

# Summarizing empirical information on between-study heterogeneity for Bayesian random-effects meta-analysis

Christian Röver<sup>1</sup>  | Sibylle Sturtz<sup>2</sup> | Jona Lilienthal<sup>2</sup> | Ralf Bender<sup>2</sup> | Tim Friede<sup>1</sup> 

<sup>1</sup>Department of Medical Statistics,  
University Medical Center Göttingen,  
Göttingen, Germany

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

## Correspondence

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

Email:

[christian.roever@med.uni-goettingen.de](mailto:christian.roever@med.uni-goettingen.de)

## Funding information

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

In Bayesian meta-analysis, the specification of prior probabilities for the between-study heterogeneity is commonly required, and is of particular benefit in situations where only few studies are included. Among the considerations in the set-up of such prior distributions, the consultation of available *empirical data* on a set of relevant past analyses sometimes plays a role. How exactly to summarize historical data sensibly is not immediately obvious; in particular, the investigation of an empirical collection of heterogeneity *estimates* will not target the actual problem and will usually only be of limited use. The commonly used normal-normal hierarchical model for random-effects meta-analysis is extended to infer a heterogeneity prior. Using an example data set, we demonstrate how to fit a distribution to empirically observed heterogeneity data from a set of meta-analyses. Considerations also include the choice of a parametric distribution family. Here, we focus on simple and readily applicable approaches to then translate these into (prior) probability distributions.

## KEYWORDS

external information, heterogeneity, hierarchical model, meta-analysis, prior distribution

## 1 | INTRODUCTION

A range of statistical modeling approaches is available for random-effects meta-analyses; here we consider meta-analysis within the common and general framework of the *normal-normal hierarchical model (NNHM)* where measurement uncertainty as well as between-study heterogeneity are modeled based on normal distributions.<sup>1,2</sup> The NNHM provides a versatile framework for meta-analysis that is applicable in a wide range of cases where the data may be summarized by a point estimate along with a measure of uncertainty (a standard error or confidence interval). Examples include many types of endpoints, for example, (standardized) mean differences, (logarithmic) odds ratios, relative risks or hazard ratios, prevalences, correlation coefficients, and many more. Analysis based on the NNHM and performed within a Bayesian framework requires the assignment of prior distributions for the overall mean effect ( $\mu$ ) and the heterogeneity standard deviation ( $\tau$ ); while the former is commonly uncontroversial and a uniform (uninformative) specification if often appropriate, the latter may require more care, in particular in cases where only few studies are included in the meta-analysis.<sup>3-5</sup>

Generally speaking, the construction of prior distributions may be approached from a range of different angles,<sup>6-9</sup> some of which include the consideration of empirical data. Besides the empirical evidence, a number of further aspects may determine details of the prior distribution's specification, such as tail behavior, robustness or conservatism.<sup>2,4</sup>

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Statistics in Medicine* published by John Wiley & Sons Ltd.

Practical application eventually also requires priors to be reasonably simple, and easily motivated, communicated and implemented.

Several authors have investigated empirical data on heterogeneity in meta-analysis applications previously.<sup>10-18</sup> In some instances, this resulted in a sample of heterogeneity *estimates*, which may provide a rough idea of plausible heterogeneity magnitudes. However, it may not provide immediate guidance in specifying a prior for the heterogeneity parameter for practical application. On the other hand, Rhodes et al<sup>19</sup> and Turner et al<sup>20</sup> fitted comprehensive, fully Bayesian models with a focus on prediction, and with the aim of yielding immediately applicable prior specifications. To this end, Turner et al<sup>20</sup> utilized a log-normal model for binary outcomes, while Rhodes et al,<sup>19</sup> after considering several alternatives, eventually settled on a log-Student- $t_5$  model for standardized mean differences.

While the results from such previous analyses are immediately applicable for certain investigations, in some cases it is desirable to perform an analogous investigation based on a specific selected data base; an example scenario also sketched in the example application below is that of a health technology assessment (HTA) authority that may be interested in evaluating the heterogeneity commonly encountered in a specific set of past analyses. The technical problem essentially is that of a *meta-analysis of heterogeneity estimates* with a focus on the heterogeneity “population” and prediction of a new (“future”) heterogeneity value. The heterogeneity estimates themselves each originate from a meta-analysis, and they are usually not well summarized simply by a point estimate and a standard error. The relevant data basis hence is a complete set of *several* meta-analyses (ie, their included effect estimates and standard errors); at least in the medical context, such original data are quite commonly reported as required by the PRISMA statement.<sup>21</sup>

In the present investigation, we formalize the approach originally established by Rhodes et al<sup>19</sup> and Turner et al<sup>20</sup> in a slightly more general way, including a range of parametric distribution families and suggestions on how to summarize the resulting posterior distributions for communication and application. Modeling here includes choice of the random-effects distribution, and the focus is on prediction. While we showcase the approach using a freely available example data set, we also provide JAGS and R code to facilitate application of the same approach to other data sets.

The remainder of this article is structured as follows. In Section 2, we outline the problem in more detail and we specify the statistical model employed to investigate and infer heterogeneity distributions. In Section 3, the approach is showcased using a small example data set. Section 4 then concludes with a discussion.

## 2 | MODELING HETEROGENEITY

### 2.1 | Preliminary considerations

In the following, we consider the *normal-normal hierarchical model (NNHM)*, in which estimation uncertainty as well as *between-study heterogeneity* are accommodated using normal variance components. The magnitude of the heterogeneity component, which empirically manifests itself as excess variability beyond what could be attributed to uncertainty alone, scales with the standard deviation  $\tau$ , which is generally unknown. In a Bayesian context, the specification of a prior distribution for  $\tau$  is then required.<sup>2,4</sup> In some instances, non-informative priors may be specified; in particular, in the context of a large number of studies included in a meta-analysis, varying the heterogeneity prior may be of little relevance to the resulting effect estimates. In general, however, consideration of prior information in the analysis will always be beneficial, and sometimes even necessary, for example, when dealing with less than about 10 studies only, or when the focus is on computing marginal likelihoods. Prior distributions then may or may not involve empirical information.<sup>4</sup>

In order to support an analysis by informing the prior choice using empirical data, first of all relevant, somehow “representative” data need to be identified. *How* exactly to arrive at such a data set is beyond the scope of the present investigation; here we assume an appropriate set of meta-analyses based on suitable effect measures to be given. In case one set out to gather a set of relevant heterogeneity estimates (and the associated meta-analyses), some guidance is given by the recommendations issued for meta-analyses in general, for example, the PRISMA statement.<sup>21</sup> Inclusion criteria should be defined, and restrictions might relate to characteristics like outcome (eg, mortality), effect measure (eg, odds ratio) or comparators (eg, pharmacological vs. placebo), as was also done in the investigations by Rhodes et al<sup>19</sup> and Turner et al.<sup>20</sup> An industrial investigator might consider a subset of their own sponsored trials, or an HTA institution might consider their own meta-analyses, and one might want to include time constraints (eg, the past decade); either way, a transparent specification of the selection process will (as usual) make the investigation more convincing. In contrast to conventional systematic reviews (searches for *studies* to be pooled), the search for meta-analyses would usually need to

have a wider scope and may commonly include different indications, different treatments, different controls, etc. In the following, we discuss how such data may then be modeled and considered in a future analysis.

It should be noted, however, that empirical data still remain only one aspect of several to guide prior specification; consideration of empirical data may in fact also be seen as “complementary” to other prior considerations. The heterogeneity parameter refers to the amount of variability *between studies*, and, based on the context and the endpoint considered, it is usually possible to provide a rough specification of what amounts are considered likely, or from which point on the heterogeneity is deemed unreasonably large.<sup>4</sup> An “empirical” prior distribution might also serve as a “sanity check”, for example, in order to confirm whether some given prior specification appears too optimistic or otherwise unrealistic. Furthermore, it will usually be possible to adapt a given prior distribution to make it “less informative”, “more conservative” or “more robust” in a certain sense.<sup>4</sup>

If no concrete (eg, empirical) information is available, one can usually still constrain the plausible magnitude of the heterogeneity to be expected, leading to *weakly informative* priors. Such priors aim to facilitate analysis by bounding the heterogeneity in a conservative way; such specifications are also closely related to regularization approaches.<sup>22,23</sup> Weakly informative priors for the heterogeneity parameter have been explored previously, and in particular their scaling depends on the context and the endpoint in question.<sup>4,24</sup>

## 2.2 | Extending the NNHM

In order to infer a heterogeneity prior for use within the NNHM context, we specify an extension of the common NNHM allowing to accommodate data from *several* meta-analyses at once, and including an additional model stage for the distribution of the associated heterogeneity parameters across meta-analyses. Suppose the external data consist of a collection of  $N$  meta-analyses, each of an individual size  $k_j$  ( $j = 1, \dots, N$ ), and providing a set of  $k_j$  estimates  $y_{ij}$  along with their associated standard errors  $\sigma_{ij}$  ( $i = 1, \dots, k_j$ ). Depending on the effect measure considered,  $y_{ij}$  and  $\sigma_{ij}$  (as well as the  $\mu_j$  and  $\tau_j$  introduced below) are all expressed in the same units.

The eventual aim of the analysis will be to determine the (*posterior*) *predictive distribution* of a “future” heterogeneity value ( $\tau^*$ ); this distribution will be relevant as a heterogeneity prior for a subsequent random-effects meta-analysis.<sup>19,20,25</sup>

## 2.3 | The normal-normal stage

Each meta-analysis has an underlying true effect  $\mu_j$  and a heterogeneity  $\tau_j$  associated. The  $j$ th meta-analysis is modeled via the “usual” NNHM:

$$y_{ij} \mid \mu_j, \tau_j \sim \text{Normal}(\mu_j, \sigma_{ij}^2 + \tau_j^2) \quad (i = 1, \dots, k_j). \quad (1)$$

Inference on different meta-analyses’ effect parameters  $\mu_j$  is effectively stratified by assuming a vague prior

$$\mu_j \mid \mu_p, \sigma_p \sim \text{Normal}(\mu_p, \sigma_p^2) \quad (2)$$

for some “neutral” prior mean  $\mu_p$  and some “large” (uninformative) prior standard deviation  $\sigma_p$ .<sup>2</sup>

## 2.4 | The heterogeneity stage

If one now chose a common “noninformative” prior for  $\tau$  across all  $N$  analyses, one would essentially recover the posteriors that would also result from independent, separate analyses of all meta-analyses individually. Instead, we specify an additional model stage by defining a joint parametric (“prior”) distribution for the heterogeneity, whose (hyper-) parameters again are to be learned from the set of meta-analyses. The purpose of this additional model stage is to capture the distribution of the heterogeneity (standard deviation-) parameters across the set of  $N$  meta-analyses.

For the heterogeneity parameters, a common parametric distribution is assumed, whose parameters in turn are to be estimated. In general, we may express this as

$$\tau_j \mid \theta \sim P(\theta) \quad (3)$$

where the heterogeneity distribution's parameter(s)  $\theta$  are assigned another hyperprior

$$\theta \sim H. \quad (4)$$

The *predictive distribution* of a “new” heterogeneity value (eg, in a future meta-analysis) is again defined through the conditional expression

$$\tau^* \mid \theta \sim P(\theta); \quad (5)$$

its (prior or posterior) marginal distribution again results from marginalizing over the corresponding distribution of  $\theta$ .

To make the specification in (3) and (4) more concrete, a common setting in practice may be to assume a half-normal distribution for  $\tau_j$ , that is,

$$\tau_j \mid s \sim \text{half-Normal}(s), \quad (6)$$

where the common scale parameter  $s$  is unknown and is assigned another hyperprior, for example,

$$s \sim \text{Uniform}(0, b) \quad (7)$$

for some fixed “large” upper bound  $b$  (eg, a bound of  $b = 10$  should generally be appropriate for a log-OR as effect measure).

In the following, inference first of all aims at inferring the common heterogeneity distribution ((3) or (6) above) and its parameters ( $\theta$  or  $s$  in the specifications above). The relevant aspect for informing future meta-analyses then is the *posterior predictive distribution* of a “new” heterogeneity value  $\tau^*$ .<sup>19,20</sup> Here this means marginalizing over (6) based on the scale parameter's posterior distribution. In the above example, this posterior predictive distribution will hence be a half-normal scale mixture.<sup>26</sup>

Many other sensible distributions may be conceivable instead of the half-normal specification in (6); the half-normal distribution, however, is probably the most common one. Inclusion of a scale parameter in the distribution's parametrization is often also sensible.<sup>4</sup> Alternatives to the half-normal would for example, be exponential, half-Cauchy or log-normal distributions. All these may be parameterized in terms of a scale parameter, the log-normal distribution in addition possesses a shape parameter. The use of gamma or inverse-gamma distributions is usually not considered appropriate.<sup>27</sup> The different prior distributions differ for example in their shapes, their tail behavior, or their complexity (number of parameters). For some guidance on their properties see for example, the extensive discussion by Röver et al (2021).<sup>4</sup> Extensions of the model's heterogeneity stage may also be conceivable, for example, in order to accommodate categorical covariables or time trends.

## 2.5 | Inference

Inference on all unknowns in the model is easily implemented via Markov chain Monte Carlo (MCMC) methods,<sup>28</sup> for example, utilizing JAGS.<sup>29,30</sup> The parameters of primary interest here are the heterogeneity distribution's parameters ( $\theta$ ) and the predictive distribution ( $\tau^*$ ). Alternative model specifications (different heterogeneity distributions in (3) or (6)) may be investigated, and a formal model comparison may be implemented based on information criteria (such as the deviance information criterion (DIC));<sup>31</sup> see also the example below).

## 2.6 | Summary and transfer

For any practical purpose, the “plain” MCMC samples alone (of  $\theta$  or  $\tau^*$ ) would only be of very limited use. In order to communicate and apply the predictive distribution as a prior in a subsequent analysis, the MCMC samples would ideally be “condensed” into a simple parametric probability distribution. There are three obvious options:

1. Derive a point estimate ( $\hat{\theta}$ ), and then use a single, “fixed” instance of (3) (ie, the conditional distribution  $p(\tau^* \mid \hat{\theta})$ ) as an approximation to the predictive distribution for later application. The point estimate here may also be chosen

“conservatively” (eg, for a scale parameter, one may prefer the mean over the median, as this will usually be larger, or one might also use some upper posterior quantile in order to be on the safe side).

2. Consider the parameter’s (or parameters’) uncertainty in addition. The predictive distribution results as a mixture distribution (eg, a *scale mixture* in case of a scale parameter  $\theta$ ). This may sometimes be approximated analytically based on the parameter’s distribution’s shape.<sup>4</sup>
3. Fit a distribution (eg, by matching moments or ML fitting) to approximate the predictive distribution (of  $\tau^*$ ) directly. Obvious choices for the distribution family may again be the heterogeneity model (3), or scale mixtures (or otherwise overdispersed versions) thereof.

In the following, we will consider these three approaches for practical application. Note that the third approach is essentially the one also implemented by Rhodes et al<sup>19</sup> and Turner et al,<sup>20</sup> or, in a related context, by Weber et al.<sup>32</sup>

When summarizing a posterior predictive distribution for practical application, the result will necessarily be a simplification, and a simple distributional form (a common, simple parametric distribution, with parameters rounded to reasonable accuracy) will foster applicability. On the other hand, care must be taken not to *oversimplify* matters, and whenever deviations are accepted, one may aim to ensure that discrepancies are on the “conservative” or “less informative” side. Generally speaking, *underestimation* of heterogeneity is commonly considered more harmful than *overestimation*, and consequently one may try and preferably push potential bias towards larger heterogeneity.<sup>4</sup> In that sense, for example, a stochastically larger distribution, a larger scale parameter, or a heavier-tailed distribution may be considered *more conservative*.<sup>2,4</sup>

### 3 | EXAMPLE APPLICATION

#### 3.1 | Example data set

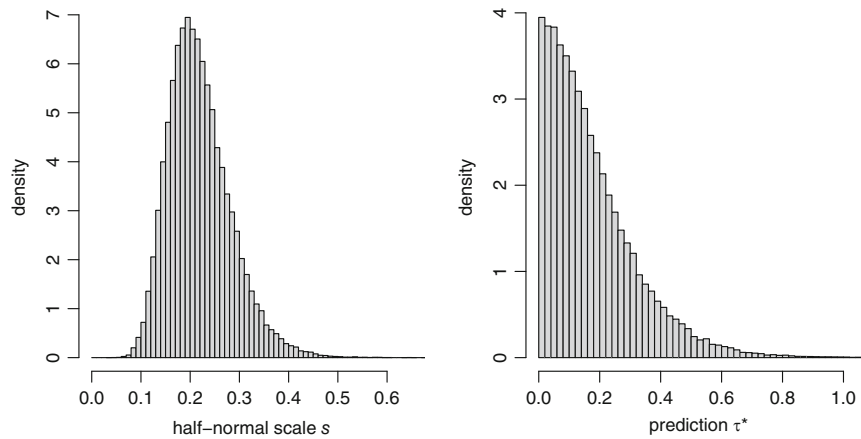
As an example, we will utilize the meta-analysis data set compiled by Seide et al.<sup>16,17</sup> This data set contains data from 40 meta-analyses and includes a total of 131 studies. The data were originally assembled as a realistic test bed for comparing different meta-analysis methods for binary data. Reports published by the German Institute for Quality and Efficiency in Health Care (*Institut für Qualität und Wirtschaftlichkeit im Gesundheitswesen, IQWiG*) were individually screened, and data from the first encountered meta-analysis utilizing a binary endpoint and not including zero counts were extracted.<sup>16,17</sup> In the following, we will consider logarithmic odds ratios (log-ORs) as effect measures. In this data set, the individual meta-analyses may be considered independent, since all of them originate from different reports and there is no overlap among the included studies. One might then consider the derived predictive distribution relevant as a prior distribution for inference in a “future” report also employing a binary endpoint.

Selection of this particular data set was somewhat opportunistic, as it was readily available, realistic and of a suitable size. However, one could easily imagine a very similar situation to come up in practice, where some health technology assessment (HTA) authority would be interested in characterizing the amount of heterogeneity encountered in a more or less confined subset of past analyses. In fact, an investigation of this kind is currently under way at the IQWiG; for more detail, see also Section 4.

Before considering the empirical evidence in the following, it may be worthwhile reviewing what the corresponding *a priori* expectations may be. Here it is important to note that the endpoint (log-OR) is defined on a logarithmic scale, so that the additive variation implied by the between-study heterogeneity translates to a *multiplicative* offset for the associated ORs. For example, if the heterogeneity was at  $\tau = 1.0$ , this would imply study-specific effects within a range of mostly  $\mu \pm 1.96$  (with 95% probability) on the log-OR scale. On the OR scale, this corresponds to effects within  $\frac{\exp(\mu)}{7.1}$  and  $\exp(\mu) \times 7.1$ , potentially implying quite dramatic differences in effects between studies.<sup>33</sup> By considering these effect scales, one would commonly assume that  $\tau$  would tend to be  $<1.0$ , for example by implementing a half-normal distribution with scale 0.18,<sup>34</sup> 0.32,<sup>35</sup> or 0.50.<sup>3</sup> A half-normal prior distribution with scale 0.50 is commonly considered a conservative choice for endpoints such as log-ORs, as first of all it rules out unreasonably large heterogeneity ranges.<sup>3-5</sup>

#### 3.2 | The analysis

The model described in Section 2, using a half-normal distribution for the heterogeneity  $\tau$  (as in (6)), a uniform prior for the scale parameter  $s$  (as in (7)) and a Normal(0, 100<sup>2</sup>) prior for the study-specific means was fitted using JAGS.<sup>29,30</sup> The



**FIGURE 1** Histograms of the posterior distribution of the half-normal scale parameter  $s$  (left panel) and of predictions  $\tau^*$  (right panel).

**TABLE 1** Summary statistics for the posterior distributions of scale parameter  $s$  and prediction  $\theta^*$  (see also Figure 1).

|                    | Scale $s$ | Prediction $\tau^*$ |
|--------------------|-----------|---------------------|
| Mean               | 0.22      | 0.17                |
| Standard deviation | 0.064     | 0.15                |
| Median             | 0.21      | 0.14                |
| 95% quantile       | 0.33      | 0.46                |
| 99% quantile       | 0.41      | 0.66                |

JAGS code implementing the model is given in Appendix A. The resulting posterior distribution of the scale parameter  $s$  and the predictive distribution of  $\tau^*$  are shown in Figure 1. Both distributions are also summarized using some descriptive statistics in Table 1. The distribution of primary relevance for a “future” analysis is the predictive distribution of  $\tau^*$ . Subsequent inference hence needs to aim at roughly capturing its shape and relevant features.

### 3.3 | Point estimation

An obvious and simple way of characterizing the predictive distribution is by simply deriving a “representative” point estimate of the heterogeneity distribution’s parameter(s) (here: the scale parameter  $s$ ) and consider the corresponding conditional distribution. In the present case, we for example have a posterior mean scale of  $\bar{s} = 0.22$  (see Table 1). The corresponding *conditional* half-normal distribution ( $p(\tau^* | s = 0.22)$ ) is shown along with the original histogram as a red line in Figure 2. Alternatively, one might also pick a more conservative (or “worst case”) point estimate, for example, the upper 95% quantile of the scale parameter’s posterior instead.

### 3.4 | Consideration of estimation uncertainty

Due to the uncertainty in the scale parameter  $s$ , the histogram of predictions  $\tau^*$  results as a (*scale*) *mixture distribution*<sup>26</sup> that generally includes more extreme (both small and large) values than a half-normal distribution could accommodate. An obvious example of a (half-) normal scale mixture is the (half-) Student- $t$  distribution, so that this may also be a natural choice for approximating the predictive distribution. The Student- $t$  distribution arises as a scale mixture of a normal distribution whose scale parameter (the normal standard deviation) follows a *scaled inverse  $\chi$ -distribution*.<sup>4</sup> By matching a scaled inverse  $\chi$ -distribution to the scale parameter’s distribution, we may then derive a corresponding half-Student- $t$  distribution as an approximation to the predictive distribution.

In the present case, the scale parameter has a mean of 0.22 and standard deviation 0.064 (see Table 1), corresponding to a coefficient of variation of 0.29. The resulting scale mixture may hence be approximated by a half-Student- $t$  distribution

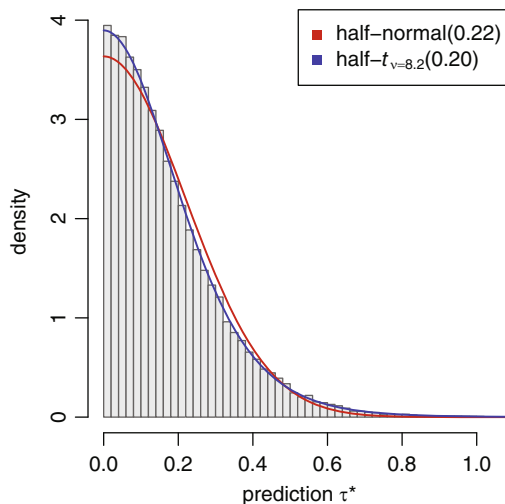


FIGURE 2 Histogram of the heterogeneity predictions  $\tau^*$  along with two approximations to its distribution.

TABLE 2 Summary statistics for the two approximations to the predictive distribution in comparison to the figures based on MCMC (see also Table 1).

|                                   | Mean | Std.dev. | Quantiles |      |      |
|-----------------------------------|------|----------|-----------|------|------|
|                                   |      |          | 50%       | 95%  | 99%  |
| Prediction $\tau^*$               | 0.17 | 0.15     | 0.14      | 0.46 | 0.66 |
| Half-normal (0.22)                | 0.18 | 0.13     | 0.15      | 0.43 | 0.57 |
| Half-Student- $t_{\nu=8.2}(0.20)$ | 0.18 | 0.15     | 0.14      | 0.46 | 0.67 |

with  $\nu = 8.2$  degrees of freedom and scale  $s = 0.20$ . This distribution is also shown in Figure 2 as a blue line, and one can see that this yields an improved fit compared to the half-normal approximation. For details on fitting a Student- $t$  distribution as a normal scale mixture, see Appendix B.

### 3.5 | Matching the predictive distribution

Instead of matching the predictive distribution based on the scale parameter distribution's shape, the predictive distribution may also be fitted to the sampled  $\tau^*$  values directly. Different distribution families may be reasonable here; more flexible (in particular: overdispersed) variations of the heterogeneity distribution (Equations (3) or (6)) are sensible choices. In the present example, a half-Student- $t$  distribution hence remains an obvious choice for the parametric family. Deriving maximum-likelihood estimates for degrees-of-freedom and scale parameters is straightforward using statistical software (via numerical optimization). Alternatively, one may also derive moment estimates; for the half-Student- $t$  distribution, one may first solve for the degrees-of-freedom and subsequently for the scale parameter (see Appendix C for more details). Another way to set up an overdispersed predictive distribution might be to fit a mixture distribution based on a small number of components, say 2–4.<sup>32</sup>

For the present example, both maximum-likelihood as well as moment estimation again yield estimates of 8.2 for the degrees-of-freedom ( $\nu$ ), and of 0.20 for the scale ( $s$ ) here; see also the corresponding distribution sketched in Figure 2. Table 2 also contrasts summary statistics of the predictive distribution (based on MCMC) to the half-normal or half-Student- $t$  fits.

### 3.6 | Model selection

Beyond consideration of common conventions or simplicity, selection of a parametric model to be fitted to the samples may also be based on designated *model selection* approaches. For example, the *deviance information criterion* (DIC) may

TABLE 3 DIC values for the comparison of three alternative models based on the example data.

| Model       | DIC   | Predictive distribution ( $\tau^*$ ) |          |      |      |      |
|-------------|-------|--------------------------------------|----------|------|------|------|
|             |       | Mean                                 | Std.dev. | 50%  | 95%  | 99%  |
| Half-normal | 163.8 | 0.17                                 | 0.15     | 0.14 | 0.46 | 0.66 |
| Exponential | 167.7 | 0.17                                 | 0.19     | 0.11 | 0.54 | 0.91 |
| Log-normal  | 178.0 | 0.26                                 | 3.46     | 0.11 | 0.61 | 1.90 |
| Half-Cauchy | 212.8 |                                      |          | 0.08 | 1.10 | 5.46 |

Note: A lower DIC value indicates a better model fit. The corresponding predictive distributions are also summarized (in analogy to Table 2).

be computed from MCMC (JAGS) output in order to judge the goodness-of-fit or predictive performance of different distribution families.<sup>31,36-39</sup> However, such model selection approaches are usually computationally challenging, and it may be questionable whether the often slight differences between a range of reasonable parametric distribution families are in fact of substantial practical relevance. For example, sensitivity analyses shown by Röver et al<sup>4</sup> seem to suggest that analyses based on priors from different families yet with matching medians may commonly lead to barely distinguishable results; this also seems to be confirmed by the alternative analyses shown in Section 3.9.

In the present example, we may compare a half-normal model with analogous ones utilizing an exponential, a half-Cauchy or a log-normal distribution for the heterogeneity (instead of the specification in (6)). Table 3 shows DIC values based on the example data; the half-normal model fits best here.

### 3.7 | Considerations beyond empirical data

Empirical data will only ever be one among several aspects in the specification of a prior distribution.<sup>6,7,40</sup> Other aspects include operating characteristics and robustness,<sup>25,41</sup> considerations of conservatism properties,<sup>2</sup> features like lower-tail and upper-tail behaviors,<sup>4</sup> or simplicity<sup>42</sup> (see also Section 2.6). In the present example, if simplicity is desired, the half-normal approximation may be preferred to the Student-*t* model. If on the other hand there is some doubt about the direct relevance of the empirical data to the given meta-analysis problem, a “more conservative” distribution (eg, a greater scale parameter) might be specified.<sup>2</sup>

Another way to “robustify” a prior, or to reconcile several prior information sources, is by implementing a *mixture prior*. The basic idea is that the prior information may be composed of several (mutually exclusive) components, for example, besides the hypothesis that the historical data are of immediate relevance to the present analysis, one may want to also consider the alternative that it is unrelated and hence some vague prior would apply. Both the “informative” and “vague” priors may be combined by attaching probabilities to both options and using a two-component mixture distribution of both as the eventual prior.<sup>25,43</sup> Inference would eventually also consider to what extent the present data appear consistent with one or another component; the effect is greater conservatism as well as a gain in *robustness* to discrepancies between the historical and current data.

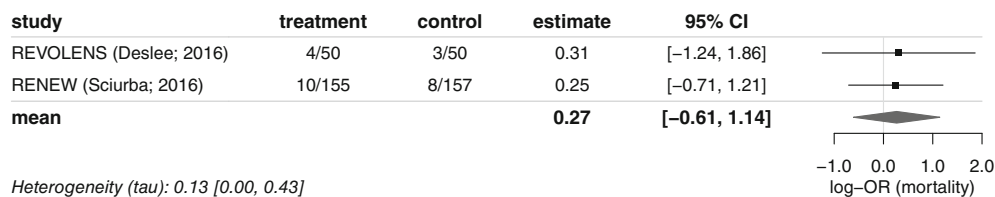
### 3.8 | Example application

We continued the original search through IQWiG publications to find the subsequent “41st” qualifying meta-analysis to apply the derived prior on.

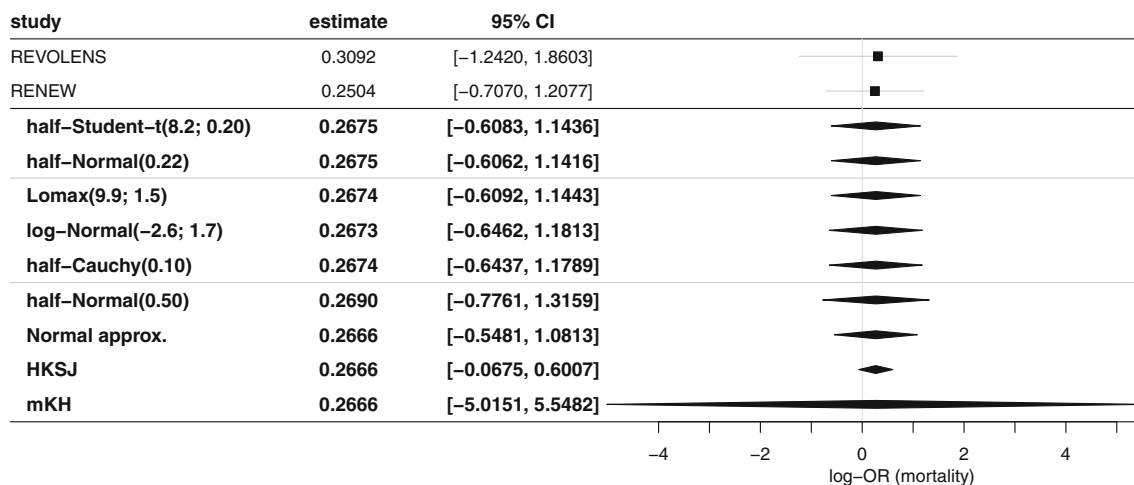
Report N14-04<sup>44</sup> (fig.32, p. 233) reports on a meta-analysis involving two studies, and comparing *endobronchial coil* vs. *no therapy* for lung volume reduction in severe pulmonary emphysema.<sup>45-47</sup> The endpoint considered here is the log-OR of *overall mortality at 12 months*. The relevant data are illustrated in Figure 3. The two studies yield very similar log-OR estimates, while the standard errors differ. The (DL) heterogeneity estimate turns out as zero in this case, so that the original (frequentist) analysis effectively was based on a “common-effect” estimate of a log-OR of 0.27 [−0.55, 1.08].<sup>44</sup>

Given that the half-normal model provides the best fit to the historical data (see Table 3) and that the half-Student-*t* yields a slightly better approximation to the predictive distribution than the half-normal (see Figure 2), here we select the half-Student-*t*<sub>v=8.2</sub>(0.20) approximation for analysis. Some properties of this distribution are sketched in Figure 2 or





**FIGURE 3** Forest plot illustrating the meta-analysis of log-OR estimates from a subsequent IQWiG report<sup>44</sup> based on the heterogeneity prior derived from previous, related meta-analyses (a half-Student- $t$  distribution with 8.2 degrees of freedom and scale 0.20).



**FIGURE 4** Forest plot illustrating different (Bayesian and frequentist) analyses of the example data from Section 3.8.

in Table 2. The resulting pooled effect estimate is shown in Figure 3. The resulting credible interval is about 8% wider than the originally reported frequentist interval due to the increased heterogeneity incorporated in the Bayesian model. The heterogeneity posterior barely differs from its prior, as one can also see when comparing the heterogeneity estimate shown in the bottom left of Figure 3 with the prior quantiles shown in Table 2 (prior and posterior densities are also shown side by side in Figure E1 in the Appendix).

### 3.9 | Variations of the example analysis

#### 3.9.1 | Alternative analyses

For comparison, we also analyzed the same data using several alternative approaches, including the alternative priors mentioned in Section 3.6 that were derived based on differing distribution families for the heterogeneity, a Bayesian approach with a common weakly informative prior, as well as three frequentist estimators.

Figure 4 illustrates the different results jointly in a forest plot. The first two (half-Student- $t$  and half-normal) are the ones also shown in Figure 2, and are based on assuming a half-normal distribution for the random effects. The following three lines show the summary estimates resulting from assuming alternative models for the external data (see also Table 4). The last four lines show common analyses not considering the external data; the half-normal heterogeneity prior with scale 0.50 is commonly considered *weakly informative* for a log-OR endpoint, as in the present context.<sup>3,4</sup> The following three confidence intervals correspond to common frequentist analysis methods and are based on a simple normal approximation (which was also utilized in the original analysis<sup>44</sup>) and the Hartung-Knapp-Sidik-Jonkman (HKSJ) and the modified Knapp-Hartung (mKH) intervals,<sup>48</sup> which are both based on Student- $t$  quantiles.

Use of the different approximations to the predictive distribution in the half-normal model (half-normal or half-Student- $t$ ; see Figure 2) barely makes a noticeable difference for the eventual analysis. Assuming alternative random-effect distributions instead of the half-normal model leads to slightly differing predictive distributions as

TABLE 4 Parametric fits to the distributions shown in Table 3.

| Model       | predictive distribution ( $\tau^*$ ) |      |          |      |      |      |
|-------------|--------------------------------------|------|----------|------|------|------|
|             | Approximation                        | Mean | Std.dev. | 50%  | 95%  | 99%  |
| Half-normal | Half-normal (0.22)                   | 0.18 | 0.13     | 0.15 | 0.43 | 0.57 |
| Half-normal | Half-Student- $t$ (8.2; 0.20)        | 0.18 | 0.15     | 0.14 | 0.46 | 0.66 |
| Exponential | Lomax (9.9; 1.5)                     | 0.17 | 0.19     | 0.11 | 0.53 | 0.89 |
| Log-normal  | Log-normal (-2.6; 1.7)               | 0.32 | 1.30     | 0.07 | 1.22 | 3.88 |
| Half-Cauchy | Half-Cauchy (0.10)                   |      |          | 0.10 | 1.27 | 6.37 |

Note: The approximations for the exponential, log-normal and half-Cauchy models are based on maximum-likelihood estimates.

TABLE 5 Predictive distributions (and their approximations) based on analyses of data subsets.

| Data                    | Predictive distribution ( $\tau^*$ ) |          |      |      |      | Approximations                                   |
|-------------------------|--------------------------------------|----------|------|------|------|--|
|                         | Mean                                 | Std.dev. | 50%  | 95%  | 99%  |  |
| All 40 analyses         | 0.17                                 | 0.15     | 0.14 | 0.46 | 0.66 | Half-normal (0.22) Half-Student- $t$ (8.2; 0.20) |
| 20 most recent analyses | 0.25                                 | 0.24     | 0.18 | 0.71 | 1.13 | Half-normal (0.31) Half-Student- $t$ (4.3; 0.26) |
| 10 most recent analyses | 0.45                                 | 0.45     | 0.32 | 1.32 | 2.17 | Half-normal (0.56) Half-Student- $t$ (3.8; 0.44) |
| 5 most recent analyses  | 0.65                                 | 0.90     | 0.38 | 2.14 | 4.28 | Half-normal (0.82) Half-Student- $t$ (2.7; 0.56) |

heterogeneity priors (see also Section 3.6 and Table 3). Despite the differences apparent in the predictive distributions (Table 3) and their corresponding parametric approximations (Table 4), all five resulting meta-analytic estimates still turn out very similar here, so one might pragmatically choose the simple half-normal prior here. The Bayesian estimate based on the weakly informative prior is more conservative, resulting in a wider interval (eg, 20% wider than for the half-normal (0.22) prior).

The different frequentist approaches suffer from the small sample size of only  $k = 2$  studies and appear to be either too optimistic or overly conservative. Since the heterogeneity point estimate is zero here, the frequentist intervals effectively correspond to “common-effect” analyses in this case. The HKSJ interval turns out roughly half as wide as the Normal interval, which may be considered counterintuitive.<sup>48</sup> The mKH interval on the other hand is extremely conservative (about six times as wide as the Normal interval), which again is not so uncommon for a meta-analysis of two studies only.<sup>5,49</sup>

### 3.9.2 | Differing amounts of historical data

When embarking on an investigation of “historical” meta-analyses as in the present example, an obvious question is *how many* meta-analyses would be required or sufficient. This is probably hard to answer in any generality, as it also very much depends on the sizes ( $k_j$ ) of involved meta-analyses, but we may at least shed some light by investigating variations of the example discussed above. To this end, besides the “full” set of  $N = 40$  studies, we restrict the data to smaller subsets including the most recent 20, 10, or 5 studies. Table 5 shows the resulting predictive distributions and corresponding approximations when considering increasingly smaller subsets of the total of 40 meta-analyses. One can see that when considering fewer heterogeneity estimates, the restricted data essentially are not able to “rule out” larger heterogeneity ranges, and so the resulting predictive distributions have larger means, medians and other quantiles. From the relative magnitude of mean and standard deviation (and the corresponding coefficient of variation) as well as from the Student- $t$  distributions’ degrees-of-freedom parameters, one can see that the predictive distributions also become increasingly heavy-tailed. The cases where the degrees-of-freedom parameter is large are those where a half-normal approximation may fit comparably well; when the degrees-of-freedom parameter is low, the half-Student- $t$  distribution may be substantially more accurate.

It is also worth noting that for few analyses (only 10 or 5 included), the corresponding half-normal approximation conveys very little information — a half-normal prior with scale 0.5 is commonly already seen as a very conservative choice

in this context,<sup>3,4</sup> and so the subsets of only 5 or 10 meta-analyses, corresponding to half-normal priors with larger scale, do not seem to add information beyond what may be assumed given already. Strictly speaking, such a-priori information could also be implemented in the hyperprior (see equations (4) or (7) in the analysis, which should then result in an implicit lower bound on the informativeness of the resulting predictive distribution. However, it may also be of interest to consider the different information sources in separation.

## 4 | DISCUSSION

In this paper, we described a method allowing to translate historical meta-analysis data into a prior distribution for the heterogeneity parameter in a subsequent meta-analysis. The approach may serve to quantify plausible ranges for a heterogeneity parameter, or simply as a cross-check whether some given prior specification appears to be over-optimistic or too conservative relative to historical data.

For illustration purposes, we considered the publicly available data set collected by Seide et al<sup>16,17</sup> as a working example. The inferred predictive distribution may be somewhat realistic, but should not be understood as a recommendation of any generality. For example, while the data might be expected to be relatively homogeneous already, many of the analyses related to mortality endpoints, which were also found to be the least heterogeneous type of outcome in the related analysis of the Cochrane Database of Systematic Reviews (CDSR) by Turner et al.<sup>20</sup> We expect the type of endpoint considered to possibly play a role for the choice of prior. Endpoint-specific priors might be set up, whereas, in view of pragmatic considerations, it may also be possible to agree on a (simpler) common setting for different endpoints (eg, odds ratios, relative risks, hazard ratios, or standardized mean differences).

The R code to reproduce the analyses shown here is included in the online supplement. We intend to subsequently apply this method to a comprehensive collection of meta-analyses published by the German Institute for Quality and Efficiency in Health Care (*Institut für Qualität und Wirtschaftlichkeit im Gesundheitswesen, IQWiG*) in order to systematically evaluate the empirical evidence on heterogeneity with respect to certain classes of effect measures or endpoints. The present results suggest that a half-normal distribution may serve the purpose well, so that the investigation would primarily yield estimates of half-normal scale parameters. These analyses could then form the basis for recommendations regarding the choice of prior in Bayesian random-effects meta-analyses for applications in health technology assessment (HTA).

While initially an obvious approach might have been to consider collections of heterogeneity *point estimates* for this purpose, it soon became obvious that these would only yield a very coarse indication of plausible heterogeneity ranges; a comparison of the predictive distribution and the distribution of point estimates for our example data is also illustrated in Appendix D. We then set up a simple extension of the NNHM to accommodate a heterogeneity distribution and to infer the predictive distribution that would be useful for prior specification. However, there may always be situations where heterogeneity estimates are available while the original studies are not;<sup>15</sup> in such cases, a collection of heterogeneity point estimates may still provide a reasonable approximation to the predictive distribution, as observed in the example. Overdispersion and bias of the heterogeneity estimates relative to the predictive distribution may be less of an issue when the original meta-analyses are reasonably large. Many model extensions or variations would be conceivable; for example, in view of binary endpoints, a binomial-normal hierarchical model similar to the one implemented by Turner et al<sup>20</sup> (which would allow to dispose of the normal approximation at the first model stage), or a model accommodating subgroups of studies as utilized by Rhodes et al<sup>19</sup> or Turner et al.<sup>20</sup>

Although the heterogeneity prior plays an important role in the setup of a meta-analysis, many other details are at least equally relevant, such as the study selection, choice of estimand, effect measure or the statistical model. Empirical information will usually only constitute one of several aspects to inform or contribute to prior specification; further aspects to be considered include a prior's tail behavior, or robustness or conservatism features.<sup>2,4</sup> While the availability of a coherent approach to accurately translate empirical data into a prior specification is convenient, we expect a rather rough summary of the resulting predictive distribution to commonly be sufficient. Unlike in usual systematic reviews, an exhaustive search may not be necessary and a smaller (but representative) sample may be sufficient.

## ACKNOWLEDGEMENT

Support from the *Deutsche Forschungsgemeinschaft (DFG)* is gratefully acknowledged (grant number FR 3070/3-1). Open Access funding enabled and organized by Projekt DEAL.

**CONFLICT OF INTEREST STATEMENT**

The authors have declared no conflict of interest.

**DATA AVAILABILITY STATEMENT**

The data that supports the findings of this study are publicly available,<sup>16</sup> R code allowing to reproduce calculations is provided in the supplementary material of this article.

**ORCID**

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

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

**REFERENCES**

- Hedges LV, Olkin I. *Statistical Methods for Meta-Analysis*. San Diego, CA, USA: Academic Press; 1985.
- Röver C. Bayesian random-effects meta-analysis using the bayesmeta R package. *J Stat Softw*. 2020;93(6):1-51. doi:10.18637/jss.v093.i06
- 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. doi:10.1002/jrsm.1217
- Röver C, Bender R, Dias S, et al. On weakly informative prior distributions for the heterogeneity parameter in Bayesian random-effects meta-analysis. *Res Synth Methods*. 2021;12(4):448-474. doi:10.1002/jrsm.1475
- 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. doi:10.1002/jrsm.1297
- Jaynes ET. Prior probabilities. *IEEE Tran Syst Sci Cybern*. 1968;SEC-4(3):227-241. doi:10.1109/TSSC.1968.300117
- Kass RE, Wasserman L. The selection of prior distributions by formal rules. *J Am Stat Assoc*. 1996;91(435):1343-1370. doi:10.1080/01621459.1996.10477003
- Zondervan-Zwijnenburg M, Peeters M, Depaoli S, Schoot R. Where do priors come from? Applying guidelines to construct informative priors in small sample research. *Res Human Develop*. 2017;14(4):305-320. doi:10.1080/15427609.2017.1370966
- Best N, Dallow N, Montague T. Prior elicitation. In: Lesaffre E, Baio G, Boulanger B, eds. *Bayesian Methods in Pharmaceutical Research*. Boca Raton: Chapman & Hall / CRC; 2020, 109:87-109.
- Higgins JPT, Whitehead A. Borrowing strength from external trials in a meta-analysis. *Stat Med*. 1996;15(24):2733-2749.
- Pullenayegum EM. An informed reference prior for between-study heterogeneity in meta-analyses of binary outcomes. *Stat Med*. 2011;30(26):3082-3094. doi:10.1002/sim.4326
- 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. doi:10.1093/ije/dys041
- Steel P, Kammeyer-Mueller J, Paterson TA. Improving the meta-analytic assessment of effect size variance with an informed Bayesian prior. *J Manage*. 2015;41(2):718-743. doi:10.1177/0149206314551964
- Langan D, Higgins JPT, Simmonds M. An empirical comparison of heterogeneity variance estimators in 12894 meta-analyses. *Res Synth Methods*. 2015;6(2):195-205. doi:10.1002/jrsm.1140
- van Erp S, Verhagen J, Grasman RPPP, Wagenmakers E-J. Estimates of between-study heterogeneity for 705 meta-analyses reported in *Psychological Bulletin* from 1990–2013. *J Open Psychol Data*. 2017;5(1):4. doi:10.5334/jopd.33
- Seide SE, Christian R, Tim F. Meta-analysis data extracted from IQWiG publications. Göttingen Research Online. 2018. doi:10.25625/BWYBNK
- 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. doi:10.1186/s12874-018-0618-3
- Günhan BK, Röver C, Friede T. Random-effects meta-analysis of few studies involving rare events. *Res Synth Methods*. 2020;11(1):74-90. doi:10.1002/jrsm.1370
- 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. doi:10.1016/j.jclinepi.2014.08.012
- 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. doi:10.1002/sim.6381
- Liberati A, Altman DG, Tetzlaff J, et al. The PRISMA statement for reporting systematic reviews and meta-analyses of studies that evaluate health care interventions: explanation and elaboration. *PLoS Med*. 2009;6(7):e1000100. doi:10.1371/journal.pmed.1000100
- Gelman A. Bayes, Jeffreys, prior distributions and the philosophy of statistics. *Stat Sci*. 2009;24(2):176-178.
- Cole SR, Chu H, Greenland S. Maximum likelihood, profile likelihood, and penalized likelihood: a primer. *Am J Epidemiol*. 2014;179(2):252-260. doi:10.1093/aje/kwt245
- Williams DR, Rast P, Bürkner P-C. Bayesian meta-analysis with weakly informative prior distributions. *PsyArXiv*. 2018. doi:10.17605/OSF.IO/7TBRM
- 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. doi:10.1111/biom.12242

26. Lindsay BG. Mixture models: theory, Geometry and Applications. *NSF-CBMS Regional Conference Series in Probability and Statistics*. Vol 5. Hayward, CA, USA: Institute of Mathematical Statistics; 1995.
27. Gelman A. Prior distributions for variance parameters in hierarchical models. *Bayesian Anal.* 2006;1(3):515-534. doi:10.1214/06-BA117A
28. Gelman A, Carlin JB, Stern H, Dunson DB, Vehtari A, Rubin DB. *Bayesian Data Analysis*. 3rd ed. Boca Raton: Chapman & Hall / CRC; 2014.
29. Plummer M. JAGS: a program for analysis of Bayesian graphical models using Gibbs sampling. In: Hornik K, Leisch F, Zeileis A, eds. *Proceedings of the 3rd International Workshop on Distributed Statistical Computing (DSC 2003)*. Vienna, Austria: R Foundation for Statistical Computing, Austrian Association for Statistical Computing (AASC); 2003.
30. Plummer M. Rjags: Bayesian graphical models using MCMC2008. *R Package*. <http://cran.r-project.org/package=rjags>. Accessed April 12, 2022.
31. Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A. Bayesian measures of model complexity and fit. *J Royal Stat Soc B*. 2002;64(4):583-639.
32. Weber S, Li Y, Seaman J, Kakizume T, Schmidli H. Applying meta-analytic-predictive priors with the R Bayesian evidence synthesis tools. *J Stat Softw.* 2021;100(19).
33. Spiegelhalter DJ, Abrams KR, Myles JP. *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. Chichester, UK: John Wiley & Sons; 2004.
34. 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. doi:10.1002/1097-0258(20001230)19:24<3359::AID-SIM710>3.0.CO;2-N
35. 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. doi:10.1177/0272989X13485157
36. Gelman A, Hill J. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge: Cambridge University Press; 2007.
37. Gelman A, Hwang J, Vehtari A. Understanding predictive information criteria for Bayesian models. *Stat Comput.* 2014;24:997-1016. doi:10.1007/s11222-013-9416-2
38. Meyer R. Deviance information criterion (DIC). In: Balakrishnan N, Colton T, Everitt B, Piegorisch W, Ruggeri F, Teugels JL, eds. *Wiley StatsRef: Statistics Reference Online*. John Wiley & Sons; 2016. doi:10.1002/9781118445112.stat07878
39. Plummer M. Penalized loss functions for Bayesian model comparison. *Biostatistics.* 2008;9(3):523-539. doi:10.1093/biostatistics/kxm049
40. Gelman A. Bayesian model building by pure thought: some principles and examples. *Stat Sinica.* 1996;6(1):215-232.
41. O'Hagan A, Pericchi L. Bayesian heavy-tailed models and conflict resolution: a review. *Brazilian J Probab Stat.* 2012;26(4):372-401. doi:10.1214/11-BJPS164
42. Jefferys WH, Berger JO. Ockham's razor and Bayesian analysis. *Am Sci.* 1992;80(1):64-72.
43. Röver C, Wandel S, Friede T. Model averaging for robust extrapolation in evidence synthesis. *Stat Med.* 2018;38(4):674-694. doi:10.1002/sim.7991
44. Institut für Qualität und Wirtschaftlichkeit im Gesundheitswesen (IQWiG). Verfahren zur Lungenvolumenreduktion beim schweren Lungenemphysem IQWiG-Berichte – Nr. 487 (N14-04)2017.
45. Deslée G, Mal H, Dutau H, et al. Lung volume reduction coil treatment vs usual care in patients with severe emphysema: the REVOLENS randomized clinical trial. *Jama.* 2016;315(2):175-184. doi:10.1001/jama.2015.17821
46. Sciruba FC, Criner GJ, Strange C, et al. Effect of endobronchial coils vs usual care on exercise tolerance in patients with severe emphysema: the RENEW randomized clinical trial. *Jama.* 2016;315(20):2178-2189. doi:10.1001/jama.2016.6261
47. Garner JL, Shah PL. Lung volume reduction in pulmonary emphysema. *Semin Respir Crit Care Med.* 2020;41(06):874-885. doi:10.1055/s-0040-1702192
48. Röver C, Knapp G, Friede T. Hartung-Knapp-Sidik-Jonkman approach and its modification for random-effects meta-analysis with few studies. *BMC Med Res Methodol.* 2015;15. doi:10.1186/s12874-015-0091-1
49. 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. doi:10.1002/bimj.201500236
50. Psarakis S, Panaretos J. The folded *t* distribution. *Commun Stat—Theory Methods.* 1990;19(7):2717-2734. doi:10.1080/03610929008830342

**How to cite this article:** Röver C, Sturtz S, Lilienthal J, Bender R, Friede T. Summarizing empirical information on between-study heterogeneity for Bayesian random-effects meta-analysis. *Statistics in Medicine*. 2023;42(14):2439-2454. doi: 10.1002/sim.9731

## APPENDIX A. JAGS CODE

The code shown in Table A1 defines the JAGS implementation of the model described in Section 2.<sup>29,30</sup> The data, a total of  $N$  effect estimates from  $n$  meta-analyses are provided in terms of the vectors of effect estimates  $y$  and variances (squared standard errors)  $v$ . Individual studies are allocated to meta-analyses via the vector of index variables  $id$ . The effects' vague prior is given through its mean `effectPriorMean` and standard deviation `effectPriorSD`.

TABLE A1 JAGS code implementing the model described in Appendix A.

```

model {
  # the normal-normal hierarchical model (NNHM):
  for (i in 1:N) { # loop over (N) individual studies:
    y[i] ~ dnorm(mu[id[i]], pow(v[i]+pow(tau[id[i]]), 2), -1)
  }
  for (i in 1:n) { # loop over (n) meta-analyses:
    mu[i] ~ dnorm(effectPriorMean, pow(effectPriorSD, -2))
  }
  # the heterogeneity prior:
  for (i in 1:n) { # loop over (n) meta-analyses:
    # half-normal distribution for tau:
    tau[i] ~ dnorm(0.0, pow(tauScale, -2)) T(0,)
  }
  tauScale ~ dunif(0.0, tauScalePriorMax)
  tauPrediction ~ dnorm(0.0, pow(tauScale, -2)) T(0,)
}

```

The model for the heterogeneity standard deviations  $\tau$  is eventually specified in the final loop; here this is a half-normal distribution with a scale parameter ( $\tau$ Scale). The scale parameter is estimated from the data after specification of a vague prior, a uniform distribution with upper bound  $\tau$ ScalePriorMax. Eventually, samples from the predictive distribution ( $\tau$ Prediction) are also generated along the way.

## APPENDIX B. SCALE MIXTURE PARAMETRIZATION

### B.1 Normal scale mixture

Suppose a half-normal distribution's scale parameter  $s$  has mean  $\mu(s)$ , variance  $\sigma^2(s)$  and coefficient of variation  $c_v(s) = \frac{\sqrt{\sigma^2(s)}}{\mu(s)}$ . Then the resulting mixture distribution may be approximated by a half-Student- $t$  distribution by assuming that the scale followed a *scaled inverse  $\chi$  distribution* with matching moments.<sup>4</sup>

The matching half-Student- $t$  distribution's parameters (degrees-of-freedom  $\nu_t$  and scale  $s_t$ ) then depend on  $c_v(s)$  and  $\mu(s)$ . First, the degrees-of-freedom may be solved for numerically based on the coefficient of variation  $c_v(s)$  by equating the scaled inverse  $\chi$  distribution's coefficient of variation (which does not depend on its scale parameter) to match  $c_v(s)$ . Then the half-Student- $t$  distribution's scale results as  $s_t = \frac{\mu(s)}{E_{\text{Inv-}\chi}(\nu_t, \sqrt{\nu_t})}$ , where  $E_{\text{Inv-}\chi}(\nu_t, \sqrt{\nu_t})$  is the expectation of a scaled inverse  $\chi$  distribution with degrees-of-freedom  $\nu_t$  and scale  $\sqrt{\nu_t}$ .<sup>4</sup>

### B.2 Exponential scale mixture

Suppose an exponential distribution's scale (or *inverse rate*) parameter  $s$  has mean  $\mu(s)$ , variance  $\sigma^2(s)$  and coefficient of variation  $c_v(s) = \frac{\sqrt{\sigma^2(s)}}{\mu(s)}$ . Then the resulting mixture distribution may be approximated by a Lomax distribution by assuming that the scale followed an inverse gamma distribution with matching moments.<sup>4</sup>

The Lomax distribution is parameterized in terms of shape  $\alpha$  and scale  $\lambda$ , with expectation  $\frac{\lambda}{\alpha-1}$  and variance  $\frac{\lambda^2 \alpha}{(\alpha-1)^2(\alpha-2)}$ . The corresponding parameter values follow from  $c_v(s)$  and  $\mu(s)$  and are given by  $\alpha = 2 + \frac{1}{c_v(s)^2}$  and  $\lambda = \mu(s)(1 + \frac{1}{c_v(s)^2})$ .<sup>4</sup>

### B.3 Log-normal distribution

Another obvious and common distribution for the heterogeneity  $\tau$  is the log-normal distribution, resulting from modeling the logarithmic heterogeneity parameter ( $\log(\tau)$ ) based on location ( $\mu$ ) and scale ( $\sigma$ ). In the context of the other models discussed here, it is useful to consider an alternative, less common parametrization of the log-normal distribution, namely, based on the exponentiated mean parameter  $\vartheta = \exp(\mu)$ . Re-writing the log-normal density  $f(x)$  in this form, it becomes apparent that  $\vartheta$  constitutes a *scale parameter* here, while  $\sigma$  is a *shape parameter*:

$$f(x) = \frac{1}{\vartheta} \frac{1}{(x/\vartheta) \sigma \sqrt{2\pi}} \exp\left(-\frac{(\log(x/\vartheta))^2}{2\sigma^2}\right). \quad (\text{B1})$$

This allows to treat the log-normal distribution analogously to the other distributions discussed for example, in Section 2. In this parametrization, the distribution is characterized by

|                          |  |
|--------------------------|--|
| Median                   | $\vartheta$  |
| Mean                     | $\vartheta \sqrt{\exp(\sigma^2)}$                      |
| Variance                 | $\vartheta^2 \exp(\sigma^2) (\exp(\sigma^2) - 1)$      |
| Standard deviation       | $\vartheta \sqrt{\exp(\sigma^2) (\exp(\sigma^2) - 1)}$ |
| Coefficient of variation | $\sqrt{\exp(\sigma^2) - 1}$                            |

When heterogeneity is modelled using a log-normal distribution in the extended NNHM (see Section 2), the same distribution may also be a reasonable choice for approximating the (usually more dispersed) predictive distribution (of  $\tau^*$ ). While a *scale mixture* of log-normal distributions does not have a simple analytic form, this may still be motivated by considering the logarithmic scale: implementing uncertainty in  $\mu = \log(\vartheta)$  using an additive normal term again leads to a log-normal distribution (with a larger  $\sigma$  parameter). Strictly speaking, the predictive distribution for a normal distribution with unknown mean and variance is a Student- $t$  distribution,<sup>28</sup> however, fitting a three-parameter distribution may well be an exaggeration in the present context.

Log-normal distributions have been used for modeling predictive distributions for example, by Turner et al<sup>20</sup> Similarly, Rhodes et al<sup>19</sup> used a log-Student- $t$ -distribution, albeit with a pre-specified, fixed degrees-of-freedom parameter ( $\nu = 5$ ).

### APPENDIX C. MOMENT ESTIMATION FOR THE HALF-STUDENT-T DISTRIBUTION

The first two moments (expectation and variance) of a half-Student- $t$  distribution with degrees-of-freedom  $\nu$  and scale  $\sigma$  are given by

$$E[X] = 2\sigma \sqrt{\frac{\nu}{\pi}} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2}) (\nu-1)} \quad (\text{for } \nu > 1) \quad (\text{C1})$$

and

$$\text{Var}(X) = \sigma^2 \left( \frac{\nu}{\nu-2} - \frac{4\nu}{\pi (\nu-1)^2} \left( \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})} \right)^2 \right) \quad (\text{for } \nu > 2). \quad (\text{C2})$$

Its coefficient of variation hence results as<sup>50</sup>

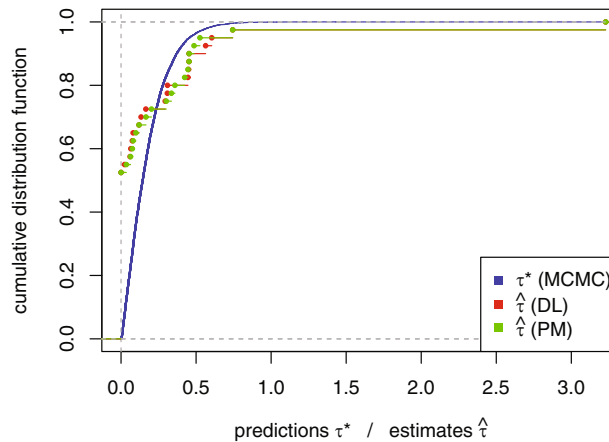
$$\frac{\sqrt{\text{Var}(X)}}{E[X]} = \sqrt{\frac{\pi (\nu-1)^2}{4 (\nu-2)} \left( \frac{\Gamma(\frac{\nu}{2})}{\Gamma(\frac{\nu+1}{2})} \right)^2 - 1} \quad (\text{for } \nu > 2), \quad (\text{C3})$$

which is independent of the scale  $\sigma$  and  $> \sqrt{\frac{\pi}{2} - 1} \approx 0.75$ . A half-Student- $t$  distribution's moment estimates may hence be computed by first matching the coefficient of variation (given that it is greater than 0.75) to yield a degrees-of-freedom estimate, and subsequently matching the expectation for a scale estimate. Solving equation (C3) for the degrees-of-freedom  $\nu$  may be done numerically.

For example, suppose a coefficient of variation of 0.8 and an expectation of 0.5 were aimed for. Numerically solving equation (C3) (eg, utilizing R's "uniroot()" function) yields a degrees-of-freedom estimate of  $\hat{\nu} = 13.6$ . A half-Student- $t$  distribution with  $\nu = 13.6$  degrees of freedom (and scale 1) would have an expectation of 0.85. The scale's moment estimator then equals  $\hat{\sigma} = \frac{0.5}{0.85} = 0.59$ , that is, the half-Student- $t$  distribution with  $\nu = 13.6$  and  $\sigma = 0.59$  has an expectation of 0.5 and coefficient of variation 0.8.

## APPENDIX D. COMPARISON TO FREQUENTIST POINT ESTIMATES

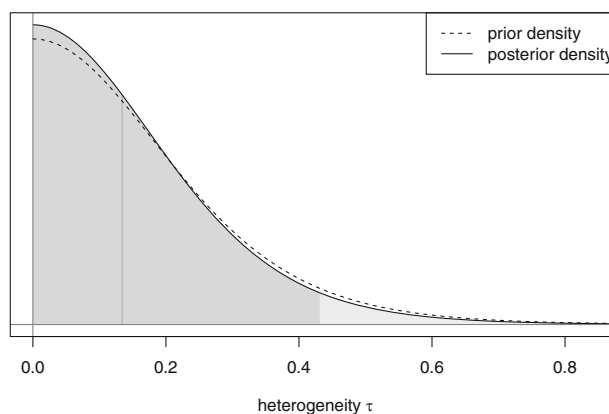
Since a heterogeneity *estimate* is associated with additional variance due to estimation uncertainty, one would expect the distribution of estimates to be overdispersed relative to the distribution of true values. Figure D1 contrasts the predictive distribution (of  $\tau^*$ , as also shown in Figure 2) with the distributions of 40 heterogeneity estimates from the example data. The DerSimonian-Laird (DL) and Paule-Mandel (PM) estimates behave similarly here, a sizeable fraction turns out as zero, and the overall picture appears to confirm the expected behavior; in particular, it becomes evident that consideration of a mere collection heterogeneity *estimates* is of only limited use for prior specification. For example, the median  $\hat{\tau}$  value would be at zero, while the average would be at 0.21 (for both DL and PM) and would to a large extent be dominated by a single outlier.



**FIGURE D1** Comparison of predictive distribution vs. the empirical distribution of 40 heterogeneity *point estimates* for the example data set.

## APPENDIX E. ANALYSES OF EXAMPLE DATA

Figure E1 illustrates the prior and posterior densities corresponding to the analysis also shown in Figure 3. As to be expected, the posterior is very similar to the prior (a half-Student-*t* distribution with 8.2 degrees of freedom and scale 0.20) in this case of only  $k = 2$  studies.



**FIGURE E1** Prior and posterior densities for the analysis shown in Figure 3. The vertical line indicates the posterior median, and the dark grey area shows the 95% credible interval.