

RESEARCH PAPER

Bounds for the weight of external data in shrinkage estimation

Christian Röver  | Tim Friede 

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

Correspondence

Christian Röver, Department of Medical Statistics, University Medical Center Göttingen, Humboldtallee 32, 37073 Göttingen, Germany.
Email: christian.roever@med.uni-goettingen.de

Funding information

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



This article has earned an open data badge “**Reproducible Research**” for making publicly available the code necessary to reproduce the reported results. The results reported in this article could fully be reproduced.

Abstract

Shrinkage estimation in a meta-analysis framework may be used to facilitate dynamical borrowing of information. This framework might be used to analyze a new study in the light of previous data, which might differ in their design (e.g., a randomized controlled trial and a clinical registry). We show how the common study weights arise in effect and shrinkage estimation, and how these may be generalized to the case of Bayesian meta-analysis. Next we develop simple ways to compute bounds on the weights, so that the contribution of the external evidence may be assessed a priori. These considerations are illustrated and discussed using numerical examples, including applications in the treatment of Creutzfeldt–Jakob disease and in fetal monitoring to prevent the occurrence of metabolic acidosis. The target study’s contribution to the resulting estimate is shown to be bounded below. Therefore, concerns of evidence being easily overwhelmed by external data are largely unwarranted.

KEYWORDS

Bayesian statistics, between-study heterogeneity, inverse-variance weights, random-effects meta-analysis, shrinkage estimation

1 | INTRODUCTION

In some situations, it is useful to support an estimate using additional external evidence, for example, when a small study in the context of a rare disease may be supplemented with data from a clinical registry or electronic health records, or when the result from a meta-analysis may be backed by an analysis in a similar field, for example, a related but somewhat different population. The involved data contributions then take on different roles, namely, that of a *source* (the external data) and a *target* (the data of primary interest). *Dynamic borrowing* refers to the class of approaches where the apparent, empirical similarity, or compatibility of the source and the target is taken into account when judging to what degree the two should be lumped together (Röver & Friede, 2020). Such approaches may be implemented, for example, via hierarchical models or informative priors; both are actually equivalent to some degree in the context of the *normal–normal hierarchical model* (NNHM) (Schmidli et al., 2014). Similarly, closely related (or partly equivalent) approaches are given by the *bias allowance* framework (Welton et al., 2012) or the *power prior* framework (Ibrahim & Chen, 2000). A recent example of such an approach is given by the EARLY PRO-TECT trial in pediatric Alport disease, where data from a randomized controlled trial (RCT) were supported by source data from an open-label arm and a clinical registry (Gross et al., 2020).

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.

© 2021 The Authors. *Biometrical Journal* published by Wiley-VCH GmbH

In the context of dynamic borrowing within the NNHM framework, the flow of information is quite commonly illustrated by quoting *weights* of data sources as these are combined to a joint estimate. As the eventual estimate may be expressed as a weighted average of the input data, the corresponding weights are a useful means of quantifying the studies' contributions to or influence on the eventual result (Hartung et al., 2008; Hedges & Olkin, 1985). Analogous weights arise for shrinkage estimates (Raudenbush & Bryk, 1985; Robinson, 1991; Viechtbauer, 2010), and, as we will show below, also in the Bayesian paradigm with prior distributions on effect and heterogeneity parameters.

When combining originally separate datasets in a meta-analysis or using shrinkage estimation, there sometimes is concern that evidence from the target data may be overwhelmed by a much larger set of source data, for example, when combining a small RCT with a large clinical registry or routine data (e.g., electronic health records) (Weber et al., 2018). In such cases, it is instructive to explicate the notion of study contributions by considering their weights. Again, we can see the dynamic nature of the approach in the changing weight of external data with varying data compatibility or discrepancy. It turns out that within the Bayesian framework we can determine the *minimum weight* of the target study (the RCT in the above example) a priori for a given analysis, and with that we are able to provide more insights into the general behavior of the meta-analysis procedure. The derived formulas show shrinkage estimation to behave reasonably and also predictably.

In the following, we will review the NNHM, and show how “study weights” arise in effect and shrinkage estimation and how the concept may be extended to the Bayesian framework. Then we take a closer look at the weights' properties and show how these may be bounded across possible prior settings and/or data realizations. The arguments are illustrated by a numerical study, and the ideas are employed in two example applications involving the joint analysis of a “small” target and a “large” source study, as well as two equally sized studies. Due to the few-studies setup (Friede et al., 2017; Röver & Friede, 2020), we will be focusing on Bayesian methods and only in between point out some connections to common analogous frequentist results. We close with a discussion of the findings and their practical implications.

2 | THE NORMAL–NORMAL HIERARCHICAL MODEL

The NNHM models a set of k estimates y_i and their standard errors σ_i as

$$y_i | \theta_i, \sigma_i \sim \text{Normal}(\theta_i, \sigma_i^2), \quad (1)$$

where θ_i are the *study-specific effects*. The θ_i are not necessarily identical for all studies, but they are also associated with a certain amount of variation, expressed as

$$\theta_i | \mu, \tau \sim \text{Normal}(\mu, \tau^2). \quad (2)$$

The mean parameter μ is the *overall mean effect*, while τ denotes the *between-study variability (heterogeneity)*. As noted elsewhere (Hartung et al., 2008; Hedges & Olkin, 1985; Röver, 2020), marginalizing over the parameters θ_i , the model may be written as

$$y_i | \mu, \tau, \sigma_i \sim \text{Normal}(\mu, \sigma_i^2 + \tau^2). \quad (3)$$

The NNHM is a random-effects (RE) model, which in the special case of $\tau = 0$ reduces to a fixed-effect (FE) (or common-effect) model. It provides a good approximation for many types of effect measures where measurement uncertainty and between-study variability may be assumed to be (approximately) normally distributed (Jackson & White, 2018). Data analysis may then aim at estimating the overall effect μ or study-specific effects θ_i (“shrinkage estimation”); in the present investigation, we will mostly be concerned with the latter.

In the following, we will denote vectors of effect estimates (y_1, \dots, y_k) and their standard errors $(\sigma_1, \dots, \sigma_k)$ by \vec{y} and $\vec{\sigma}$, respectively. Furthermore, we will be mostly concerned with the special case of only two studies ($k = 2$) and a noninformative (improper) uniform prior for the overall effect ($p(\mu) \propto 1$).

3 | STUDY WEIGHTS

3.1 | Conditional weights

Assuming an (improper) uniform prior for the overall effect μ , the conditional posterior distribution of μ (given τ) is normal with mean

$$\tilde{\mu}(\tau) = E[\mu|\tau, \vec{y}, \vec{\sigma}] = \sum_{i=1}^k w_i(\tau) y_i, \quad (4)$$

where the *inverse-variance (IV) weights* $w_j(\tau)$ are given by

$$w_j(\tau) = \frac{\frac{1}{\sigma_j^2 + \tau^2}}{\sum_{i=1}^k \frac{1}{\sigma_i^2 + \tau^2}} \quad (5)$$

as in the frequentist framework (Friede et al., 2017; Hartung et al., 2008; Hedges & Olkin, 1985). A similar formula also applies to a normal prior (Röver, 2020). These two (conditionally conjugate) priors are computationally simple, readily motivated, and, because of that reason, probably also the most commonly used priors for μ for this model. Informative priors might, for example, be motivated by general plausibility considerations or empirical data (Günhan et al., 2020) or determined through expert elicitation (Best et al., 2020; Hampson et al., 2014).

The conditional posterior of the study-specific effect θ_j (the *shrinkage estimate*) is also normal with mean $\tilde{\theta}_j(\tau)$ depending on y_j and $\tilde{\mu}(\tau)$, namely

$$\tilde{\theta}_j(\tau) = E[\theta_j|\tau, \vec{y}, \vec{\sigma}] = b_j(\tau) y_j + (1 - b_j(\tau)) \tilde{\mu}(\tau), \quad (6)$$

where the corresponding weight (Röver, 2020; Wandel et al., 2017) is

$$b_j(\tau) = \frac{\frac{1}{\sigma_j^2}}{\frac{1}{\sigma_j^2} + \frac{1}{\tau^2}}. \quad (7)$$

The formulation in (6) shows to which degree the estimate is *shrunk* towards the common overall mean $\tilde{\mu}(\tau)$ (depending on the amount of heterogeneity). Equation (6) may be rewritten as

$$\tilde{\theta}_j(\tau) = \left[b_j(\tau) + (1 - b_j(\tau)) w_j(\tau) \right] y_j + \sum_{i \neq j} \left[(1 - b_j(\tau)) w_i(\tau) \right] y_i \quad (8)$$

$$= c_{jj}(\tau) y_j + \sum_{i \neq j} c_{ij}(\tau) y_i \quad (9)$$

$$= \sum_{i=1}^k c_{ij}(\tau) y_i, \quad (10)$$

so that the actual *shrinkage weights* $c_{ij}(\tau)$ (the weight of the i th study in the estimate of the j th study) become more explicit. In the special case of only two studies ($k = 2$), the coefficients $c_{ij}(\tau)$ simplify to

$$c_{11}(\tau) = \frac{\sigma_2^2 + 2\tau^2}{\sigma_1^2 + \sigma_2^2 + 2\tau^2}, \quad c_{12}(\tau) = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2 + 2\tau^2}, \quad (11)$$

and analogously for c_{22} and c_{21} .

The conditional mean $\tilde{\mu}(\tau)$ commonly also arises in frequentist approaches as an overall effect estimator, where usually a heterogeneity estimate $\hat{\tau}$ is plugged in for τ (Hartung et al., 2008; Hedges & Olkin, 1985). Similarly, plug-in estimates of $\tilde{\theta}_j(\tau)$ are widely used and commonly known as “best linear unbiased prediction” (BLUP) (Raudenbush & Bryk, 1985; Robinson, 1991; Viechtbauer, 2010). The weights ($w_j(\tau)$ or $c_{ij}(\tau)$) are then often quoted along with the results in order to illustrate the individual studies’ contributions to the overall result. Note that while weights may be appealing, they still constitute an ultimately somewhat heuristic notion of the concept of *a study’s contribution*, as these only relate to the posterior expectation. A more complete picture might actually be obtained by considering the corresponding meta-analytic-predictive (MAP) prior (Schmidli et al., 2014), which comprehensively describes the information conveyed by the source study.

3.2 | Marginal weights

In a Bayesian multiparameter model, the *conditional* expectations (of effects μ or θ_j) as derived above are commonly of limited interest; what is usually more interesting are the *marginal* posterior expectations, as these refer to the posterior distribution integrated over other parameters such as the heterogeneity τ in the considered model. Marginal posterior expectations here result from the conditional expectations as expected values with respect to the heterogeneity’s marginal posterior distribution $p(\tau|\vec{y}, \vec{\sigma})$, that is,

$$E[\mu|\vec{y}, \vec{\sigma}] = E_{p(\tau|\vec{y}, \vec{\sigma})} [E[\mu|\tau, \vec{y}, \vec{\sigma}]] \quad \text{and} \quad E[\theta_j|\vec{y}, \vec{\sigma}] = E_{p(\tau|\vec{y}, \vec{\sigma})} [E[\theta_j|\tau, \vec{y}, \vec{\sigma}]]. \quad (12)$$

In both cases, the conditional expectations result in convex combinations of the form $\sum_i \alpha_i(\tau) y_i$ (see Equations 4 and 10). For convex (or, more generally, linear) combinations, we may rewrite the expectations as

$$E_{p(\tau|\vec{y}, \vec{\sigma})} \left[\sum_i \alpha_i(\tau) y_i \right] = \sum_i E_{p(\tau|\vec{y}, \vec{\sigma})} [\alpha_i(\tau)] y_i, \quad (13)$$

so that it becomes apparent that the marginal expectation may again be expressed as a weighted average of the effects y_i , where the study weights now arise as the *posterior expected weights*. These constitute straightforward generalizations of the common *conditional* weights to the Bayesian context. The weights can be obtained from one-dimensional integrals (expectations) involving the heterogeneity’s marginal posterior distribution and may easily be computed numerically; they are returned by default by the `bayesmeta` R package (Röver, 2015, 2020).

3.3 | Properties

For $\tau = 0$, the NNHM reduces to the FE model, in which all study effects θ_i coincide with the overall mean μ . As τ is varied between the two extremes of $\tau = 0$ and $\tau \rightarrow \infty$, several effects may be observed for the conditional weights:

- The IV-weights $w_j(\tau)$ move (not necessarily monotonically) from “fixed-effect” weights $w_j(0) = \frac{\frac{1}{\sigma_j^2}}{\sum_i \frac{1}{\sigma_i^2}}$ that depend on the study’s precision towards “average” weights $w_i(\infty) = \frac{1}{k}$, where all studies have the same weight.
- The weights $b_j(\tau)$ increase monotonically from 0 towards 1.
- The shrinkage weights $c_{jj}(\tau)$ (the contribution of the j th study to its own shrinkage estimate) increase monotonically from the FE weight towards 1.

For the conditional expectations, this implies:

- The conditional effect estimate $\tilde{\mu}(\tau)$ moves from the FE estimate towards an unweighted average.
- The conditional shrinkage estimates $\tilde{\theta}_j(\tau)$ move from the FE estimate towards the ‘unpooled’ original estimates y_j .

TABLE 1 Limit values of relevant meta-analysis shrinking weights and estimates, under homogeneity ($\tau = 0$) and large heterogeneity ($\tau \rightarrow \infty$)

Expression		Heterogeneity extreme	
		$\tau = 0$	$\tau \rightarrow \infty$
Inverse-variance (IV) weight $w_j(\tau)$	(4)	FE weight $\left(\frac{\sigma_j^{-2}}{\sum_i \sigma_i^{-2}}\right)$	Average weight $\left(\frac{1}{k}\right)$
Weight $b_j(\tau)$	(6)	0	1
Shrinkage weight $c_{jj}(\tau)$	(10)	FE weight	1
Conditional overall mean $\tilde{\mu}(\tau)$	(4)	FE estimate	Unweighted average
Conditional shrinkage estimate $\tilde{\delta}_j(\tau)$	(6)	FE estimate	Original estimate y_j

These effects are also summarized in Table 1. Posterior expectations of the weights depend on the heterogeneity’s posterior distribution $p(\tau|\vec{y}, \vec{\sigma})$. For a uniform effect prior, a given heterogeneity prior $p(\tau)$ and standard errors σ_i , the posterior density is given by

$$p(\tau|\vec{y}, \vec{\sigma}) \propto p(\tau) f_{\vec{\sigma}}(\tau) g_{\vec{y}}(\tau) \tag{14}$$

with

$$g_{\vec{y}}(\tau) = \exp\left(-\frac{1}{2} \left[\frac{(y_1 - \tilde{\mu}(\tau))^2}{\sigma_1^2 + \tau^2} + \frac{(y_2 - \tilde{\mu}(\tau))^2}{\sigma_2^2 + \tau^2} \right]\right) = \exp\left(-\frac{1}{2} \frac{(y_2 - y_1)^2}{\sigma_1^2 + \sigma_2^2 + 2\tau^2}\right) \tag{15}$$

(see, e.g., Equation in Röver, 2020), where $p(\tau)$ is the heterogeneity’s prior density, and $f_{\vec{\sigma}}(\tau)$ is a lengthier term involving τ and $\vec{\sigma}$. From (15), one can see that the heterogeneity’s posterior depends on the data (y_1, y_2) only via the absolute difference $|y_2 - y_1|$, which in a sense constitutes the “empirical” or “observed” amount of heterogeneity, through the exponential term $g_{\vec{y}}(\tau)$.

A closer look at $g_{\vec{y}}(\tau)$ shows that it always remains between zero and one ($0 < g_{\vec{y}}(\tau) \leq 1$). For $y_2 = y_1$, it is constant at $g_{\vec{y}}(\tau) = 1$. For a given difference $|y_2 - y_1| > 0$, it takes its minimum at $\tau = 0$ and then increases monotonically with τ . For any given τ , it decreases monotonically in $|y_2 - y_1|$. One might think of $g_{\vec{y}}(\tau)$ as “ruling out” smaller τ values in the heterogeneity posterior and pushing the posterior mode towards higher τ values as $|y_2 - y_1|$ increases.

The functional form of the posterior (15) implies that for increasing $|y_2 - y_1|$ the resulting marginal heterogeneity posterior becomes *stochastically larger* (Shaked & Shanthikumar, 2007); see also the Appendix for a derivation. When varying the prior distribution $p(\tau)$ in (14), we may to some extent also predict the effect on the heterogeneity posterior: in particular, choosing a stochastically larger heterogeneity prior will imply a stochastically larger posterior as well (see also the Appendix).

4 | BOUNDS FOR THE STUDY WEIGHTS

4.1 | Lower bounds

The above conditions imply that we can derive bounds for the shrinkage weights. As mentioned previously, concerns are sometimes raised that the target estimates may be overwhelmed by the source data, that is, that certain weights may become *too small* (Weber et al., 2018). In the following, we will describe the conditions under which we can derive *lower bounds* on weights, that is, where we can make sure that weights remain above a certain minimum. Important consequences for the weights, valid quite generally or for certain heterogeneity priors $p(\tau)$, are derived below. Note that while we assume the standard errors σ_i to be given (a common assumption to be made in meta-analysis or study design considerations), the data (estimates y_i) or the prior ($p(\tau)$) may be varied.

4.2 | A study’s minimum contribution to its own shrinkage estimate: the “FE weight”

The (conditional) shrinkage weight $c_{jj}(0)$, that is the j th study’s contribution to its own shrinkage estimate evaluated at $\tau = 0$, constitutes a lower bound for the posterior mean weights. Any heterogeneity prior $p(\tau)$ may attach prior proba-

bility to τ values larger than zero, for which the weights are only increasing. These “FE weights” may simply be computed as the common study weights in a fixed-effect meta-analysis. This property holds independent of the actual data (y_i) or the heterogeneity prior ($p(\tau)$).

4.3 | Minimum posterior mean shrinkage weight: The “coincidence weight”

For any prior distribution $p(\tau)$, the *coincidence* case of $y_1 = y_2$ is the data realization yielding the lowest possible posterior mean shrinkage weight. Any data with $|y_2 - y_1| > 0$ will imply a stochastically larger heterogeneity posterior that will (due to monotonicity of weights $c_{jj}(\tau)$ as a function of τ) lead to larger posterior mean shrinkage weights. The coincidence weights may simply be computed by performing the meta-analysis with the data (y_1 and y_2) substituted by two identical numbers. This property holds independent of the data (y_i) and for any given heterogeneity prior ($p(\tau)$).

4.4 | Stochastically ordered priors and their posterior mean weights

Considering stochastically ordered families of heterogeneity priors allows to vary the posterior mean shrinkage weight. For properly chosen stochastically smaller priors, the posterior mean may approach the FE weight, while for stochastically larger priors the posterior mean weight may approach 100%. An obvious, simple way to yield a stochastically ordered family of prior distributions for the heterogeneity is by using (or introducing) a scale parameter (Mood et al., 1974, Section VII.6.2). An example would be given by varying the scale parameter within the half-normal family of prior distributions (Röver et al., 2021). This property holds for given data (y_i) and a stochastically ordered family of heterogeneity priors ($p(\tau)$).

5 | NUMERICAL ILLUSTRATION

In order to demonstrate the shrinkage weights’ properties, we consider an illustrative case motivated by a scenario involving a log-odds-ratio (log-OR) endpoint, analogous to the simulation scenario discussed by Röver and Friede (2020). For a study of size n_i featuring two treatment arms and a binary endpoint, the results may be summarized in a 2×2 contingency table. Assuming an even distribution of events and nonevents across table cells implies a log-OR estimate with a standard error of approximately $\frac{4}{\sqrt{n_i}}$ (Röver, 2020). Considering a combination of a “small” and a “large” study with sizes $n_1 = 25$ and $n_2 = 400$ then leads to standard errors of $\sigma_1 = 0.8$ and $\sigma_2 = 0.2$, respectively. We will then derive the smallest RCT’s shrinkage estimate (for the study-specific effect θ_1) that is of course primarily informed by y_1 , but supported by the external data y_2 . The present case of $\sigma_1 \gg \sigma_2$ is the kind of setting in which we expect to see larger gains from shrinkage estimation, and this is exactly where using historical data is the most attractive in practice.

For the analysis, we choose a half-normal heterogeneity prior with scale 0.5 (HN(0.5)), which constitutes a conservative choice for the present scenario (Friede et al., 2017; Röver et al., 2021). For illustration purposes, we also utilize a (stochastically larger) HN(1.0) prior. We then fix the target y_1 (arbitrarily) at zero and vary the source y_2 in order to investigate the effect on the resulting shrinkage estimates and weights.

Figure 1 illustrates estimates’ and weights’ dependence on the difference between estimates (y_1 and y_2). The top row of forest plots shows three example cases of (a) coinciding target and source estimates, (b) some moderate, and (c) larger discrepancy between the two; the resulting shrinkage estimate for the target is shown in blue. The second row shows the posterior means of θ_1 (solid lines) and the corresponding 95% CIs (dashed lines) across the continuum of source data values. At the top of the plot, the three cases (a)–(c) are marked, and the blue lines correspond to the estimates also shown above. The red lines show analogous estimates, but corresponding to a (stochastically larger) HN(1.0) prior. Note that “large” $|y_2 - y_1|$ values (here, e.g., $|y_2 - y_1| > 1.96(\sigma_1 + \sigma_2) = 1.96$) would imply nonoverlapping CIs for source and target studies (as in case (c)), which in reality may mean that estimates would not actually be pooled at all. The practically most relevant bit of the plot is hence in the neighborhood of zero.

Finally, the bottom plot shows the posterior expected weights to illustrate the first (target) study’s contribution to its own shrinkage estimate. The minimum (for both heterogeneity priors) is attained in the “coincidence case” (a) of $y_2 - y_1 = 0$; for example, for the HN(0.5) prior the coincidence weight is at 29%. Increasing the observed effect difference $|y_2 - y_1|$ (i.e.,




FIGURE 1 Effect of varying the difference between quoted estimates ($y_2 - y_1$) on the first shrinkage estimate (for θ_1). The top row shows three data examples of (a) coinciding and (b)–(c) increasingly diverging estimates, along with the resulting shrinkage estimate for the target study. The second row illustrates the estimates across the continuum of increasing y_2 values relative to the “plain” interval ($y_1 \pm 1.96\sigma_1$). The bottom panel shows the posterior mean shrinkage weight ($E[c_{11}(\tau)|\vec{y}, \vec{\sigma}]$) for the first study, based on two different priors and for varying $y_2 - y_1$. Note that $y_2 - y_1 = 0$ constitutes the “coincidence case”

the “observed heterogeneity”) then yields increasing weights for y_1 , implying less borrowing from the source. In cases (b) and (c), the shrinkage weight amounts to 38% and 63%, respectively. Also, the choice of a stochastically larger prior, here realized by a larger scale parameter in the same family of distributions, leads to larger weights for y_1 , for any $|y_2 - y_1|$, including the minimum at $|y_2 - y_1| = 0$. The first study’s absolute minimum shrinkage weight, the “FE weight,” in this case is at $c_{11}(0) = \frac{\sigma_1^2}{\sigma_1^2 + \sigma_2^2} = \frac{1}{17} = 5.9\%$.

TABLE 2 Data from Varges et al. (2017) on an observational and a randomized study investigating the effect of doxycycline on survival in CJD

i	Study	Patients		log(HR)	
		Treatment	Control	y_i	σ_i
1	Observational	55	33	-0.499	0.249
2	Randomized	7	5	-0.173	0.631

TABLE 3 Estimates for the CJD example

τ prior	Mean weight		Effect estimate θ_2	
	Minimum	Actual	Mean	95% CI
HN(0.5)	38.9%	39.5%	-0.370	[-1.157, 0.477]
HN(1.0)	52.1%	53.1%	-0.326	[-1.232, 0.664]
		(100.0%)	-0.173	[-1.410, 1.064]

For different heterogeneity priors (HN(0.5) or HN(1.0)), the corresponding minimum (coincidence) weight is given, as well as the resulting weight for the actual data along with the corresponding shrinkage estimates. The very last line shows the estimate based only on y_2 and σ_2 for comparison.

Note that while “ $y_1 = y_2$ ” constitutes a “worst case” in a certain sense (leading to the lowest shrinkage weight), it also still is the most desirable case, in the sense that this is when the data are in agreement and one would expect to learn the most from the source study.

6 | APPLICATIONS

6.1 | Creutzfeldt–Jakob example

A small RCT was conducted in order to investigate the effect of doxycycline on survival in patients suffering from Creutzfeldt–Jakob disease (CJD). In this ultra rare condition, only 12 patients could be recruited, and so data from an observational study were considered as complementing evidence (Varges et al., 2017). Both studies quote estimated hazard ratios (HRs), and these estimates along with their standard errors are jointly analyzed in a meta-analysis; the data are also shown in Table 2. With the focus being on the evidence from the RCT, a shrinkage estimate for this study is derived (Röver & Friede, 2020). Both studies are in agreement, suggesting a beneficial treatment effect, while the absolute effect magnitude is larger for the observational data.

Since the larger observational study provides a much more precise estimate (smaller standard error), one might fear that the randomized evidence will be overwhelmed by the external data in a joint analysis. The FE weight in this case amounts to $c_{22}(0) = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} = 13.5\%$; this would be the RCT’s weight in an FE analysis, and it constitutes a lower bound on the RCT’s weight for any data realization (y_1, y_2) or any heterogeneity prior $(p(\tau))$.

For a log-HR, we may then assume a half-normal prior with scale 0.5 (HN(0.5)) for the heterogeneity (Friede et al., 2017; Röver & Friede, 2020; Röver et al., 2021). For this prior, we get a minimum posterior mean weight (coincidence weight) for the randomized study of 38.9%, which may already be considered reassuringly large, in view of the sample sizes involved and compared to the FE weight. Any data realization (y_1, y_2) will hence yield an eventual weight $\geq 38.9\%$ for the RCT. Also, a larger scale of the heterogeneity prior (i.e., a larger expected amount of heterogeneity) will increase the minimum weight for y_2 ; for example, a HN(1.0) prior would yield a minimum expected shrinkage weight of 52.1%. For the actual data (Table 2), we then get a weight of 39.5%, slightly above the minimum, for the RCT. Table 3 shows weights and estimates corresponding to the two different heterogeneity priors. In both cases, the actual weights are not far from their minimum value, and for both analyses there is a sizeable gain in precision for the shrinkage estimate when compared to the original estimate (y_2, σ_2) alone.

TABLE 4 Data from Rietbergen et al. (2011) on two gynecological RCTs investigating whether fetal monitoring using CTG combined with ECG ST-segment analysis was associated with a reduced risk of metabolic acidosis, compared to CTG alone

<i>i</i>	Study	Treatment		Control		log(RR)	
		Events	Total	Events	Total	y_i	σ_i
1	Amer-Wählin (2001)	15	2159	31	2079	−0.764	0.313
2	Westerhuis (2007)	20	2827	30	2840	−0.401	0.287

TABLE 5 Estimates for the metabolic acidosis example

τ prior	Mean weight		Effect estimate θ_2	
	Minimum	Actual	Mean	95% CI
HN(0.5)	72.5%	74.0%	−0.495	[−0.986, 0.005]
HN(1.0)	78.7%	80.5%	−0.472	[−0.983, 0.051]
		(100.0%)	−0.401	[−0.964, 0.163]

For different heterogeneity priors (HN(0.5) or HN(1.0)), the corresponding minimum (coincidence) weight is given, as well as the resulting weight for the actual data along with the corresponding shrinkage estimates. The very last line shows the estimate based only on y_2 and σ_2 for comparison.

6.2 | Metabolic acidosis example

A gynecological RCT investigated whether fetal monitoring using cardiotocography (CTG) combined with ECG ST-segment analysis (ST) reduced the occurrence of metabolic acidosis, compared to CTG alone (Westerhuis et al., 2007). Here the relative risk (RR) of metabolic acidosis comparing the two treatment groups is of interest. When analyzing the data, evidence from an earlier, similar RCT (Amer-Wählin et al., 2001) may be utilized to support parameter estimation. This example dataset was originally investigated by Rietbergen et al. (2011); the corresponding data are shown in Table 4.

Primary interest focuses on the more recent target study by Westerhuis et al. (2007) and on a shrinkage estimate of its study-specific effect θ_2 . The two trials are of roughly comparable size (5667 vs. 4238 participants), and from the “FE weight” of $c_{22}(0) = \frac{\sigma_2^2}{\sigma_1^2 + \sigma_2^2} = 54.3\%$ one can already see that the second study will definitely contribute the majority of weight when estimating its own effect θ_2 .

For a log-RR, we may again use a half-normal prior with scale 0.5 for the heterogeneity (Friede et al., 2017; Röver et al., 2021); this yields a minimum (coincidence) mean shrinkage weight of 72.5%. A larger heterogeneity prior scale again leads to an increased shrinkage weight; for example, for a HN(1.0) prior, the minimum weight is at 78.7%.

Table 5 shows the corresponding weights and estimates. Compared to the previous example, the precision gain is not quite as large here.

7 | CONCLUSIONS

Bayesian meta-analysis provides a transparent means for extrapolation or *borrowing of strength* from external data (Röver & Friede, 2020). Also within a Bayesian inference framework, study weights for overall and shrinkage effect estimates may be derived as posterior expected weights, for any number of studies k . The FE weights (conditional on $\tau = 0$) constitute the absolute minimum shrinkage weights across all heterogeneity priors and data realizations. In case of $k = 2$ studies, the heterogeneity posterior depends on the data only via the absolute difference in both estimates ($|y_2 - y_1|$). A larger difference leads to a stochastically larger heterogeneity posterior. When the estimates coincide, that is, $y_2 = y_1$, the smallest possible shrinkage weight for a given heterogeneity prior (across all possible data realizations) is obtained. Concerning the choice of heterogeneity prior, a stochastically larger prior leads to a stochastically larger posterior and with that to increased (minimum and actual) shrinkage weights.

The above findings have important implications for the weightings that may occur within a meta-analysis. The shrinkage weight is bounded below (irrespective of the prior and data) by the FE weight. For any particular given prior, the (posterior mean) shrinkage weight is also bounded below across possible data realizations by the “coincidence weight.” Having a bound on the weight effectively means bounding the “leverage” of the external data for the shrinkage estimate. A lower bound of, say, 50% means that the resulting shrinkage estimate will not move more than halfway from the effect y_i

towards the external data in case of near-concordant evidence. For greater discrepancies, the target study's weight will even be larger, or, conversely, the influence of the source data will be smaller.

The Food and Drug Administration guidance on “Leveraging existing clinical data for extrapolation to pediatric uses of medical devices” (U.S. Department of Health and Human Services & Food and Drug Administration, 2016), for example, elaborates on issues commonly encountered in extrapolation endeavors. One concern raised here is the exchangeability assumption (3) commonly made in hierarchical models. In the common case of only $k = 2$ studies, however, the same model (as far as shrinkage estimation is concerned) may alternatively be motivated via the *reference model* (Röver & Friede, 2020). This is similar to the *bias allowance model* framework (Welton et al., 2012), where the target study is estimating the parameter of interest “directly,” while the source is associated with a potential bias term of unknown direction and magnitude. Moreover, the advantages of using (informative) priors on the heterogeneity parameter are acknowledged in the guidance document, in particular as this facilitates dynamical borrowing based on the empirically observed compatibility of source and target data.

We would like to encourage consideration of minimum weights as a *diagnostic tool* of the evidence constitution and of implications of prior settings for a given or anticipated data scenario. The study of weights should, however, not be used for guiding the selection of the heterogeneity prior. The choice of prior should primarily be driven by considerations of prior information on between-study variability. Different amounts of heterogeneity might however be anticipated in different contexts—for example, in the two examples discussed above, greater heterogeneity may be plausible between observational and randomized data than for the case of two RCTs (Röver et al., 2021).

A closely related approach to investigating the contributions of data sources to a joint analysis works via the consideration of *effective sample sizes* (Neuenschwander et al., 2020). Target and source data may be assessed based on their “share” of the total sample size and also the effect of heterogeneity on the resulting MAP prior may be quantified via *prior maximum sample sizes* (Neuenschwander et al., 2010). Robustification of a MAP prior may be achieved by implementing a mixture of vague and informative prior components (Schmidli et al., 2014). In case the source data's weight is considered *too large*, a simple remedy might be to artificially inflate its associated standard error, similarly to the idea behind a *power prior* approach (Ibrahim & Chen, 2000). Alternatively, the target sample size might also be increased, in case that is an option. For an overview of approaches to downweighting of external data, see also Viele et al. (2014).

The considerations which provide some insights into the inner workings of shrinkage estimation facilitate diagnostics even before considering actual data. A requirement is that the standard errors (σ_i) need to be known beforehand. Quite often these may be approximated based on a *unit information standard deviation* and the sample size (Kass & Wasserman, 1995; Röver et al., 2021). The need to make assumptions about anticipated standard errors is a common issue in similar design-of-experiment contexts (Cohen, 1988). Fears of external evidence “overruling” the target data (Weber et al., 2018) may be unwarranted, or may be checked before carrying out the target study, as the NNHM behaves predictably and reasonably within a Bayesian analysis. Potential problems arise or are amplified when using frequentist methods: the concerning common occurrence of zero heterogeneity estimates means that analyses may fall back to an FE approach, which here is the least cautious or least conservative analysis. For the case of few studies, the probability of obtaining a zero heterogeneity estimate is alarmingly high—approaching 50% even for moderate amounts of heterogeneity (Friede et al., 2017), which may actually render frequentist heterogeneity estimation for small k a somewhat questionable exercise. Within a Bayesian framework, marginalization over the plausible range of heterogeneity values will lead to a more conservative behavior. In summary, with the target study's contribution to the resulting Bayesian shrinkage estimate being bounded below, concerns of evidence being easily overwhelmed by external source data can be addressed a priori and may be shown to be largely unwarranted.

ACKNOWLEDGMENT

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


CONFLICTS OF INTEREST

The authors have declared no conflict of interest.

DATA AVAILABILITY STATEMENT

The data used in the examples are shown in Tables 2 and 4.

OPEN RESEARCH BADGES

 This article has earned an Open Data badge for making publicly available the digitally-shareable data necessary to reproduce the reported results. The data is available in the Supporting Information section.

This article has earned an open data badge **“Reproducible Research”** for making publicly available the code necessary to reproduce the reported results. The results reported in this article could fully be reproduced.

ORCID

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

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

REFERENCES

- Amer-Wählin, I., Hellsten, C., Norén, H., Herbst, A., Kjellmer, I., Lilja, H., Lindoff, C., Månsson, M., Olofsson, P., Sundström, A.-K., & Maršál, K. (2001). Cardiocography only versus cardiocography plus ST analysis of fetal electrocardiogram for intrapartum fetal monitoring: A Swedish randomised controlled trial. *The Lancet*, *358*, 534–538.
- Bartoszewicz, J., & Skolimowska, M. (2006). Preservation of classes of life distributions and stochastic orders under weighting. *Statistics & Probability Letters*, *76*, 587–596.
- Best, N., Dallow, N., & Montague, T. (2020). Prior elicitation. In E. Lesaffre, G. Baio, & B. Boulangier (Eds.), *Bayesian methods in pharmaceutical research* (Ch. 5, pp. 87–109). Chapman & Hall/CRC.
- Cohen, J. (1988). *Statistical power analysis for the behavioural sciences* (2nd ed.). Routledge.
- Friede, T., Röver, C., Wandel, S., & Neuenschwander, B. (2017a). Meta-analysis of few small studies in orphan diseases. *Research Synthesis Methods*, *8*, 79–91.
- Friede, T., Röver, C., Wandel, S., & Neuenschwander, B. (2017b). Meta-analysis of two studies in the presence of heterogeneity with applications in rare diseases. *Biometrical Journal*, *59*, 658–671.
- Gross, O., Tönshoff, B., Weber, L. T., Pape, L., Latta, K., Fehrenbach, H., Lange-Sperandio, B., Zappel, H., Hoyer, P., Staude, H., König, S., John, J., und Gellermann, U., Hoppe, B., Galiano, M., Hoecker, B., Ehren, R., Lerch, C., Kashtan, C. E., .. Friede, T. (2020). A multicenter, randomized, placebo-controlled, double-blind phase 3 trial with open-arm comparison indicates safety and efficacy of nephroprotective therapy with ramipril in children with Alport’s syndrome. *Kidney International*, *97*, 1275–1286.
- Günhan, B. K., Röver, C., & Friede, T. (2020). Random-effects meta-analysis of few studies involving rare events. *Research Synthesis Methods*, *11*, 74–90.
- Hampson, L. V., Whitehead, J., Eleftheriou, D., & Brogan, P. (2014). Bayesian methods for the design and interpretation of clinical trials in very rare diseases. *Statistics in Medicine*, *33*, 4186–4201.
- Hartung, J., Knapp, G., & Sinha, B. K. (2008). *Statistical meta-analysis with applications*. Wiley.
- Hedges, L. V., & Olkin, I. (1985). *Statistical methods for meta-analysis*. Academic Press.
- Ibrahim, J. G., & Chen, M.-H. (2000). Power prior distributions for regression models. *Statistical Science*, *15*, 46–60.
- Jackson, D., & White, I. R. (2018). When should meta-analysis avoid making hidden normality assumptions? *Biometrical Journal*, *60*, 1040–1058.
- Kass, R. E., & Wasserman, L. (1995). A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *Journal of the American Statistical Association*, *90*, 928–934.
- Lehmann, E. L. (1955). Ordered families of distributions. *The Annals of Mathematical Statistics*, *26*, 399–419.
- Mezcarski, M. (2015). 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, Poland. <https://EconPapers.repec.org/RePEc:sgl:annals:i:37:y:2015:p:339-360>.
- Mood, A. M., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics* (3rd ed.). McGraw-Hill.
- Neuenschwander, B., Capkun-Niggli, G., Branson, M., & Spiegelhalter, D. J. (2010). Summarizing historical information on controls in clinical trials. *Clinical Trials*, *7*, 5–18.
- Neuenschwander, B., Weber, S., Schmidli, H., & O’Hagan, A. (2020). Predictively consistent prior effective sample sizes. *Biometrics*, *76*, 578–587.
- Raudenbush, S. W., & Bryk, A. S. (1985). Empirical Bayes meta-analysis. *Journal of Educational and Behavioural Statistics*, *10*, 75–98.
- Rietbergen, C., Klugkist, I., Janssen, K. J. M., Moons, K. G. M., & Hoijtink, H. J. A. (2011). Incorporation of historical data in the analysis of randomized therapeutic trials. *Contemporary Clinical Trials*, *32*, 848–855.
- Robinson, G. K. (1991). That BLUP is a good thing: The estimation of random effects. *Statistical Science*, *6*, 15–32.
- Röver, C. (2015). bayesmeta: Bayesian random-effects meta-analysis. R package. <http://cran.r-project.org/package=bayesmeta>.
- Röver, C. (2020). Bayesian random-effects meta-analysis using the bayesmeta R package. *Journal of Statistical Software*, *93*, 1–51.
- Röver C., Bender R., Dias S., Schmid C. H., Schmidli H., Sturtz S., Weber S., Friede T. (2021). On weakly informative prior distributions for the heterogeneity parameter in Bayesian random-effects meta-analysis. *Research Synthesis Methods*, <https://doi.org/10.1002/jrsm.1475>.
- Röver, C., & Friede, T. (2020). Dynamically borrowing strength from another study. *Statistical Methods in Medical Research*, *29*, 293–308.
- Schmidli, H., Gsteiger, S., Roychoudhury, S., O’Hagan, A., Spiegelhalter, D., & Neuenschwander, B. (2014). Robust meta-analytic-predictive priors in clinical trials with historical control information. *Biometrics*, *70*, 1023–1032.
- Shaked, M., & Shanthikumar, J. G. (2007). *Stochastic orders*. Springer-Verlag.

- U.S. Department of Health and Human Services, & Food and Drug Administration (2016). Leveraging existing clinical data for extrapolation to pediatric uses of medical devices. Guidance for Industry and Food and Drug Administration Staff. <https://www.fda.gov/media/91889/download>.
- Varges, D., Manthey, H., Heinemann, U., Ponto, C., Schmitz, M., Schulz-Schaeffer, W. J., Krasnianski, A., Breithaupt, M., Fincke, F., Kramer, K., Friede, T., & Zerr, I. (2017). Doxycycline in early CJD – a double-blinded randomized phase II and observational study. *Journal of Neurology, Neurosurgery and Psychiatry*, 88, 119–125.
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36, 1–48.
- Viele, K., Berry, S., Neuenschwander, B., Amzal, B., Chen, F., Enas, N., Hobbs, B., Ibrahim, J. G., Kinnersley, N., Lindborg, S., Micallef, S., Roychoudhury, S., & Thompson, L. (2014). Use of historical control data for assessing treatment effects in clinical trials. *Pharmaceutical Statistics*, 13, 41–54.
- Wandel, S., Neuenschwander, B., Röver, C., & Friede, T. (2017). Using phase II data for the analysis of phase III studies: an application in rare diseases. *Clinical Trials*, 14, 277–285.
- Weber, K., Hemmings, R., & Koch, A. (2018). How to use prior knowledge and still give new data a chance? *Pharmaceutical Statistics*, 17, 329–341.
- Welton, N. J., Sutton, A. J., Cooper, N. J., Abrams, K. R., & Ades, A. E. (2012). *Evidence synthesis for decision making in healthcare*. Wiley.
- Westerhuis, M. E. M. H., Moons, K. G. M., van Beek, E., Bijvoet, S. M., Drogtop, A. P., van Geijn, H. P., van Lith, J. M. M., Mol, B. W. J., Nijhuis, J. G., Oei, S. G., Porath, M. M., Rijnders, R. J. P., Schuitemaker, N. W. E., van der Tweel, I., Visser, G. H. A., Willekes, C., & Kwee, A. (2007). A randomised clinical trial on cardiocography plus fetal blood sampling versus cardiocography plus ST-analysis of the fetal electrocardiogram (STAN) for intrapartum monitoring. *BMC Pregnancy and Childbirth*, 7, 13.

SUPPORTING INFORMATION

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

How to cite this article: Röver C, Friede T. Bounds for the weight of external data in shrinkage estimation. *Biometrical Journal*. 2021;63:1131–1143. <https://doi.org/10.1002/bimj.202000227>

APPENDIX

A.1 | Stochastic ordering of heterogeneity posteriors

Consider two parameter sets \vec{y}_a and \vec{y}_b for which $0 \leq |y_{a;2} - y_{a;1}| < |y_{b;2} - y_{b;1}|$. Then the ratio of the heterogeneity's marginal posterior densities is given by (cf. (14))

$$\frac{p(\tau|\vec{y}_b, \vec{\sigma})}{p(\tau|\vec{y}_a, \vec{\sigma})} = \frac{c_{\vec{y}_b} p(\tau) g_{\vec{y}_b}(\tau)}{c_{\vec{y}_a} p(\tau) g_{\vec{y}_a}(\tau)} = \frac{c_{\vec{y}_b}}{c_{\vec{y}_a}} \frac{g_{\vec{y}_b}(\tau)}{g_{\vec{y}_a}(\tau)} \propto \frac{g_{\vec{y}_b}(\tau)}{g_{\vec{y}_a}(\tau)}, \quad (\text{A1})$$

where $c_{\vec{y}_a}$ and $c_{\vec{y}_b}$ are the densities' normalizing constants, and the where the latter ratio of “ $g_{\vec{y}}(\tau)$ ” terms is monotonically increasing in τ . With that, condition (C) in Lehmann (1955) is fulfilled, and the posterior corresponding to \vec{y}_b is *stochastically larger* than the one associated with \vec{y}_a .

A.2 | Stochastic ordering of posteriors for different priors

Consider two heterogeneity priors with densities $p_1(\tau)$ and $p_2(\tau)$, where p_2 is stochastically larger than p_1 . A posterior distribution constitutes a special case of a “weighted distribution” (Męczarski, 2015). For the posterior distributions corresponding to p_1 and p_2 follows that these will inherit the same stochastic ordering (Bartoszewicz & Skolimowska, 2006).

A.3 | R code for CJD example

```
# specify data:
cjd <- cbind.data.frame("study" =c("observational", "randomized"),
"logHR" =c(-0.499, -0.173),
"logHR.se" =c(0.249, 0.631))
# analyze:
library("bayesmeta")
bm <- bayesmeta(y=cjd$logHR, sigma=cjd$logHR.se, labels=cjd$study,
tau.prior=function(t) dhalfnormal(t, scale=0.5))
# show posterior mean shrinkage weights:
```

```
bm$weights.theta
# show shrinkage estimates:
bm$theta
# derive FE weights (percentages, using "metafor" library):
weights(rma.uni(yi=cjd$logHR, sei=cjd$logHR.se, slab=cjd$study,
measure="GEN", method="FE"))
# alternatively, compute directly:
cjd$logHR.se^-2 / sum(cjd$logHR.se^-2)
# determine coincidence (minimum) posterior mean weights:
bayesmeta(y=c(0,0), sigma=cjd$logHR.se, labels=cjd$study,
tau.prior=function(t) dhalfnormal(t, scale=0.5))$weights.theta
```