effects within a factor of 3 around the overall mean log-OR μ . For the present investigation, we would then suggest a half-normal(0.5) prior as a reasonably conservative choice, which also agrees roughly with the empirical evidence (see Figure 3). A meta-analysis based on this prior is shown in Figure 5. In this example we have two studies only, demonstrating the somewhat speculative nature of inferring heterogeneity based on sparse data, and highlighting the value of considering a-priori probabilities. In the present case, the two studies involved are not very large, and their resulting CIs are overlapping, which makes the data consistent with a wide range of heterogeneity values, from homogeneity ($\tau = 0$) up to magnitudes of $\tau = 10$ or $\tau = 20$. Including the weakly informative heterogeneity prior, and effectively down-weighting unreasonably large heterogeneity values, then leads to an estimate of -1.81 for the log-OR, corresponding to a reduction in the odds of an AR event down to $\exp(-1.81) = 16\%$. While the uncertainty still is large (ranging roughly from 5% up to 50%), the analysis clearly indicates a substantial reduction in AR events here. The heterogeneity's posterior density is also shown in Figure 8; here we can see that for the present example constellation, the posterior is very similar to the prior. With the very uncertain original estimates (due to the small sample sizes), the overall mean's CI is wide, but the additional width of the prediction interval is limited due to the (prior and empirical) information on the heterogeneity, and a noticeable shrinkage effect is also observable.

TABLE 9 Log-OR example data¹¹⁰

<i>i</i>	Study	Treatment group		Control group		log-OR	
		Events (<i>a_i</i>)	Total (<i>n_{1,i}</i>)	Events (<i>c_i</i>)	Total (<i>n_{2,i}</i>)	<i>y_i</i>	σ_i
1	Heffron (2003)	14	61	15	20	−2.31	0.60
2	Spada (2006)	4	36	11	36	−1.26	0.64

Note: *a* and *n₁* as well as *c* and *n₂* denote the event counts and total numbers of patients in treatment and control groups, which together summarize the trial outcome in terms of a 2 × 2 table. The *y_i* are the derived logarithmic odds ratios and σ_i are the associated standard errors that eventually go into the analysis (see Section 2.1). Negative values here indicate a reduction of the event odds, that is, a beneficial treatment effect.




FIGURE 5 Forest plots for the two examples discussed in Sections 5.3.1 and 5.3.2. In both cases, a half-normal(0.5) prior for the heterogeneity τ was used

5.3.2 | Log incidence rate ratio

Four studies investigating the effect of ferric carboxymaltose versus placebo in heart-failure patients with iron deficiency were jointly analyzed by Anker et al.¹¹² The main outcome was the *incidence rate ratio* (IRR) with respect to the composite endpoint of recurrent cardiovascular (CV) hospitalizations or CV death. The relevant available data are shown in Table 10. The eventual analysis is based on the logarithmic ratio of the event rates (per 100 patient-years of follow-up) of treatment over placebo group.

As in the previous example, the outcome is analyzed on the logarithmic scale, so that many arguments apply essentially analogously here. Regarding empirical evidence on previously encountered amounts of heterogeneity, there are no studies available that would be directly applicable for log-IRRs, however, odds ratios and rate ratios have quite some similarity, so that these findings also have some bearing here. The UISD here is at $\sigma_1 = 2$ *per event* (see Supporting Information Appendix A.3); with a total of 114 events observed among a total of 839 patients¹¹² (a rate of ≈ 0.14 events *per patient*), this would correspond to $\sigma_1 \approx \frac{2}{\sqrt{0.14}} = 5.4$.

per patient. For the present data, we empirically get an average of $s_1 = 6.6$.

For this example, we would again suggest a half-normal(0.5) prior. A meta-analysis based on this prior is shown in Figure 5. While the data *look* homogeneous (all intervals have some overlap, also because some studies

TABLE 10 Log-IRR example data¹¹²

<i>i</i>	Study	Rate ratio [95% CI]	<i>n_i</i>	log-IRR	
				<i>y_i</i>	<i>σ_i</i>
1	FAIR-HF (2009)	0.44 [0.22, 0.90]	459	−0.82	0.36
2	CONFIRM-HF (2015)	0.68 [0.38, 1.21]	301	−0.39	0.30
3	EFFICACY-HF (2015)	1.09 [0.21, 5.54]	34	0.09	0.83
4	FER-CARS-01 (2018)	0.87 [0.07, 10.4]	45	−0.14	1.28

Note: The incidence rate ratios for the composite endpoint of recurrent cardiovascular (CV) hospitalizations and CV mortality are given for each study. For the analysis, the logarithmic rate ratio is considered. Negative values here indicate a reduction of incidence rates, that is, a beneficial treatment effect.

TABLE 11 Log-odds example data due to Neuenschwander et al.⁷⁷

<i>i</i>	Study	Remission		Proportion $p_i = \frac{x_i}{n_i}$	Odds $\frac{x_i}{n_i - x_i} = \frac{p_i}{1 - p_i}$	log-odds	
		Events (<i>x_i</i>)	Total (<i>n_i</i>)			<i>y_i</i>	<i>σ_i</i>
1	Feagan (2005)	9	63	0.143	0.167	−1.79	0.36
2	Rutgeerts (2005a)	18	121	0.149	0.175	−1.74	0.26
3	Rutgeerts (2005b)	7	123	0.057	0.060	−2.81	0.39
4	Van Assche (2006)	6	56	0.107	0.120	−2.12	0.43

Note: The *n_i* and *x_i* here denote total numbers and the numbers of remitting patients among these. Analysis is done based on the derived log-odds *y_i* and their standard errors *σ_i*.

are very small and intervals are correspondingly wide), we would still anticipate the possibility of heterogeneity—since from experience we know that heterogeneity is frequently present, and because we know that heterogeneous circumstances are still likely to produce data that may still “look homogeneous.”⁷ Compared to our a-priori expectations of τ values up to 0.98 (see Table 3), the posterior then suggests a slightly lower heterogeneity range of up to 0.75, but the data do not provide very much evidence in this regard (see also the posterior in Figure 8). The mean treatment effect eventually is at a log-IRR of −0.49, corresponding to an IRR of 61% (i.e., a reduction in the event rate), with a CI ranging from 33% up to 116%. For these somewhat homogeneous estimates, one can see that the ones with very large associated standard errors eventually have shrinkage estimates close to the overall prediction interval. A sensitivity analysis investigating alternative prior choices for this example is also shown in Supporting Information Appendix D.4.

5.3.3 | Log odds

Neuenschwander et al. investigated the use of historical data in order to inform the analysis of a new data set.⁷⁷ A meta-analysis of several trials in ulcerative colitis was performed in order to support the analysis of a subsequent phase II trial. The figure of interest here was the

probability for *clinical remission at week 8* in placebo-treated patients, and the main interest was in a prediction for the new study's event probability, to then formally integrate this in a subsequent analysis using a meta-analytic-predictive (MAP) approach.⁷² Four previous randomized controlled trials reporting this endpoint were available, their data are shown in Table 11. Instead of working directly on the estimated probabilities *p*, the analysis here is done based on the odds $\frac{p}{1-p}$, and a subsequent log-transformation.⁹²

Homogeneity of placebo rates is not expected—differences between control rates are among the main reasons for requiring a control arm for each RCT, and for pursuing a contrast-based analysis.^{113,114} The studies were designed aiming for an estimate of the treatment effect, and the placebo rate originally was mostly a nuisance parameter here. However, *some* amount of similarity still is anticipated, and the aim of this exercise is to carefully derive the predictive distribution, which of course depends on the amount of heterogeneity τ .

The earliest of the four studies was planned anticipating a remission rate of 10% for the placebo group,¹¹⁵ and hence a USD of $\sigma_1 \approx \sqrt{\frac{1}{0.1} + \frac{1}{0.9}} = 3.33$ may be expected. Empirically, we get an estimate of $s_1 = 3.2$ from the present data set.

As the endpoint are *logarithmic* odds, we may again apply similar reasoning as in the previous subsections, regarding the anticipated ratios of odds. However, a

major difference here is that while clinical trials are usually carefully designed to provide reliable estimates of treatment effects (treatment/control contrasts), this is not necessarily the case for the event rates that we are considering here; we may expect the log-odds to be more variable than the log-ORs. With this in mind, and considering conservatism and robustness particularly desirable in the present context, we would suggest a half-normal(1.0) prior here. From Table 3, we can see that the implied 95% prior predictive interval then spans a range of roughly a factor 9 around the median μ . Given the context, it may be of particular interest to consider the associated prior maximum sample size n_{∞}^* (see Section 3.4.5); for the prior median of $\tau = 0.67$, we have $\frac{\tau}{\sigma_1} = \frac{0.67}{3.2} = 0.21$, corresponding to a maximum size of $n_{\infty}^* = 23$ (compared to an original total of 363 subjects included in the analysis). The prior's 95% quantile is (approximately) at $\tau = 2$, and larger values would effectively imply (with $n_{\infty}^* < 3$) an almost noninformative posterior predictive distribution.

The eventual analysis is illustrated in Figure 6. Looking at the heterogeneity's posterior (Figure 8), one can see that heterogeneity here appeared to be less than anticipated. The prediction interval is relatively wide, and on the back-transformed scale is centered at a probability of 0.11 with its 95% posterior predictive interval ranging from 0.03 to 0.34. The posterior predictive distribution's standard error is 0.70, and relative to the UISD, this roughly corresponds to an effective sample size of $\frac{3.2^2}{0.70^2} = 21$ subjects.

5.4 | Regression slopes

Bergau et al.¹¹⁶ investigated predictors of all-cause mortality among patients with an implantable cardioverter-defibrillator (ICD) device. Several potential covariables were considered, among these the *left ventricular ejection fraction (LVEF)*, which is a measure of the efficiency of

heart function that is usually determined via echocardiography. LVEF is commonly expressed in percent, where 52%–72% are normally observed in healthy individuals, while values below 30% are considered abnormal.¹¹⁷ Criteria for an indicated ICD therapy include various conditions, including thresholds on the LVEF in the range 30–40%.¹¹⁸ Five studies were found that had reported on survival analyses including LVEF as a predictor, and a meta-analysis was performed based on the coefficients standardized to a 5% point decrease in LVEF; the data are shown in Table 12. The different studies also included different sets of additional covariates in their analyses.¹¹⁶

The regressor, LVEF, here is expressed in percentages (between 0 and 100), which might just as well have been expressed as a fraction (between 0 and 1), while for the analysis a unit of a 5 percentage point decrease was used — this highlights the importance of clarifying the scale of the increment Δ_x that heterogeneity considerations are to be based on. Table 12 also shows the distributions of LVEF within studies; these are roughly similar and have standard deviations of the order of 10 percentage points. For the “reference” increment Δ_x for judging plausible heterogeneity magnitudes, we will then consider a difference of 20 percentage points, which roughly spans the bulk of LVEF values encountered in each of the studies. This also coincides with the range of values considered “normal” (52%–72%) or the difference between “normal” and “abnormal” ranges ($\geq 52\%$ vs. $< 30\%$) here. The (empirical) UISD for the present data is at $s_1 = 1.9$ (for the 5% increments shown in Table 12, corresponding to $s_1 = 7.5$ for a 20% difference).

Since the regression coefficient is to be interpreted as a logarithmic HR, we will assume a half-normal(0.5) prior for the effect corresponding to a $\Delta_x = 20$ percentage point increment (analogously to the arguments made in Sections 5.3.1 and 5.3.2). For the 5% point decreases considered in the analyses, this then implies a four-fold smaller heterogeneity, that is, a half-normal(0.125) prior.




FIGURE 6 Forest plots for the examples discussed in Sections 5.3.3 and 5.4. For the log-odds, a half-normal(1.0) prior was used, and for the log-HR regression slopes, a half-normal(0.125) prior was used for the heterogeneity τ

TABLE 12 Regression example data¹¹⁶

<i>i</i>	Study	LVEF (%)		HR [95% CI]	<i>n_i</i>	log-HR	
		Mean	SD			<i>y_i</i>	<i>σ_i</i>
1	Maciag (2012)	28.0	4.0	1.16 [0.90, 1.51]	121	0.148	0.132
2	Hage (2013)	28.0	15.0	1.23 [1.08, 1.40]	696	0.207	0.066
3	Demirel (2014)	31.9	9.3	1.29 [1.02, 1.64]	99	0.255	0.121
4	Konstantino (2016)	31.6	11.1	1.16 [1.02, 1.32]	1125	0.148	0.066
5	Rodríguez-Mañero (2016)	26.2	7.6	1.23 [1.10, 1.36]	1174	0.207	0.054

Note: Regression slopes result from survival analyses and are expressed in terms of hazard ratios (HRs) and with reference to a 5% point decrease in LVEF. The baseline means and standard deviations of LVEF values are also shown.

Analysis results for a half-normal(0.125) prior are illustrated in Figure 6. The estimates are very homogeneous, which is evident from the forest plot as well as from the estimated heterogeneity (see also Figure 8). The overall log-HR estimate is at 0.19, corresponding to 1.21-fold increased mortality hazard for a 5% point decrease (worsening) in LVEF.

5.5 | Correlations

Molloy et al.¹¹⁹ investigated the relationship between conscientiousness and medication adherence. A total of 16 relevant studies reporting correlation coefficients of the two factors were found, which were also graded according to their methodological quality. Three of the studies were rated with the highest quality score; their data are shown in Table 13. The data are also available as part of the `metafor` R package.⁹¹ In order to avoid problems due to the bounded parameter space of correlations r_i (between -1 and $+1$), we will use the Fisher- z transformed values instead. Note that, since in the present example the reported correlations (r_u) are relatively close to zero, the corresponding Fisher- z values (y_i) are almost identical here (see Table 13; r_i and y_i values only differ in their third decimal place) and the transformation eventually makes little difference.

As elaborated in Section 4.5, we expect smaller magnitudes of heterogeneity for correlation endpoints (say, mostly $\tau \leq 0.3$); the UISD is at $\sigma_1 = 1.0$, which also matches the figures we see empirically in the present data set ($s_1 = 1.004$). Van Erp et al.⁸² report a median and 95% quantile of 0.12 and 0.29, respectively, for empirically observed heterogeneity estimates from published studies. Meta-analyzing the remaining set of 13 studies from the present data set¹¹⁹ (using a uniform prior), in order to quantify the evidence “external” to the example data, yields a heterogeneity estimate of 0.07 with 95% CI [0.00, 0.17].

Heterogeneity values of $\tau = 0.1$ or $\tau = 0.2$ would imply differences between a random pair of studies of a

TABLE 13 Correlation example data^{91,119}

<i>i</i>	Study	Correlation r_i	<i>n_i</i>	Fisher's z	
				<i>y_i</i>	<i>σ_i</i>
1	Stilley (2004)	0.24	158	0.245	0.080
2	Ediger (2007)	0.05	326	0.050	0.056
3	Jerant (2011)	0.01	771	0.010	0.036

Note: r_i and n_i here denote the empirical correlation coefficients and the underlying sample sizes. The y_i are the Fisher- z transformed correlations and σ_i the associated standard errors that eventually go into the analysis (see Section 2.1). A positive effect size y_i here indicates a positive correlation.

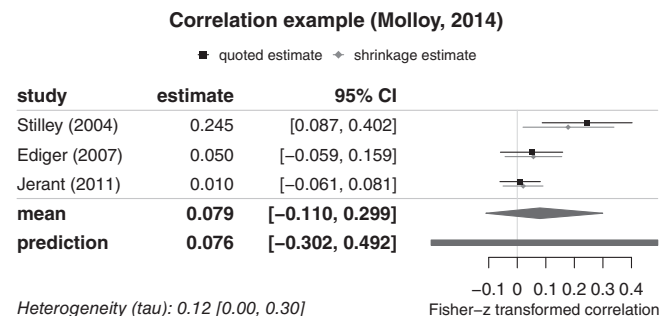


FIGURE 7 Forest plot for the example discussed in Section 5.5. For the (Fisher- z transformed) correlations, a half-normal(0.2) prior was used for the heterogeneity τ

similar order of magnitude (see Table 1). A half-normal (0.2) prior for the heterogeneity would cover values mostly in the range below 0.4, with a prior median at $\tau = 0.13$ (see Table 3).

For the present analysis, we would then suggest a half-normal(0.2) prior for the heterogeneity. A meta-analysis of the example data based on this prior is illustrated in Figure 7. The two traits were originally measured using differing scales, so that complete homogeneity might be considered especially unlikely. The heterogeneity's resulting posterior median is at $\tau = 0.12$ (with the 95% CI ranging up to 0.30), its posterior distribution is also

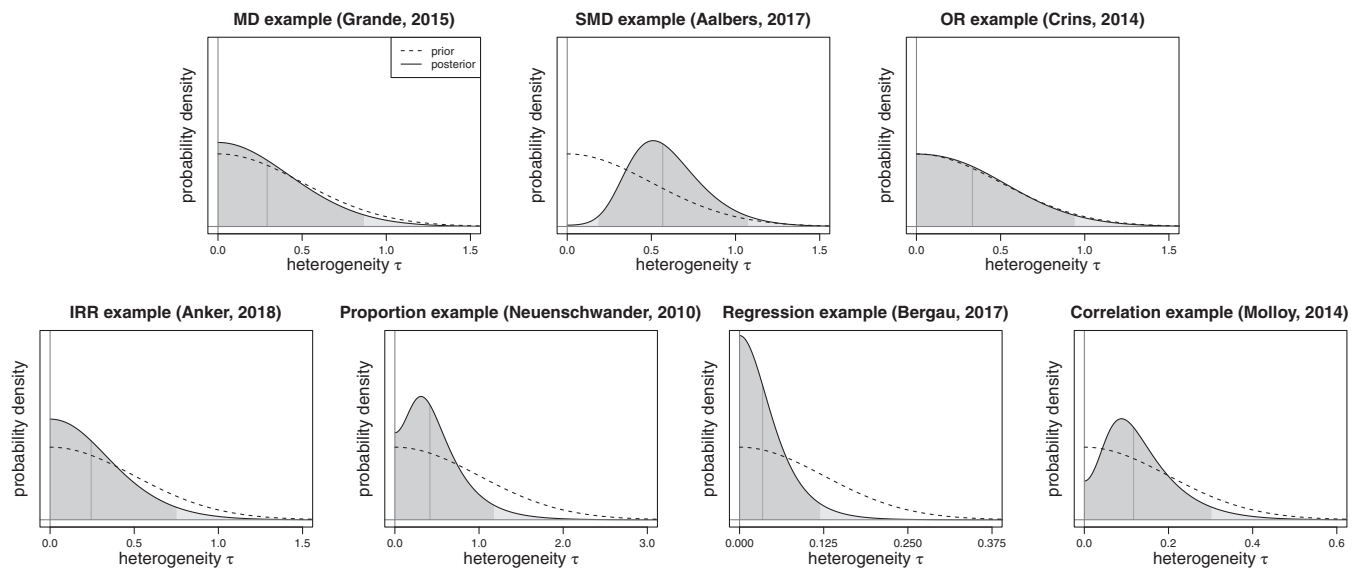


FIGURE 8 Marginal prior and posterior densities for the heterogeneity parameter τ in the seven examples discussed in Section 5. The dashed lines show the prior densities, the solid lines show the posteriors. The area shaded in dark gray indicates the 95% credible interval, the vertical line is the posterior median

illustrated in Figure 8. The three studies are of differing size and suggest neutral to slightly positive correlation between conscientiousness and medication adherence. The resulting mean estimate is positive at about 0.08, while the CI ranges from negative to positive (−0.1 to +0.3).

6 | DISCUSSION

While executing a Bayesian meta-analysis is not technically difficult, specifying a widely acceptable prior remains a challenge, especially when it comes to the heterogeneity parameter τ . Although the problem may appear complex at first, it is usually possible to break down the specification into a number of more specific questions that are easier to approach one-by-one. These steps are summarized in Table 6 and may be outlined as follows: (i) what is the effect's scale? (ii) what is the probable magnitude of other effects? (iii) how large is the unit information standard deviation (UISD)? (iv) is relevant empirical information available? The information may then be related to more concrete prior specifications by constraining (v) prior quantiles (of τ) (vi) prior predictive quantiles (of θ_i), and (vii) other prior properties. We have demonstrated the prior specification in seven applications involving few studies and covering a range of common effect scales and application areas, leading to sensible prior distributions and results in all examples. Besides the case of few studies, another context in which (weakly) informative priors are useful is whenever

marginal likelihoods (or Bayes factors) need to be computed.⁷³ Calculation of marginal likelihoods requires proper prior distributions, and special care must be taken in their selection in order to avoid (seemingly) paradoxical results.^{1,120}

In many applications, the results will be robust to variations of the prior, which may also be checked in sensitivity analyses. The prior specification will usually not be the most crucial or influential among the line of assumptions being made, which include normality,⁵ exchangeability, the selection of estimates to be pooled, or the choice between effect measures.¹²¹ Different prior specifications will of course leave their imprint on the posterior distribution, for example, results based on short- or heavy-tailed priors will reflect the differing assumptions, which may be based on emphasizing regularization or robustness aspects. There usually is no unique “correct” prior, and “skeptical” or “enthusiastic” results may be derived by implementing corresponding prior assumptions.⁴² Even uncertainty in the prior distribution itself (or its scale) may be accommodated by using mixture priors. Consideration of the stochastic ordering of heterogeneity priors may help assessing more or less conservative settings, which may be useful for the definition of sensitivity analyses. However, we would also like to warn against inflationary default specification and execution of multiple analyses here, as the resulting alternative estimates may lead to unnecessary ambiguity or inconsistent (flip-flopping) conclusions. In Supporting Information Appendix D.4, sensitivity analyses are discussed in the context of the two examples from Sections 5.1 and

5.3.2. Pre-specification of analyses (and their intended consequences) may help here. In case there is genuine a-priori uncertainty about the heterogeneity's magnitude, this might better be reflected in a single prior (e.g., in terms of a mixture distribution). Either way, one needs to be prepared and willing to base the eventual analysis results on the posterior also when the data have little information on heterogeneity to add to the weakly informative prior, as was the case for some of the examples discussed here (see Figure 8). If it is not possible to specify a suitable (weakly) informative prior for the expected heterogeneity, then one might have to resort to a more conservative approach using uninformative priors (or possibly even to a qualitative synthesis only).

Another central assumption crucial to the validity of inference is the *exchangeability* (see Section 2.1). This might be compromised by selection effects, for example, publication bias¹²² or reporting bias.¹²³ Especially in the case of only few studies, such effects might be hard to detect from the data, and information on the presence of selection effects may need to come from considerations of the context.

Choice of heterogeneity priors has consequences for estimation of the overall mean parameter, but in particular also in prediction and shrinkage applications, as the inferred heterogeneity directly impacts on the amount of borrowing-of-strength^{20,33,72}; smaller heterogeneity will lead to stronger pooling of estimates, and larger heterogeneity will imply that individual estimates are only loosely connected through the model.

Especially in regulatory settings such as drug approval or health technology assessment (HTA) the definition of a standard prior distribution for the heterogeneity parameter is important to avoid post hoc discussions in case the use of different prior distributions leads to results suggesting conflicting interpretations. The Institute for Quality and Efficiency in Health Care (IQWiG) in Germany is currently looking into determining the empirical distribution for the between-study heterogeneity parameter from all published IQWiG reports with the goal to motivate a suitable prior distribution for HTA applications.

While in the present manuscript we focused on the NNHM, some of the arguments laid out here are analogously transferable to other models for pairwise meta-analysis, for example, a Binomial-Normal model. Additional parameters and their priors may need to be specified in regard to baselines (which are often nuisance parameters and assigned vague priors).^{85,113,114,124} More complex applications in evidence synthesis such as meta-regression or network-meta-analysis would again require similar prior specifications regarding between-study heterogeneity in the effects, but would then entail additional

model components, for example, in order to accommodate individual-patient data (IPD).¹²⁵⁻¹²⁷ Analogous arguments also extend more generally to hierarchical or multilevel models, such as generalized linear mixed models (GLMMs).^{2,128} The sensitivity analyses shown in Supporting Information Appendix D.4 suggest that (for a given prior median) the prior distribution's shape has little impact on the results, as compared to the scaling of the prior. As it might simplify prior specification further, it will be interesting to investigate whether or to what extent this feature holds more generally. In summary, the application of Bayesian methods with weakly informative prior distribution for the heterogeneity parameter can be recommended for meta-analyses with random effects especially in the common case of only few studies. This paper provides guidance on the choice of useful prior distributions for various effect measures and data situations.

ACKNOWLEDGMENT

Support from the *Deutsche Forschungsgemeinschaft* (DFG) is gratefully acknowledged (grant number FR 3070/3-1).

CONFLICT OF INTEREST

The authors have declared no conflict of interest.

ORCID

Christian Röver  <https://orcid.org/0000-0002-6911-698X>

Ralf Bender  <https://orcid.org/0000-0002-2422-4362>

Sofia Dias  <https://orcid.org/0000-0002-2172-0221>

Christopher H. Schmid  <https://orcid.org/0000-0002-0855-5313>

Tim Friede  <https://orcid.org/0000-0001-5347-7441>

REFERENCES

1. Gelman A, Carlin JB, Stern H, Dunson DB, Vehtari A, Rubin DB. *Bayesian Data Analysis*. 3rd ed. Boca Raton, FL: Chapman & Hall; 2014.
2. Gelman A, Hill J. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press. 2007.
3. Hedges LV, Olkin I. *Statistical Methods for Meta-Analysis*. San Diego, CA: Academic Press; 1985.
4. Hartung J, Knapp G, Sinha BK. *Statistical Meta-Analysis with Applications*. Hoboken, NJ: John Wiley & Sons; 2008.
5. Jackson D, White IR. When should meta-analysis avoid making hidden normality assumptions? *Biom J*. 2018;60(6):1040-1058. <https://doi.org/10.1002/bimj.201800071>.
6. Röver C, Friede T. Contribution to the discussion of "When should meta-analysis avoid making hidden normality assumptions?": a Bayesian perspective. *Biom J*. 2018;60(6):1068-1070. <https://doi.org/10.1002/bimj.201800179>.
7. Friede T, Röver C, Wandel S, Neuenschwander B. Meta-analysis of few small studies in orphan diseases. *Res Synth Methods*. 2017;8(1):79-91. <https://doi.org/10.1002/jrsm.1217>.

8. Friede T, Röver C, Wandel S, Neuenschwander B. Meta-analysis of two studies in the presence of heterogeneity with applications in rare diseases. *Biom J*. 2017;59(4):658-671. <https://doi.org/10.1002/bimj.201500236>.
9. Bender R, Friede T, Koch A, et al. Methods for evidence synthesis in the case of very few studies. *Res Synth Methods*. 2018; 9(3):382-392. <https://doi.org/10.1002/jrsm.1297>.
10. Gonnemann A, Framke T, Großhennig A, Koch A. No solution yet for combining two independent studies in the presence of heterogeneity. *Stat Med*. 2015;34(16):2476-2480. <https://doi.org/10.1002/sim.6473>.
11. Smith TC, Spiegelhalter DJ, Thomas A. Bayesian approaches to random-effects meta-analysis: a comparative study. *Stat Med*. 1995;14(24):2685-2699. <https://doi.org/10.1002/sim.4780142408>.
12. Sutton AJ, Abrams KR. Bayesian methods in meta-analysis and evidence synthesis. *Stat Methods Med Res*. 2001;10(4):277-303. <https://doi.org/10.1177/096228020101000404>.
13. Schmid CH. Using Bayesian inference to perform meta-analysis. *Eval Health Prof*. 2001;24(2):165-189. <https://doi.org/10.1177/01632780122034867>.
14. Spiegelhalter DJ. Incorporating Bayesian ideas into health-care evaluation. *Statistical Science*. 2004;19(1):156-174. <https://doi.org/10.1214/088342304000000080>.
15. Spiegelhalter DJ, Abrams KR, Myles JP. *Bayesian approaches to clinical trials and health-care evaluation*. Chichester, England: John Wiley & Sons; 2004.
16. Higgins JPT, Thompson SG, Spiegelhalter DJ. A re-evaluation of random-effects meta-analysis. *J Roy Stat Soc A*. 2009;172(1): 137-159. <https://doi.org/10.1111/j.1467-985X.2008.00552.x>.
17. Lunn D, Barrett J, Sweeting A, Thompson S. Fully Bayesian hierarchical modelling in two stages with application to meta-analysis. *J Roy Stat Soc C*. 2013;62(4):551-572. <https://doi.org/10.1111/rssc.12007>.
18. Röver C, Friede T. Discrete approximation of a mixture distribution via restricted divergence. *J Comput Graph Stat*. 2017; 26(1):217-222. <https://doi.org/10.1080/10618600.2016.1276840>.
19. Röver C. bayesmeta: Bayesian random-effects meta analysis. R package; 2015. <http://cran.r-project.org/package=bayesmeta>.
20. Röver C. Bayesian random-effects meta-analysis using the bayesmeta R package. *J Stat Software*. 2020;93(6):1-51. <https://doi.org/10.18637/jss.v093.i06>.
21. Ding T, Baio G. *bmeta: Bayesian meta-analysis and meta-regression*. CRAN. R package; 2015.
22. Senn S. Trying to be precise about vagueness. *Stat Med*. 2007; 26(7):1417-1430. <https://doi.org/10.1002/sim.2639>.
23. van Dongen S. Prior specification in Bayesian statistics: three cautionary tales. *J Theor Biol*. 2006;242(1):90-100. <https://doi.org/10.1016/j.jtbi.2006.02.002>.
24. Williams DR, Rast P, Bürkner PC. Bayesian meta-analysis with weakly informative prior distributions. *PsyArXiv*. 2018. <https://doi.org/10.17605/OSF.IO/7TB8M>.
25. Gelman A. Prior distributions for variance parameters in hierarchical models. *Bayesian Anal*. 2006;1(3):515-534. <https://doi.org/10.1214/06-BA117A>.
26. Neuenschwander B, Schmidli H. Use of historical data. In: Lesaffre E, Baio G, Boulanger B, eds. *Bayesian Methods in Pharmaceutical Research*. Chapman & Hall / CRC; 2020: 111-137.
27. Prevost TC, Abrams KR, Jones DR. Hierarchical models in generalized synthesis of evidence: an example based on studies of breast cancer screening. *Stat Med*. 2000;19(24):3359-3376. [https://doi.org/10.1002/1097-0258\(20001230\)19:24<3359::AID-SIM710>3.0.CO;2-N](https://doi.org/10.1002/1097-0258(20001230)19:24<3359::AID-SIM710>3.0.CO;2-N).
28. Dias S, Sutton AJ, Welton NJ, Ades AE. Evidence synthesis for decision making 3: Heterogeneity—subgroups, meta-regression, bias, and bias-adjustment. *Med Decis Making*. 2013;33(5):618-640. <https://doi.org/10.1177/0272989X13485157>.
29. Debray TPA, Moons KGM, Valkenhoef VG, et al. Get real in individual participant data (IPD) meta-analysis: a review of the methodology. *Res Synth Methods*. 2015;6(4):293-309. <https://doi.org/10.1002/jrsm.1160>.
30. Stewart L, Moher D, Shekelle P. Why prospective registration of systematic reviews makes sense. *Syst Rev*. 2012;1:7. <https://doi.org/10.1186/2046-4053-1-7>.
31. Stevens JW. A note on dealing with missing standard errors in meta-analyses of continuous outcome measures in WinBUGS. *Pharm Stat*. 2011;10(4):374-378. <https://doi.org/10.1002/pst.491>.
32. Borenstein M, Hedges LV, Higgins JPT, Rothstein HR. A basic introduction to fixed-effect and random-effects models for meta-analysis. *Res Synth Methods*. 2010;1(2):97-111. <https://doi.org/10.1002/jrsm.12>.
33. Wandel S, Neuenschwander B, Röver C, Friede T. Using phase II data for the analysis of phase III studies: an application in rare diseases. *Clin Trials*. 2017;14(3):277-285. <https://doi.org/10.1177/1740774517699409>.
34. Pullenayegum EM. An informed reference prior for between-study heterogeneity in meta-analyses of binary outcomes. *Stat Med*. 2011;30(26):3082-3094. <https://doi.org/10.1002/sim.4326>.
35. Jaynes ET. *Probability Theory: the Logic of Science*. Cambridge: Cambridge University Press; 2003.
36. Gelman A. Bayes, Jeffreys, prior distributions and the philosophy of statistics. *Stat Sci*. 2009;24(2):176-178.
37. Cole SR, Chu H, Greenland S. Maximum likelihood, profile likelihood, and penalized likelihood: A primer. *Am J Epidemiol*. 2014;179(2):252-260. <https://doi.org/10.1093/aje/kwt245>.
38. Chung Y, Rabe-Hesketh S, Choi IH. Avoiding zero between-study variance estimates in random-effects meta-analysis. *Stat Med*. 2013;32(23):4071-4089. <https://doi.org/10.1002/sim.5821>.
39. Klein N, Kneib T. Scale-dependent priors for variance parameters in structured additive distributional regression. *Bayesian Anal*. 2016;11(4):1071-1106. <https://doi.org/10.1214/15-BA983>.
40. Bayarri MJ, Berger JO. The interplay of Bayesian and frequentist analysis. *Stat Sci*. 2004;19(1):58-80. <https://doi.org/10.1214/088342304000000116>.
41. Kass RE. Statistical inference: the big picture. *Stat Sci*. 2011;26(1):1-9. <https://doi.org/10.1214/10-STS337>.
42. Spiegelhalter DJ, Freedman L. Bayesian approaches to randomized trials. *J Roy Stat Soc A*. 1994;157(3):357-416. <https://doi.org/10.2307/2983527>.
43. Jaynes ET. Confidence intervals vs. Bayesian intervals. In: Harper WL, Hooker CA, eds. *Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*. Dordrecht: D. Reidel; 1976:175-257.
44. Datta GS, Sweeting TJ. Probability matching priors. *Handbook of Statistics*. 2005;25(3):91-114. [https://doi.org/10.1016/S0169-7161\(05\)25003-4](https://doi.org/10.1016/S0169-7161(05)25003-4).

45. Dawid AP. The well-calibrated Bayesian. *J Am Stat Assoc.* 1982;77(379):605-610. <https://doi.org/10.1080/01621459.1982.10477856>.
46. Mandelkern M. Setting confidence intervals for bounded parameters (with discussion). *Stat Sci.* 2002;17(2):149-172. <https://doi.org/10.1214/ss/1030550859>.
47. Gneiting T, Balabdaoui F, Raftery AE. Probabilistic forecasts, calibration and sharpness. *J Roy Stat Soc B.* 2007;69(2):243-268. <https://doi.org/10.1111/j.1467-9868.2007.00587.x>.
48. Gelman A. Prior choice recommendations; 2018. <https://github.com/stan-dev/stan/wiki/Prior-Choice-Recommendations#prior-for-scale-parameters-in-hierarchical-models>
49. O'Hagan A, Pericchi L. Bayesian heavy-tailed models and conflict resolution: a review. *Brazil J Prob Stat.* 2012;26(4):372-401. <https://doi.org/10.1214/11-BJPS164>.
50. Jaynes ET. Prior probabilities. *IEEE Trans Syst Sci Cybernet.* 1968;4(3):227-241. <https://doi.org/10.1109/TSSC.1968.300117>.
51. Morita S, Thall P, Müller P. Determining the effective sample size of a parametric prior. *Biometrics.* 2008;64(2):595-602. <https://doi.org/10.1111/j.1541-0420.2007.00888.x>.
52. Neuenschwander B, Weber S, Schmidli H, O'Hagan A. Predictively consistent prior effective sample sizes. *Biometrics.* 2020;76(2):578-587. <https://doi.org/10.1111/biom.13252>.
53. Evans M, Jang GH. Weak informativity and the information in one prior relative to another. *Stat Sci.* 2011;26(3):423-439. <https://doi.org/10.1214/11-STS357>.
54. Kass RE, Wasserman L. A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *J Am Stat Assoc.* 1995;90(431):928-934. <https://doi.org/10.2307/2291327>.
55. Lazar N. Ockham's razor. *Wiley Interdiscip Rev Comput Stat.* 2010;12(1):243-246. <https://doi.org/10.1002/wics.75>.
56. Ren S, Oakley JE, Stevens JW. Incorporating genuine prior information about between-study heterogeneity in random effects pairwise and network meta-analyses. *Med Decis Making.* 2018;38(4):531-542. <https://doi.org/10.1177/0272989X18759488>.
57. Hampson LV, Whitehead J, Eleftheriou D, Brogan P. Bayesian methods for the design and interpretation of clinical trials in very rare diseases. *Stat Med.* 2014;33(24):4186-4201. <https://doi.org/10.1002/sim.6225>.
58. Polson NG, Scott JG. On the half-Cauchy prior for a global scale parameter. *Bayesian Anal.* 2012;7(4):887-902. <https://doi.org/10.1214/12-BA730>.
59. Cunanan KM, Iasonos A, Shen R, Gönen M. Variance prior specification for a basket trial design using Bayesian hierarchical modeling. *Clin Trials.* 2018;16:142-153. <https://doi.org/10.1177/1740774518812779>.
60. Schmid CH, Carlin BP, Welton NJ. Bayesian methods for meta-analysis. In: Schmid CH, White I, Stijnen T, eds. *Handbook of meta-analysis*. New York: Chapman and Hall; 2021.
61. Ades AE, Lu G, Higgins JPT. The interpretation of random-effects meta-analysis in decision models. *Med Decis Making.* 2005;25(6):646-654. <https://doi.org/10.1177/0272989X05282643>.
62. Röver C, Friede T. Dynamically borrowing strength from another study. *Stat Methods Med Res.* 2020;29(1):293-308. <https://doi.org/10.1177/0962280219833079>.
63. Röver C, Friede T. Bounds for the weight of external data in shrinkage estimation. *arXiv Preprint.* 2020;2004:02525.
64. Shaked M, Shanthikumar JG. *Stochastic Orders*. New York: Springer-Verlag; 2007.
65. Mood AM, Graybill FA, Boes DC. *Introduction to the Theory of Statistics*. 3rd ed. New York: McGraw-Hill; 1974.
66. Męczarski M. Stochastic orders in the Bayesian framework. *Collegium of Economic Analysis Annals 37*, Instytut Ekonometrii, Szkoła Główna Handlowa w Warszawie (Institute of Econometrics, Warsaw School of Economics); Warsaw; 2015.
67. Bartoszewicz J, Skolimowska M. Preservation of classes of life distributions and stochastic orders under weighting. *Stat Prob Lett.* 2006;76(6):587-596. <https://doi.org/10.1016/j.spl.2005.09.003>.
68. Berger J, Berliner LM. Robust Bayes and empirical Bayes analysis with ϵ -contaminated priors. *Ann Stat.* 1986;14(2):461-486. <https://doi.org/10.1214/aos/1176349933>.
69. Berger JO. *Statistical Decision Theory and Bayesian Analysis*. 2nd ed. New York, NY: Springer-Verlag; 1985.
70. DuMouchel W. Predictive cross-validation of Bayesian meta-analyses. In: Bernardo JM, Berger JO, Dawid AP, AFM S, eds. *Bayesian Statistics*. Oxford, England: Oxford University Press; 1996:107-127.
71. DuMouchel WH, Normand SL. Computer modeling strategies for meta-analysis. In: Stangl DK, Berry DA, eds. *Meta-analysis in Medicine and Health Policy*. New York, NY: Marcel Dekker; 2000:127-178.
72. Schmidli H, Gsteiger S, Roychoudhury S, O'Hagan A, Spiegelhalter D, Neuenschwander B. Robust meta-analytic-predictive priors in clinical trials with historical control information. *Biometrics.* 2014;70(4):1023-1032. <https://doi.org/10.1111/biom.12242>.
73. Röver C, Wandel S, Friede T. Model averaging for robust extrapolation in evidence synthesis. *Stat Med.* 2018;38(4):674-694. <https://doi.org/10.1002/sim.7991>.
74. Rhodes KM, Turner RM, Higgins JPT. Predictive distributions were developed for the extent of heterogeneity in meta-analyses of continuous outcome data. *J Clin Epidemiol.* 2015;68(1):52-60. <https://doi.org/10.1016/j.jclinepi.2014.08.012>.
75. Turner RM, Jackson D, Wei Y, Thompson SG, Higgins PT. Predictive distributions for between-study heterogeneity and simple methods for their application in Bayesian meta-analysis. *Stat Med.* 2015;34(6):984-998. <https://doi.org/10.1002/sim.6381>.
76. Rücker G, Schwarzer G, Carpenter J, Okin I. Why add anything to nothing? The arcsine difference as a measure of treatment effect in meta-analysis with zero cells. *Stat Med.* 2009;28(5):721-738. <https://doi.org/10.1002/sim.3511>.
77. Neuenschwander B, Capkun-Niggli G, Branson M, Spiegelhalter DJ. Summarizing historical information on controls in clinical trials. *Clin Trials.* 2010;7(1):5-18. <https://doi.org/10.1177/1740774509356002>.
78. Higgins JPT, Whitehead A. Borrowing strength from external trials in a meta-analysis. *Stat Med.* 1996;15(24):2733-2749. [https://doi.org/10.1002/\(SICI\)1097-0258\(19961230\)15:24<2733::AID-SIM562>3.0.CO;2-O](https://doi.org/10.1002/(SICI)1097-0258(19961230)15:24<2733::AID-SIM562>3.0.CO;2-O).
79. Turner RM, Davey J, Clarke MJ, Thompson SG, Higgins JPT. Predicting the extent of heterogeneity in meta-analysis, using empirical data from the Cochrane Database of Systematic Reviews. *Int J Epidemiol.* 2012;41(3):818-827. <https://doi.org/10.1093/ije/dys041>.

80. Kontopantelis E, Springate DA, Reeves D. A re-analysis of the Cochrane Library data: the dangers of unobserved heterogeneity in meta-analyses. *PLoS One*. 2013;8(7):e69930. <https://doi.org/10.1371/journal.pone.0069930>.
81. Steel P, Kammeyer-Mueller J, Paterson TA. Improving the meta-analytic assessment of effect size variance with an informed Bayesian prior. *J Manag*. 2015;41(2):718-743. <https://doi.org/10.1177/0149206314551964>.
82. van Erp S, Verhagen J, Grasman RPPP, Wagenmakers EJ. Estimates of between-study heterogeneity for 705 meta-analyses reported in *Psychological Bulletin* from 1990–2013. *J Open Psychol Data*. 2017;5(4):4. <https://doi.org/10.5334/jopd.33>.
83. Seide SE, Röver C, Friede T. Meta-analysis data extracted from IQWiG publications. Göttingen Research Online; 2018.
84. Seide SE, Röver C, Friede T. Likelihood-based random-effects meta-analysis with few studies: empirical and simulation studies. *BMC Med Res Methodol*. 2019;19:16. <https://doi.org/10.1186/s12874-018-0618-3>.
85. Günhan BK, Röver C, Friede T. Random-effects meta-analysis of few studies involving rare events. *Res Synth Methods*. 2020; 11(1):74-90. <https://doi.org/10.1002/jrsm.1370>.
86. Hsu H, Lachenbruch PA. Paired *t* test. In: Armitage P, Colton T, eds. *Encyclopedia of Biostatistics*. 2nd ed. Chichester, England: Wiley & Sons; 2005.
87. Baker RD, Jackson D. Meta-analysis inside and outside particle physics: two traditions that should converge? *Res Synth Methods*. 2013;4(2):109-124. <https://doi.org/10.1002/jrsm.1065>.
88. Menard S. Standardized regression coefficients. In: Lewis-Beck MS, Bryman A, Liao TF, eds. *The Sage Encyclopedia of Social Science Research Methods*. Thousand Oaks, CA: Sage Publications; 2004:1069-1070.
89. Cohen J. *Statistical Power Analysis for the Behavioural Sciences*. 2nd ed. New York: Routledge; 1988.
90. Sawilowsky SS. New effect size rules of thumb. *Journal of Modern Applied Statistical Methods*. 2009;8(2):597-599. <https://doi.org/10.22237/jmasm/1257035100>.
91. Viechtbauer W. Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*. 2010;36(3). <https://doi.org/10.18637/jss.v036.i03>.
92. Trikalinos TA, Trow P, Schmid CH. *Simulation-based Comparison of Methods for Meta-analysis of Proportions and Rates*. AHRQ Publication 13(14)-EHC084-EF. Rockville, MD: Agency for Healthcare Research and Quality; 2013.
93. Becker B, Wu MJ. The synthesis of regression slopes in meta-analysis. *Stat Sci*. 2007;22(3):414-429. <https://doi.org/10.1214/07-STS243>.
94. Gelman A, Jakulin A, Pittau MG, Su YS. A weakly informative default prior distribution for logistic and other regression models. *Ann Appl Stat*. 2008;2(4):1360-1383. <https://doi.org/10.1214/08-AOAS191>.
95. Newman TB, Browner WS. In defense of standardized regression coefficients. *Epidemiology*. 1991;2(5):383-386. <https://doi.org/10.1097/00001648-199109000-00014>.
96. Greenland S, Maclure M, Schlesselman JJ, Poole C, Morgenstern H. Standardized regression coefficients: a further critique and review of some alternatives. *Epidemiology*. 1991;2(5):387-392. <https://doi.org/10.1097/00001648-199109000-00016>.
97. Bring J. How to standardize regression coefficients. *Am Stat*. 1994;48(3):209-213.
98. Schulze R. *Meta-Analysis: A Comparison of Approaches*. Hogrefe & Huber: Göttingen; 2004.
99. Correlation AP. In: Armitage P, Colton T, eds. *Encyclopedia of Biostatistics*. 2nd ed. Chichester, England: Wiley & Sons; 2005.
100. Grande AJ, Keogh J, Hoffmann TC, Beller EM, Del Mar CB. Exercise versus no exercise for the occurrence, severity and duration of acute respiratory infections. *Cochrane Database Syst Rev*. 2015;6:CD010596. <https://doi.org/10.1002/14651858.CD010596.pub2>.
101. Del Mar C. Understanding the burden of acute respiratory infections. *Annales Nestlé*. 2000;58(2):41-48.
102. Ambrosino N et al. Acute lower respiratory infections. In: Gibson GJ, Loddenkemper R, Sibille Y, Lundbäck B, eds. *European Lung White Book*. Sheffield, England: European Respiratory Society (ERS); 2013:210-223.
103. Sanders S, Doust J, Del Mar C. *Acute Respiratory Infections*. *Health Sciences & Medicine Papers* 86. Robina, Qld, Australia: Bond University; 2008.
104. Aalbers S, Fusar-Poli L, Freeman RE, et al. Music therapy for depression. *Cochrane Database Syst Rev*. 2017;11:CD004517. <https://doi.org/10.1002/14651858.CD004517.pub3>.
105. Albornoz Y. The effects of group improvisational music therapy on depression in adolescents and adults with substance abuse: a randomized controlled trial. *Nord J Music Ther*. 2011;20(3):224. <https://doi.org/10.1080/08098131.2010.522717>.
106. Hamilton M. A rating scale for depression. *J Neurol Neurosurg Psychiatry*. 1960;23:56-62. <https://doi.org/10.1136/jnnp.23.1.56>.
107. Kriston L, von Wolff A. Not as golden as standards should be: interpretation of the Hamilton rating scale for depression. *J Affect Disord*. 2011;128(1):175-177. <https://doi.org/10.1016/j.jad.2010.07.011>.
108. Furukawa T, Akechi T, Azuma H, Okuyama T, Higuchi T. Evidence-based guidelines for interpretation of the Hamilton rating scale for depression. *J Clin Psychopharmacol*. 2007;27(5):531-534. <https://doi.org/10.1097/JCP.0b013e31814f30b1>.
109. Masson SC, Tejani AM. Minimum clinically important differences identified for commonly used depression rating scales. *J Clin Epidemiol*. 2013;66(7):807. <https://doi.org/10.1016/j.jclinepi.2013.01.010>.
110. Crins ND, Röver C, Goralczyk AD, Friede T. Interleukin-2 receptor antagonists for pediatric liver transplant recipients: a systematic review and meta-analysis of controlled studies. *Pediatr Transplant*. 2014;18(8):839-850. <https://doi.org/10.1111/ptr.12362>.
111. Goralczyk AD, Hauke N, Bari N, Tsui TY, Lorf T, et al. Interleukin-2 receptor antagonists for liver transplant recipients: A systematic review and meta-analysis of controlled studies. *Hepatology*. 2011;54(2):541-554. <https://doi.org/10.1002/hep.24385>.
112. Anker SD, Kirwan BA, van Veldhuisen DJ, et al. Effects of ferri-carboxymaltose on hospitalisations and mortality rates in iron-deficient heart failure patients: an individual patient data meta-analysis. *Eur J Heart Fail*. 2018;20(1):125-133. <https://doi.org/10.1002/ehf.823>.
113. Dias S, Ades AE. Absolute or relative effects? Arm-based synthesis of trial data. *Res Synth Methods*. 2016;7(1):23-28. <https://doi.org/10.1002/jrsm.1184>.

114. White I, Turner RM, Karahalios A, Salanti G. A comparison of arm-based and contrast-based models for network meta-analysis. *Stat Med*. 2019;38(27):5197-5213. <https://doi.org/10.1002/sim.8360>.
115. Feagan BG et al. Treatment of ulcerative colitis with a humanized antibody to the $\alpha_4\beta_7$ integrin. *N Engl J Med*. 2005;352(24):2499-2507. <https://doi.org/10.1056/NEJMoa042982>.
116. Bergau L, Tichelbäcker T, Kessel B, et al. Predictors of mortality and ICD shock therapy in primary prophylactic ICD patients—a systematic review and meta-analysis. *PLoS One*. 2016;12(10):e0186387. <https://doi.org/10.1371/journal.pone.0186387>.
117. Lang RM, Badano LP, Mor-Avi V, Afilalo J, et al. Recommendations for cardiac chamber quantification by echocardiography in adults: an update from the American Society for Echocardiography and the European Association for Cardiovascular Imaging. *Europ Heart J*. 2015;16(3):233-271. <https://doi.org/10.1093/ehjci/jev014>.
118. Tracy CM, Epstein AE, Darbar D, others. ACCF/AHA/HRS focused update incorporated into the ACCF/AHA/HRS 2008 guidelines for device-based therapy of cardiac rhythm abnormalities: a report of the American College of Cardiology Foundation / American Heart Association task force on practice guidelines and the Heart Rhythm Society. *J Am College Cardiol*. 2012;61(3):e6-e75. <https://doi.org/10.1016/j.jacc.2012.11.007>.
119. Molloy GJ, O'Carroll RE, Ferguson E. Conscientiousness and medication adherence: a meta-analysis. *Ann Behaviour Med*. 2014;47(1):92-101. <https://doi.org/10.1007/s12160-013-9524-4>.
120. Lindley DV. A statistical paradox. *Biometrika*. 1957;44(1/2):187-192. <https://doi.org/10.1093/biomet/44.1-2.187>.
121. Deeks JJ, Altman D. Effect measures for meta-analysis of trials with binary outcomes. In: Egger M, Davey Smith G, Altman D, eds. *Systematic Reviews in Health Care: Meta-Analysis in Context*. 2nd ed. London: BMJ Publishing; 2001: 313-335.
122. Rothstein HR, Sutton AJ, Borenstein M. *Publication Bias in Meta-analysis*. Chichester, England: Wiley & Sons; 2005.
123. Williamson PR, Gamble C, Altman DG, Hutton JL. Outcome selection bias in meta-analysis. *Stat Methods Med Res*. 2005;14(5):515-524.
124. Wang Z, Lin L, Hodges JS, Chu H. The impact of covariance priors on arm-based Bayesian network meta-analyses with binary outcomes. *Stat Med*. 2020;39:2883-2900. <https://doi.org/10.1002/sim.8580>.
125. Debray T, Moons KGM, Abo-Zaid GMA, Koffijberg H, Riley RD. Individual participant data meta-analysis for a binary outcome: one-stage or two-stage? *PLoS One*. 2013;8(4):e60650. <https://doi.org/10.1371/journal.pone.0060650>.
126. Burke DL, Ensor J, Riley RD. Meta-analysis using individual participant data: one-stage and two-stage approaches, and why they may differ. *Stat Med*. 2017;36(5):855-875. <https://doi.org/10.1002/sim.7141>.
127. Kontopantelis E. A comparison of one-stage vs two-stage individual patient data meta-analysis methods: a simulation study. *Res Synth Methods*. 2018;9(3):417-430. <https://doi.org/10.1002/jrsm.1303>.
128. Brown H, Prescott R. *Applied Mixed Models in Medicine*. 3rd ed. Chichester, England: John Wiley & Sons; 2014.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Röver C, Bender R, Dias S, et al. On weakly informative prior distributions for the heterogeneity parameter in Bayesian random-effects meta-analysis. *Res Syn Meth*. 2021;1–27. <https://doi.org/10.1002/jrsm.1475>