Research Synthesis Methods (2026), **17**: 797–815
doi:10.1017/rsm.2026.10081

Research Synthesis Methods

RESEARCH ARTICLE

Meta-analytic-predictive priors based on a single study

Christian Röver [1] and Tim Friede [1,2,3]

[1]Department of Medical Statistics, University Medical Center Göttingen, Göttingen, Germany
[2]DZHK (German Center for Cardiovascular Research), Partner Site Lower Saxony, Göttingen, Germany
[3]DZKJ (German Center for Child and Adolescent Health), Göttingen, Germany

Corresponding author: Christian Röver; Email: christian.roever@med.uni-goettingen.de

Received: 21 May 2025; **Revised:** 10 November 2025; **Accepted:** 4 February 2026; published online 24 March 2026

Keywords: bias allowance; dynamic borrowing; MAP prior; power prior; random-effects meta-analysis; shrinkage estimation

Abstract

Meta-analytic-predictive (MAP) priors have been proposed as a generic approach to deriving informative prior distributions, where external empirical data are processed to learn about certain parameter distributions. The use of MAP priors is also closely related to shrinkage estimation (also sometimes referred to as *dynamic borrowing*). A potentially odd situation arises when the external data consist only of *a single study*. Conceptually, this is not a problem, it only implies that certain prior assumptions gain in importance and need to be specified with particular care. We outline this important, not uncommon special case and demonstrate its implementation and interpretation based on the normal–normal hierarchical model. The approach is illustrated using example applications in clinical medicine.

Highlights

What is already known?

- Shrinkage estimation may be used to effectively and robustly borrow information between related data sources.
- Shrinkage estimation may alternatively be motivated via a meta-analytic-predictive (MAP) approach.

What is new?

- A MAP approach remains sensible down to the extreme case of only a single study.
- The MAP prior's usual features are retained, in addition, there are connections to power prior and bias allowance approaches.

Potential impact for RSM readers

- MAP priors are useful for constructing empirically motivated priors based on external/historical data.
- MAP priors may serve as an additional motivation for related approaches (bias allowance models and power priors).
- Practical application is straightforward using existing software packages.

This article was awarded Open Data and Open Materials badges for transparent practices. See the Data availability statement for details.

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Society for Research Synthesis Methodology. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (https://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

1. Introduction

The potential of clinical research is commonly limited by data sparsity issues; such problems particularly arise in the context of rare diseases, where the number of potential study subjects is small, or in pediatric indications, where ethical considerations may limit the recruitment of patients. The large variety of rare diseases still means that a sizeable proportion of the population is affected by rare diseases, posing a considerable economic burden. Even in more common indications, data sparsity problems may arise, for example, when the focus is on smaller sub-populations, or when novel treatments or standards of care emerge. In any of these cases, the careful consideration of all potentially relevant evidence available is essential.[1–3] When evidence from a single experiment, such as a clinical trial, is not sufficiently conclusive on its own, it may sometimes help to view the data in the context of related instances (similar experiments) in order to yield more confident conclusions. This idea is explicitly implemented in *shrinkage estimation*, where a hierarchical model is set up accounting for estimation uncertainly at the study level as well as for variability (and similarity) between studies[4,5]; models of this kind are commonly also used in the context of meta-analysis.[6,7] The borrowing-of-information taking place between the study of primary interest and the external data may be viewed in terms of the overarching joint model as a *meta-analytic-combined (MAC)* approach, or, equivalently, by formulating the *meta-analytic predictive (MAP)* prior that explicates the information contributed by the external data to the shrinkage estimate.[8] A particular special case is given when a single ("target") study is supported by a single ("source") study; such situations are not uncommon, and shrinkage estimation here has proven useful.[9–11] Application of a hierarchical model makes it behave *dynamically* in the sense that more or less information is borrowed, depending on the apparent similarity of target and source data.[9] When considering this case in terms of the implied MAP prior, the "meta-analysis" involved here is based on a single study, which may appear somewhat counterintuitive at first. It is this perceived contradiction that we aim to address here; while a *meta-analysis* is commonly thought of as involving larger amounts of data, we will see that a hierarchical model may essentially be fit also to a single data point, and sensible predictions may be derived. Such a smallest-possible meta-analysis does not pose a conceptual problem, and there is no reason to abandon the general concept even if the amount of historical data drops below a couple of studies. It only implies that, due to the particular sparsity of data, prior specification within the model receives special importance, a problem which, however, is common in meta-analysis of few studies in general,[12] and which analogously applies for alternative (and closely related) borrowing methods, such as power priors or bias allowance models.[11,13] Since others appear to have struggled with or shied away from the idea of a single-study meta-analysis where in fact it may have been a viable option,[14,15] it seems worthwhile to investigate this special case a bit closer. Closer inspection of this particular case then also highlights how the properties of this MAP prior materialize, as well as its close connection to bias allowance and power prior approaches.

The remainder of this article is structured as follows: in Section 2, the normal–normal hierarchical model (NNHM) is introduced, the meta-analysis model which then is the basis for shrinkage estimation between a pair of studies, and for the MAP prior based on a single study. The ideas will be illustrated in two practical examples in Section 3. Section 3.1 discusses an application in paediatric Alport syndrome that was originally formulated in terms of a shrinkage estimation problem. Section 3.2 introduces a trial design application in cardiology, where information from a similar past study is designated for consideration in the eventual analysis via an informative MAP prior. Section 4 then closes with a brief discussion.

2. Shrinkage estimation using two studies

2.1. The normal–normal hierarchical model

The most common model for random-effects meta-analysis is given by the NNHM. It implements sampling error as well as between-study heterogeneity using normal distributions. The data are given

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

in terms of k estimates y_i and associated standard errors s_i $(i = 1, \ldots, k)$. Each individual study aims to quantify a parameter θ_i, so that

$$y_i|\theta_i, s_i \sim \text{Normal}(\theta_i, s_i^2). \tag{2.1}$$

The underlying parameters θ_i are not necessarily identical for all studies, instead some amount of *(between-study) heterogeneity* is allowed for, expressed as

$$\theta_i|\mu, \tau \sim \text{Normal}(\mu, \tau^2). \tag{2.2}$$

Often the overall mean μ is the aim of the analysis, while sometimes the *study-specific parameters* θ_i are also of interest.[9,16] The heterogeneity, while important, usually remains a nuisance parameter. In the context of "shrinkage estimation" of the θ_i, an interesting aspect is that the problem may be motivated in two ways; classically, one may think of shrinkage estimation as a joint analysis of all (k) estimates, which also returns estimates of any θ_i parameter along the way; this is also denoted as the *meta-analytic-combined* (*MAC*) approach. The problem, however, may also be factored into the evidence stemming from the ith study alone, as well as the information provided by the remaining $(k - 1)$ estimates. Shrinkage estimation then may be interpreted as the analysis of the ith study, based on a prior distribution that results as the predictive distribution derived from a meta-analysis of the other $(k - 1)$ studies; this prior is denoted as the *meta-analytic-predictive* (*MAP*) *prior*. Both MAC and MAP approaches are equivalent and yield identical shrinkage estimates.[8]

In the following, we will focus on the special case of only two studies $(k = 2)$. For the shrinkage estimate (θ_2), this implies a MAP prior that is based on a single study (i.e., the data provided through y_1 and s_1). While this may appear odd at first, the idea readily applies also in this special case, as will be demonstrated in the following. Analysis may generally be performed based on informative or uninformative priors for μ, while a proper, informative prior is required for τ.[7,12] The case of only $k = 2$ studies is also closely connected to the related concepts of *power prior*[9] or *bias allowance models*.[13,17]

2.2. Uniform prior for the overall mean effect (μ)

Priors for the overall mean parameter (μ) in the NNHM may be specified as informative or as uninformative. For (more or less informative) priors, normal distributions are an obvious choice, also since these lead to analytically simple inference. Quite commonly, effect priors however are chosen as uninformative and (improper) uniform, not least due to certain analogies to frequentist meta-analysis procedures.[7] In case of an (improper, non-informative) uniform prior for μ, certain expressions turn out particularly simple, which is also why we will focus on this particular, yet insightful, common and practically relevant case in the following. As the (improper) uniform prior constitutes the limiting case of an increasingly uninformative effect prior, the following considerations may also be viewed as relating to the limiting behavior for increasingly uninformative priors (e.g., for normal priors when their variance approaches infinity). The uniform effect prior leads to a normal conditional posterior for the overall mean effect μ, with moments given by

$$\text{E}[\mu|y_1, s_1, \tau] = y_1 \quad \text{and} \quad \text{Var}(\mu|y_1, s_1, \tau) = s_1^2 + \tau^2, \tag{2.3}$$

and a marginal heterogeneity likelihood that is constant (independent of τ), so that the heterogeneity's posterior equals its prior.[7] This seems reasonable, since (as long as the overall mean effect prior is uniform) a single observation y_1 does not provide information on the heterogeneity τ.

2.3. The MAP prior for the effect in a new study

The MAP prior results as the posterior predictive distribution for a "new," second study's (study-specific) effect θ_2 given the data from the first study (y_1, s_1). In the NNHM framework, the *conditional predictive distribution* again is normal with mean

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

$$\mathrm{E}[\theta_2|y_1, s_1, \tau] \;=\; \mathrm{E}[\mu|y_1, s_1, \tau] \;=\; y_1 \tag{2.4}$$

and variance

$$\mathrm{Var}(\theta_2|y_1, s_1, \tau) \;=\; \mathrm{Var}(\mu|y_1, s_1, \tau) + \tau^2 \;=\; s_1^2 + 2\tau^2 \tag{2.5}$$

(see also (2.2) and (2.3), and the more detailed derivation in Appendix A.1). These expressions make sense in the present context: We know θ_1 with accuracy given by the standard error s_1, and we know that the difference between θ_1 and θ_2 is normally distributed with variance $2\tau^2$, so that the (conditional) variance expression results as a corresponding sum.[9] As pointed out above, information on heterogeneity (τ) so far is based on the prior *only*; the heterogeneity in this context generally requires a proper, informative prior (since $k < 3$).[7,12]

The eventual (marginal) predictive distribution (marginalized over the distribution of τ) hence results as a normal *scale mixture*[18,19] with fixed mean (2.4) and with variance as given in (2.5), where τ is distributed according to the specified prior. In particular, one may think of the MAP prior as the first study's point estimate $p(\theta_1|y_1, s_1)$ convolved with the prior predictive distribution $p(\theta_2|\theta_1, \tau)$ and then marginalized over τ. The MAP prior is symmetric around y_1, and its (marginal) variance results from (2.5) as

$$\mathrm{Var}(\theta_2|y_1, s_1) \;=\; \mathrm{E}[s_1^2 + 2\tau^2] \;=\; s_1^2 + 2\,\mathrm{E}[\tau^2], \tag{2.6}$$

where the expectation $\mathrm{E}[\tau^2]$ depends on the assigned heterogeneity prior. For a range of common prior specifications, this expectation may be derived analytically; Table A1 in Appendix A.4 lists some popular cases. From this expression, one can see that the relative magnitudes of s_1^2 and $\mathrm{E}[\tau^2]$ determine whether the resulting MAP prior's variance is dominated by estimation uncertainty (regarding θ_1) or anticipated heterogeneity (τ). These two variance components may in fact also be considered to reflect so-called "type A" and "type B" uncertainties relating to measurement uncertainty and background knowledge, respectively, which together sum up to form the *combined standard uncertainty* u_c (the square root of (2.6)).[20,21]

Since the MAP prior results as a normal scale mixture, it is generally heavier-tailed than a normal distribution, which has implications for the resulting operating characteristics. A heavy-tailed (MAP-) prior means that in combination with a ("shorter-tailed") normal likelihood, the likelihood will dominate in case of a prior-data conflict.[22] Such robustness properties have in fact been noted and demonstrated in the meta-analysis context, as these lead to a *dynamic* borrowing behavior.[9] The MAP priors' tail behavior will also be illustrated for some examples below (Figures 4 and 5).

Another way to quantify the precision of a prior distribution is by relating it to a number of observations that would in a certain sense convey an equivalent amount of information. In the following, we will use two approaches to this effect. Firstly, a prior may be assessed in terms of a corresponding absolute *effective sample size (ESS)* (here: number of patients), this will be done by quoting ESSs based on the *expected local-information-ratio* ($\mathrm{ESS}_{\mathrm{ELIR}}$). This measure is based on the prior density's curvature and it ensures *predictive consistency*, that is, on expectation, the posterior's ESS will be the sum of the prior's $\mathrm{ESS}_{\mathrm{ELIR}}$ *plus* the actual sample size.[23] Secondly, the added information from including the prior in a specific analysis may be expressed in terms of the (relative) *gain in ESS*.[9] This is based on comparing the relative width of a confidence interval *with* and *without* considering the informative prior, and then determining by what factor the sample size would have needed to be increased to yield the same precision gain (see also Appendix A.2).

2.4. The bias allowance model connection

In the 2-study case, there is a one-to-one correspondence between the NNHM and a simple *bias allowance model*[13]; instead of the NNHM assumption (2.2) in combination with a uniform prior for the overall mean effect μ and heterogeneity prior $p_\star(\tau)$ as in Section 2.1, one may specify

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

$$\theta_2|\alpha, \beta = \alpha \tag{2.7}$$

$$\theta_1|\alpha, \beta \sim \text{Normal}(\alpha,\ \beta^2) \tag{2.8}$$

with prior $p(\beta) = \frac{1}{\sqrt{2}}\, p_\star\left(\frac{\beta}{\sqrt{2}}\right)$ for the standard deviation β.[9] This "reference model" is different in that one estimate (the reference, or "target" y_2) directly relates to α, while the other one (the "source" y_1) is associated with an additional offset to account for potential bias. The *shrinkage estimates* of θ_i, however, can be shown to be identical in both models as long as a uniform prior for the overall mean effect (μ) is used.[9] The reference model may be considered a variation of *Pocock's bias model* or, more generally, a *bias allowance model*.[13,17,24] The β parameter, which only differs from τ by a scaling factor of $\sqrt{2}$, may also help motivating a heterogeneity prior, as it directly relates to the expected difference between θ_2 and θ_1, without reference to a common overall mean μ.

2.5. The power prior connection

There is also a connection to a so-called *power prior*, which has been proposed as an approach for deliberate down-weighting of prior information. It is intended for a prior distribution that itself results as a posterior, and the power prior results from applying an exponent a_0 (with $0 \leq a_0 \leq 1$) to its likelihood contribution.[24,25] When conditioning on a fixed τ value, the (conditional) MAP prior is normal with moments given in (2.4) and (2.5); in particular, note that τ^2 acts additively on the "plain" variance (s_1^2). In this context, a power prior with fixed exponent a_0 on the other hand would correspond to a $\text{Normal}\left(y_1,\ \frac{s_1^2}{a_0}\right)$ distribution, where the (inverse) exponent acts multiplicatively on the variance. Both MAP and power prior then are identical if $a_0 = \left(2\frac{\tau^2}{s_1^2} + 1\right)^{-1}$.[9,26,27] It is interesting to note that the relationship between a_0 and τ here depends on the ratio $\frac{\tau}{s_1}$; Pawel et al. (2024)[27] point out that the exponent a_0 directly relates to the "relative" heterogeneity as expressed though the popular I^2 statistic,[28] which in this case simply equals $I^2 = \frac{\tau^2}{\tau^2+s_1^2}$. The exponent may then directly be expressed as a function of the corresponding I^2 as $a_0 = \frac{1-I^2}{1+I^2}$. While a prior probability distribution for τ is readily motivated with reference to the effect scale of y_i and θ_i,[12] specification of a fixed α_0 value remains tricky, and a prior distribution may be even harder to motivate, as it would implicitly relate to the I^2 scale. On the other hand, through the above functional correspondence, any prior for the heterogeneity τ immediately implies a corresponding distribution for the exponent a_0; an example is illustrated in Appendix A.5.

3. Two practical applications

3.1. Pediatric Alport example

3.1.1. Application

Gross et al. (2020)[29] performed a randomized controlled trial (RCT) in Alport syndrome to investigate the effects of ramipril, an angiotensin-converting enzyme inhibitor (ACEi). Recruitment of participants to the RCT was hampered by the rare and pediatric nature of the disease, and so the analysis of the RCT had been planned with the inclusion of observational data from an open-label arm and a natural disease cohort.[30] *Time to disease progression* was a co-primary endpoint, and from the observational data, a hazard ratio (HR) of 0.53 [0.22, 1.29] was estimated based on 70 patients. Only 20 patients entered into the RCT, and an HR of 0.51 [0.12, 2.20] was estimated. The data are also summarized in Table 1.

The analysis was then performed by jointly considering both log-HR estimates in an NNHM, anticipating a reasonable amount of heterogeneity between them (expressed through a half-Normal(0.5) prior for τ), and deriving a shrinkage estimate for the RCT effect (θ_2). The resulting estimate then was substantially more precise than if the RCT data were considered in isolation; the HR was estimated at 0.52 [0.19, 1.39].[29]

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

Table 1. *Alport example data from Gross et al. (2020).*[29]

				log-HR	
i	Estimate	n_i	Hazard ratio (HR)	y_i	s_i
1	Observational	70	0.53 [0.22, 1.29]	−0.635	0.451
2	RCT	20	0.51 [0.12, 2.20]	−0.673	0.742

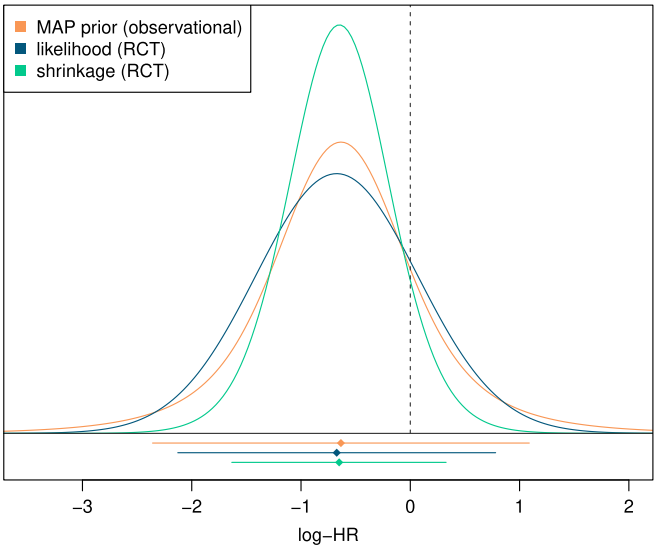


Figure 1. *Illustration of MAP-prior, likelihood, and (shrinkage-) posterior for the Alport example discussed in Section 3.1.*[29] *The horizontal lines at the bottom indicate point estimates and corresponding 95% intervals.*

To make the flow of information transparent, we may now derive the corresponding MAP prior reflecting the information contributed by the observational data. Figure 1 illustrates MAP-prior, likelihood, and posterior (shrinkage estimate) for the Alport example. The MAP prior here has a mean of $y_i = -0.63$ and a variance of $s_1^2 + 2\,\mathrm{E}[\tau^2] = 0.45^2 + 2 \times 0.5^2 = 0.84^2$. While the observational sample size was 70 patients, the MAP prior's effective sample size ($\mathrm{ESS}_{\mathrm{ELIR}}$) is at only 26 patients (i.e., 37% of originally 70 actual patients).[23] The eventual shrinkage interval is only 67% as wide as the original, implying a substantial "effective gain in sample size"[9]; such a precision increase would otherwise have required more than doubling the sample size (by the addition of 24 extra patients). So in this case, the absolute ($\mathrm{ESS}_{\mathrm{ELIR}}$) estimate matches well the observed precision gain.

3.1.2. Variations of the MAP prior

The heterogeneity (τ) prior was specified as a half-Normal(0.5) distribution, which is a reasonably conservative choice for endpoints such as HRs, as it covers "reasonable" and up to "fairly high" levels of heterogeneity ($\tau \leq 1$) and leaves a small prior probability for "fairly extreme" amounts ($\tau > 1$).[12,31] Since conclusions heavily depend on the heterogeneity prior settings, it may however be interesting to investigate the effects of a range of reasonable alternative specifications; in particular, we will consider different prior scales and different distribution families. Among the various assumptions implemented in the analysis, a "too optimistic" heterogeneity prior (favoring small heterogenity) might yield

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

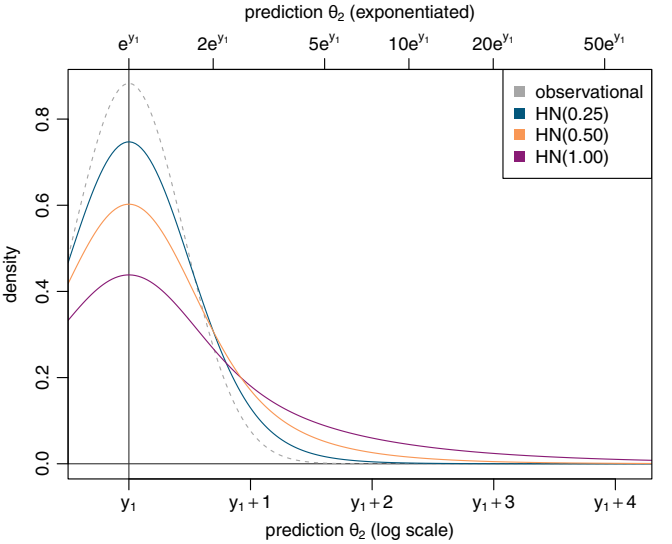


Figure 2. *Illustration of the resulting MAP-prior for varying heterogeneity prior scales. The dashed line indicates the likelihood of the observational data alone for comparison.*

results inappropriately close to a common-effect analysis, while overly "pessimistic" or "conservative" assumptions may on the other hand eventually lead to very little borrowing of information.

Assuming that $s_1 = 0.451$ (as in the present example, see Table 1), we can illustrate the resulting MAP prior when varying the heterogeneity prior scale. Figure 2 shows the likelihood of the observational estimate along with the corresponding MAP priors for half-normal heterogeneity priors with scales 0.25, 0.50, and 1.00. Increasing the heterogeneity prior scale yields a MAP prior that becomes increasingly wider than the plain likelihood alone. In the present case, the effect scale was a logarithmic HR, and on the logarithmic scale, the original MAP prior (based on a half-Normal(0.5) heterogeneity prior) covers a range of $y_1 \pm 1.72$ with 95% probability. On the exponentiated scale (see also the top axis in Figure 2), a difference of 1.72 in log-HR would correspond to a 5.6-fold larger HR. If one switched to a half-Normal(0.25) or half-Normal(1.0) prior instead, the range would change to $y_1 \pm 1.13$ or $y_1 \pm 3.18$ instead, corresponding to multiplicative factors of 3.1 or 24.0, respectively.

Half-normal distributions are a common and obvious choice as heterogeneity priors, possible reasons may be familiarity and availability, as well as a "flat" shape near the origin and a rather short tail.[12] Variations of the distribution family commonly do not alter conclusions dramatically as long as they cover a similar range, as manifested, for example, in a common prior median.

MAP priors corresponding to alternative specifications to a half-Normal(0.5) for the heterogeneity are illustrated in Figure 3. A range of distribution families is used, with their scale parameters specified such that all correspond to a common prior median for τ (of 0.34). These different heterogeneity prior families are also shown in Appendix A.3. The resulting MAP prior densities themselves are hard to distinguish. Differences are more noticeable when focusing on the tail behavior, for example, considering cumulative distribution functions (as shown in Figure 4) or logarithmic densities (shown in Figure 5). One can see that heavier-tailed heterogeneity priors also yield correspondingly heavier-tailed MAP-priors. The heterogeneity and the corresponding MAP priors' properties are also summarized and compared in Table 2 in terms of prior quantiles and effective sample sizes (ESS_{ELIR}).

It is sometimes also instructive to observe the effects of variations of the prior on the resulting estimates; for example, varying the heterogeneity prior scale allows for a sensitivity (or tipping point) analysis. Such an analysis is shown in Appendix A.6; the amount of borrowing is reflected in the shrinkage interval's width, but in the present example, inference would not change qualitatively, and a log-HR of zero always remains included.

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

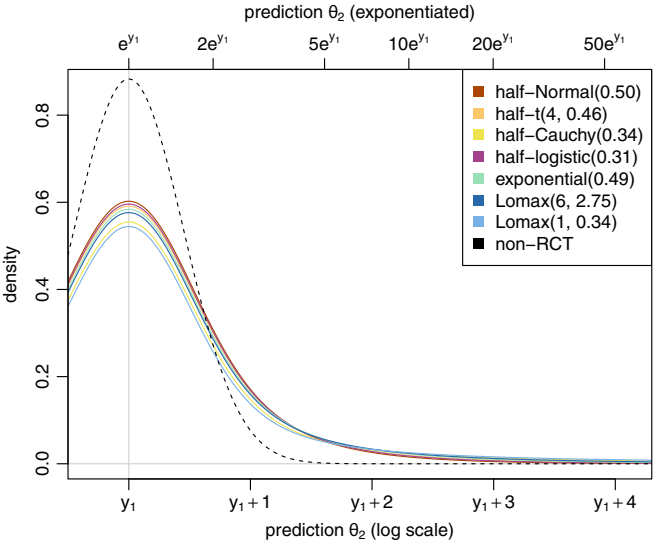


Figure 3. *Illustration of MAP-prior's dependence on the heterogeneity prior distribution family. The different heterogeneity priors shown here all share the same prior median.*

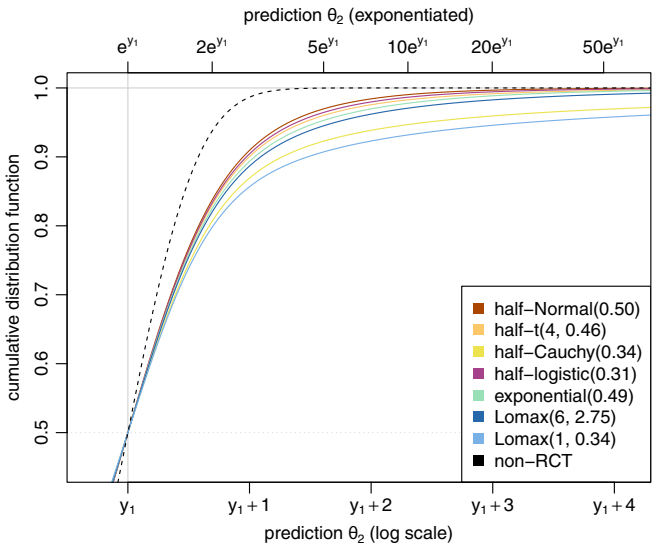


Figure 4. *The MAP-priors' cumulative distribution functions corresponding to the densities shown in Figure 3.*

3.2. Heart failure example

The SPIRIT-HF trial has been designed in order to test the efficacy of spironolactone in patients with heart failure (HF).[32] Spironolactone is expected to reduce cardiovascular mortality as well as hospitalizations due to HF, and had previously been investigated in the TOPCAT trial.[33] Both studies refer to the composite of (recurrent) HF hospitalization and cardiovascular death as the primary endpoint to evaluate treatment efficacy. Despite a sizeable sample size of 3,445 patients and a mean follow-up duration of more than three years, the TOPCAT trial failed to demonstrate statistical significance; the estimated HR was at 0.89 (0.77, 1.04) ($p = 0.14$).[33]

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

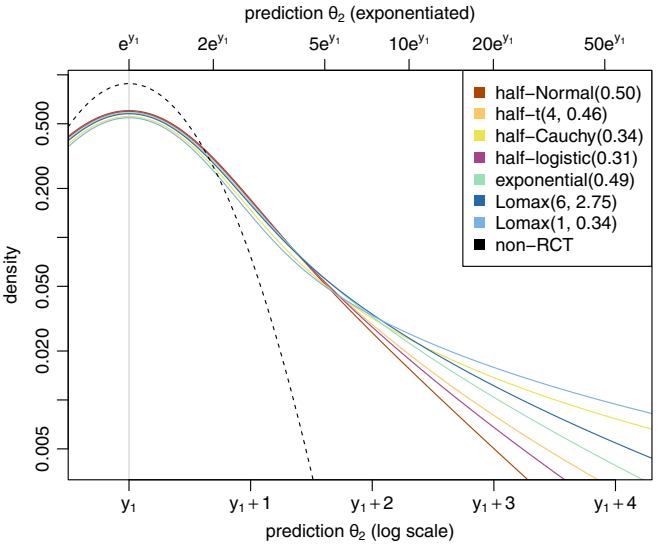


Figure 5. *The MAP-priors' densities on a logarithmic scale (see also Figure 3). Note that the likelihood for the observational data alone follows a parabola shape here, while the corresponding MAP priors are clearly much heavier-tailed.*

Table 2. *Summaries of MAP priors resulting from several settings for the heterogeneity (τ) prior. The half-normal(0.5) prior is contrasted with half-normal priors of differing scale, as well as with priors of differing distributional families, but with matching prior medians. Note that in the context of the present example, the MAP prior's domain corresponds to logarithmic hazard ratios (log-HRs). Quantiles are centered at y_1.*

τ Prior				Standard	Quantile ($\theta_2 - y_1$)		
Family	Scale	Median	ESS_{ELIR}	Deviation	95.0%	97.5%	99.5%
Half-normal	0.50	0.34	26.6	0.84	1.32	1.72	2.72
Half-normal	0.25	0.17	45.7	0.57	0.93	1.13	1.62
Half-normal	1.00	0.67	12.8	1.48	2.35	3.18	5.19
Half-$t_{\nu=4}$	0.46	0.34	25.3	1.02	1.45	1.98	3.58
Half-Cauchy	0.34	0.34	23.4		2.45	4.85	24.02
Half-logistic	0.31	0.34	25.8	0.91	1.39	1.85	3.09
Exponential	0.49	0.34	24.5	1.07	1.56	2.19	3.96
Lomax($\alpha = 6$)	2.75	0.34	24.0	1.31	1.70	2.50	5.05
Lomax($\alpha = 1$)	0.34	0.34	23.1		3.29	7.05	37.17

The analysis of the new SPIRIT-HF trial meanwhile is being planned, and may take into consideration the evidence already generated in the TOPCAT trial. One idea may be to derive a *shrinkage estimate*, anticipating some between-study heterogeneity, and dynamically borrowing information from the earlier study based on the corresponding MAP prior.[9,10] For the between-study heterogeneity τ, use of a half-Normal(0.25) prior may be appropriate. The heterogeneity prior may be motivated referring to anticipated levels of heterogeneity based on general considerations[12] or using empirical evidence, in particular in view of the similar study designs and the effect measure being a log-HR.[34]

In the present case, analysis is based on the logarithmic HR; the HR estimated in the TOPCAT trial corresponds to a log-HR of −0.117 with a standard error of 0.077. The corresponding TOPCAT

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

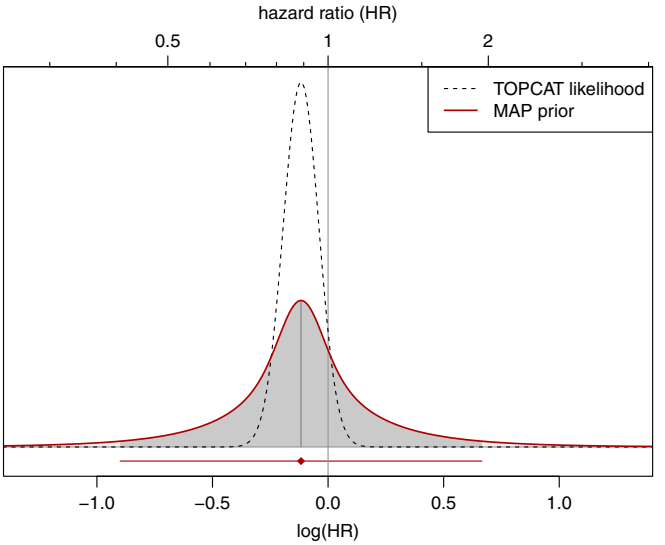


Figure 6. *Illustration of likelihood and corresponding MAP-prior for the heart failure example, using a half-Normal*(0.25) *prior for* τ. *The horizontal line at the bottom indicates the 95% prediction interval.*

likelihood along with the resulting MAP prior is illustrated in Figure 6. The variance (squared standard error) of the TOPCAT study's estimate was $s_1^2 = 0.077^2$ while for the assumed heterogeneity prior the expected heterogeneity variance is $\mathrm{E}[\tau^2] = 0.25^2$ (see also Table A1), so that the resulting MAP prior's variance (2.6) is $s_1^2 + 2\,\mathrm{E}[\tau^2] = 0.362^2$, and the majority of the variance is due to epistemic uncertainty relating to the anticipated similarity of the TOPCAT and SPIRIT-HF parameters. The 95% prediction interval for the MAP prior is centered at the TOPCAT log-HR estimate and ranges from −0.899 to +0.665, corresponding to HRs in the range [0.407, 1.945]. According to the MAP prior, the probability of a beneficial treatment effect (a log-HR below zero) is 71%. The MAP prior has an $\mathrm{ESS}_{\mathrm{ELIR}}$ of 399, that is, only 12% of the 3,445 actual TOPCAT patients, and 31% of the estimated enrolment of 1,300 SPIRIT-HF patients.[32] This means that the prior derived from the TOPCAT data will not enter the eventual analysis as an additional 3,445 patients (as would be the case if both study populations were pooled naïvely), but instead we expect an accuracy corresponding to a total of some $1,300 + 399$ patients. The SPIRIT-HF study's contribution to its own shrinkage estimate may also be assessed[10]; assuming that both studies show the same dependence of standard error and sample size, the SPIRIT-HF study will account for a minimum of 61% in weight to the eventual effect estimate.

4. Discussion

Despite the seemingly odd notion of a *meta-analysis of a single study*, the use of MAP priors remains completely consistent down to the extreme case of only one data point. The "usual" toolbox remains available, including common prior specifications,[12] computation of ESSs,[23] robustification,[8] as well as common meta-analysis software (e.g., the `bayesmeta` or `RBesT` R packages)[7,35] for practical implementation. In addition, for $k = 1$, there are connections to *bias allowance* and *power prior* models (see Section 2) that may help motivating a MAP approach (or vice versa). MAP priors based on a few estimates are generally rather heavy-tailed, which will ensure robust operating characteristics.[9,10,22] For a few data points in general, and in particular for only a single data point, the prior specification for the heterogeneity parameter τ gains in importance and needs to be particularly well-founded and convincing.[12]

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

While only the normal model (NNHM) was discussed here, the idea also extends to other model families; for example, derivation of a MAP prior would also work for a binomial-normal model (as implemented in the `RBesT` package).[35] Another related approach (with some similarity to the power prior) is given by the *commensurate prior*,[36] which, however, does not constitute a special case of the MAP prior. Empirical MAP priors may then be utilized in different ways, either to simply motivate a reasonable sample size (or other design aspects), or to implement explicit borrowing of historical information.[37,38]

When MAP priors are used to also inform the analysis, it is important to approach the evaluation of operating characteristics from a sensible angle; the naive application of classically "frequentist" measures to judge a Bayesian procedure, in particular when informative priors are involved, will often not provide a meaningful assessment of its actual features.[39–41]

A common concern in the context of the use of historical data is that an informative prior might unduly dominate the eventual analysis; for example, in the HF example application, one might be worried that the much larger TOPCAT trial would swamp the data from the smaller SPIRIT-HF study. However, for the shrinkage estimate of interest here, the second study's contribution is bounded by a minimum of 61% within the suggested setup. This proportion would increase for a more conservative heterogeneity prior specification[12] or when implementing robustification,[8] however, such modeling decisions should probably rather be based on considerations of prior information than on deduced operating characteristics.

Besides considerations of the value of "borrowed" information for a given parameter estimate, MAP priors based on historical data may also be interesting for the *design* of subsequent trials, with or without the eventual use of shrinkage estimation in the final analysis. Historical information may then help determining sensible ranges for nuisance parameters[37] or sample sizes,[42,43] for interim decisions,[8,44] or it may be used in a more comprehensive fashion to ensure a positive joint outcome.[44,45]

Van Zwet et al. (2024) argue that the analysis of a single study should also account for heterogeneity of the treatment effect across studies. Therefore, they propose to consider analyses of individual studies also within an overarching NNHM framework similar to our approach presented here; using informative, empirically motivated priors for both μ and τ, inference may then be focused on the overall mean effect (μ) rather than the study-specific θ_1 even in the analysis of only a single study.[46,47]

Author contributions. Conceptualization: C.R. and T.F.; Methodology: C.R.; Writing original draft: C.R. Both authors approved the final submitted draft.

Competing interest statement. The authors declare that no competing interests exist.

Data availability statement. The data supporting the findings of this study are openly available at *Zenodo* under the URL https://doi.org/10.5281/zenodo.18633334.

Funding statement. Support from the German Centre for Cardiovascular Research *(Deutsches Zentrum für Herz-Kreislauf-Forschung e.V., DZHK)* is gratefully acknowledged (Grant No. 81Z0300108).

References

[1] Gagne JJ, Thompson L, O'Keefe K, Kesselheim AS. Innovative research methods for studying treatments for rare diseases: methodological review. *BMJ*. 2014;349: g6802.
[2] Tudur Smith C, Williamson PR, Beresford MW. Methodology of clinical trials for rare diseases. *Best Pract Res Clin Rheumatol*. 2014;28: 247–262.
[3] Gamalo-Siebers M, Savic J, Basu C, et al. Statistical modeling for Bayesian extrapolation of adult clinical trial information in pediatric drug evaluation. *Pharm Stat*. 2017;16: 232–249.
[4] Morris CN, Lysy M. Shrinkage estimation in multilevel normal models. *Stat Sci*. 2012;7: 115–134.
[5] Gelman A, Hill J. *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Cambridge University Press; 2007.
[6] Fleiss JL. The statistical basis of meta-analysis. *Stat Methods Med Res*. 1993;2: 121–145.
[7] Röver C. Bayesian random-effects meta-analysis using the Bayesmeta R package. *J Stat Softw*. 2020;93: 1–51.
[8] 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: 1023–1032.

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

[9] Röver C, Friede T. Dynamically borrowing strength from another study through shrinkage estimation. *Stat Methods Med Res*. 2020;29: 293–308.
[10] Röver C, Friede T. Bounds for the weight of external data in shrinkage estimation. *Biom J*. 2021;65: 1131–1143.
[11] Lesaffre E, Qi H, Banbeta A, van Rosmalen J. A review of dynamic borrowing methods with applications in pharmaceutical research. *Braz J Probab Stat*. 2024;38: 1–31.
[12] 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: 448–474.
[13] Welton NJ, Sutton AJ, Cooper NJ, Abrams KR, Ades AE. *Evidence Synthesis for Decision Making in Healthcare*. Wiley; 2012.
[14] Iglesias JF, Muller O, Zaugg S, et al. A comparison of an ultrathin-strut biodegradable polymer sirolimus-eluting stent with a durable polymer everolimus-eluting stent for patients with acute ST-segment elevation myocardial infarction undergoing primary percutaneous coronary intervention: rationale and design of the BIOSTEMI trial. *EuroIntervention*. 2018;18: 692–699.
[15] Harari O, Soltanifar M, Verhoek A, Heeg B. Alone, together: on the benefits of Bayesian borrowing in a meta-analytic setting. *Pharm Stat*. 2023;22: 903–920.
[16] Wandel S, Neuenschwander B, Röver C, Friede T. Using phase II data for the analysis of phase III studies: an application in rare diseases. *Clin Trials*. 2017;14: 277–285.
[17] Pocock SJ. The combination of randomized and historical controls in clinical trials. *J Chronic Dis*. 1976;29: 175–188.
[18] Lee SX, McLachlan GJ. Scale mixture distribution. In: Balakrishnan N, Colton T, Everitt B, Piegorsch W, Ruggeri F, Teugels JL, eds. *Wiley StatsRef: Statistics Reference Online*. John Wiley & Sons; 2019. https://doi.org/10.1002/9781118445112.stat08201.
[19] Lindsay BG. Mixture models: theory, geometry and applications. In: Collins J, Kelly P, Donoho MG eds. *NSF-CBMS Regional Conference Series in Probability and Statistics*. Vol. 5. Institute of Mathematical Statistics; 1995. http://www.jstor.org/stable/4153184.
[20] Kirkup L. A guide to GUM. *Eur J Phys*. 2002;23: 483–487.
[21] van der Bles AM, van der Linden S, Freeman ALJ, et al. Communicating uncertainty about facts, numbers and science. *R Soc Open Sci*. 2019;6: 181870.
[22] O'Hagan A, Pericchi L. Bayesian heavy-tailed models and conflict resolution: a review. *Braz J Probab Stat*. 2012;26: 372–401.
[23] Neuenschwander B, Weber S, Schmidli H, O'Hagan A. Predictively consistent prior effective sample sizes. *Biometrics*. 2020;76: 578–587.
[24] Neuenschwander B, Schmidli H. Use of historical data. In: Lesaffre E, Baio G, Boulanger B, eds. *Bayesian Methods in Pharmaceutical Research*. Chapman & Hall / CRC; 2020: 111–137. Chap. 6. https://doi.org/10.1201/9781315180212.
[25] Ibrahim JG, Chen MH. Power prior distributions for regression models. *Stat Sci*. 2000;15: 46–60.
[26] Chen MH, Ibrahim JG. The relationship between the power prior and hierarchical models. *Bayesian Anal*. 2006;1: 551–574.
[27] Pawel S, Aust F, Held L, Wagenmakers EJ. Power priors for replication studies. *TEST*. 2024;33: 127–154.
[28] Higgins JPT, Thompson SG. Quantifying heterogeneity in a meta-analysis. *Stat Med*. 2002;21: 1539–1558.
[29] Gross O, Tönshoff B, Weber LT, et al. 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 Int*. 2020;97: 1275–1286.
[30] Gross O, Friede T, Hilgers R, et al. Safety and efficacy of the ACE-inhibitor ramipril in Alport syndrome: the double-blind, randomized, placebo-controlled, multicenter phase III EARLY PRO-TECT Alport trial in pediatric patients. *ISRN Pediatr*. 2012;2012: 436046.
[31] Friede T, Röver C, Wandel S, Neuenschwander B. Meta-analysis of few small studies in orphan diseases. *Res Synth Methods*. 2017;8: 79–91.
[32] Pieske B. Spironolactone in the treatment of heart failure (SPIRIT-HF). NCT 04727073. ClinicalTrials.gov. 2021. https://www.clinicaltrials.gov/study/NCT04727073.
[33] Pitt B, Pfeiffer MA, Assmann SF, et al. Spironolactone for heart failure with preserved ejection fraction. *N Engl J Med*. 2014;370: 1383–1392.
[34] Lilienthal J, Sturtz S, Schürmann C, et al. Bayesian random-effects meta-analysis with empirical heterogeneity priors for application in health technology assessment with very few studies. *Res Synth Methods*. 2024;15: 275–287.
[35] Weber S, Li Y, Seaman JW, Kakizume T, Schmidli H. Applying meta-analytic-predictive priors with the R Bayesian evidence synthesis tools. *J Stat Softw*. 2021;100: 1–32.
[36] Hobbs BP, Carlin BP, Mandrekar SJ, SD. J. Hierarchical commensurate and power prior models for adaptive incorporation of historical information in clinical trials. *Biometrics*. 2011;67: 1047–1056.
[37] Schmidli H, Neuenschwander B, Friede T. Meta-analytic-predictive use of historical variance data for the design and analysis of clinical trials. *Comput Stat Data Anal*. 2017;113: 100–110.
[38] Muehlemann N, Zhou T, Mukherjee R, Hossain MI, Roychoudhury S, Russek-Cohen E. A tutorial on modern Bayesian methods in clinical trials. *Ther Innov Regul Sci*. 2023;57: 402–416.
[39] Gneiting T, Balabdaoui F, Raftery AE. Probabilistic forecasts, calibration and sharpness. *J Royal Stat Soc B*. 2007;69: 243–268.

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

[40] Cook SR, Gelman A, Rubin DB. Validation of software for Bayesian models using posterior quantiles. *J Comput Graph Stat* 2006;15: 675–692.
[41] Best N, Ajimi M, Neuenschwander B, Saint-Hilary G, Wandel S. Beyond the classical type I error: Bayesian metrics for Bayesian designs using informative priors. *Stat Biopharm Res*. 2025;17: 183–196. DOI: https://doi.org/10.1080/19466315.2024.2342817.
[42] Lindley DV. The choice of sample size. *J Royal Stat Soc D*. 1997;46: 129–138.
[43] Brutti P, De Santis F, Gubbiotti S. Bayesian-frequentist sample size determination: a game of two priors. *Metro*. 2014;72: 133–151.
[44] Neuenschwander B, Roychoudhuri S, Schmidli H. On the use of co-data in clinical trials. *Stat Biopharm Res*. 2016;8: 345–354.
[45] Pawel S, Consonni G, Held L. Bayesian approaches to designing replication studies. *Psychol Methods*. 2023. Advance online publication. https://doi.org/10.1037/met0000604.
[46] Gelman A. Meta-analysis with a single study. In: *Statistical Modeling, Causal Inference, and Social Science*. 2024. https://statmodeling.stat.columbia.edu/2024/11/11/meta-analysis-with-a-single-study/.
[47] van Zwet E, Wiecek W, Gelman A. Meta-analysis with a single study. *Stat Methods Med Res*. 2025;34: 2302–2312. DOI: https://doi.org/10.1177/09622802251380628.

A. Appendix

A.1. Posterior predictive distribution

When an informative Normal$(\mu_\mathrm{p}, {\sigma_\mathrm{p}}^2)$ prior is assumed for the overall mean μ, the posterior predictive distribution for a "new" study-specific mean θ_{k+1} (conditional on a given heterogeneity value) is again normal with moments

$$\mathrm{E}[\theta_{k+1}|y_1,\ldots,y_k,s_1,\ldots,s_k,\tau] = \frac{\frac{\mu_\mathrm{p}}{{\sigma_\mathrm{p}}^2} + \sum_{i=1}^{k}\frac{y_i}{s_i^2+\tau^2}}{\frac{1}{{\sigma_\mathrm{p}}^2} + \sum_{i=1}^{k}\frac{1}{s_i^2+\tau^2}} \tag{A.1}$$

$$\mathrm{Var}(\theta_{k+1}|y_1,\ldots,y_k,s_1,\ldots,s_k,\tau) = \frac{1}{\frac{1}{{\sigma_\mathrm{p}}^2} + \sum_{i=1}^{k}\frac{1}{s_i^2+\tau^2}} + \tau^2 \tag{A.2}$$

implying for the specific case of $k = 1$ that

$$\mathrm{E}[\theta_2|y_1,s_1,\tau] = \mu_\mathrm{p}\frac{s_1^2+\tau^2}{{\sigma_\mathrm{p}}^2+s_1^2+\tau^2} + y_1\frac{{\sigma_\mathrm{p}}^2}{{\sigma_\mathrm{p}}^2+s_1^2+\tau^2} \tag{A.3}$$

$$\mathrm{Var}(\theta_2|y_1,s_1,\tau) = \frac{1}{\frac{1}{{\sigma_\mathrm{p}}^2}+\frac{1}{s_1^2+\tau^2}} + \tau^2 \tag{A.4}$$

(see Röver (2020)[7]). One can already see that in the limiting case of an increasingly vague effect prior ($\sigma_\mathrm{p} \to \infty$), the prior's influence vanishes, and the (conditional) variance increases.

Specification of an improper uniform prior for the overall mean effect μ also leads to a proper posterior; the posterior predictive moments then are of the slightly simpler form

$$\mathrm{E}[\theta_{k+1}|y_1,\ldots,y_k,s_1,\ldots,s_k,\tau] = \frac{\sum_{i=1}^{k}\frac{y_i}{s_i^2+\tau^2}}{\sum_{i=1}^{k}\frac{1}{s_i^2+\tau^2}} \tag{A.5}$$

$$\mathrm{Var}(\theta_{k+1}|y_1,\ldots,y_k,s_1,\ldots,s_k,\tau) = \frac{1}{\sum_{i=1}^{k}\frac{1}{s_i^2+\tau^2}} + \tau^2 \tag{A.6}$$

(see Röver (2020)[7]). In the case of a single study ($k = 1$), these expressions then simplify to

$$\mathrm{E}[\theta_2|y_1,s_1,\tau] = y_1 \quad \text{and} \quad \mathrm{Var}(\theta_2|y_1,s_1,\tau) = s_1^2 + 2\tau^2. \tag{A.7}$$

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

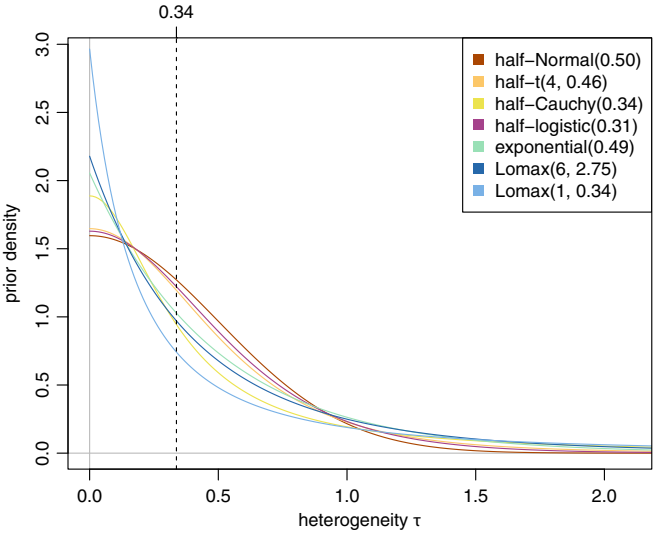


Figure A1. *Illustration of the several heterogeneity priors compared in Section 3.1.2 in terms of their probability density functions. All priors are scaled such that they have a common median of 0.34 (the median of a half-normal(0.5) prior; dashed line).*

A.2. Gain in effective sample size

The *relative* gain in information from a prior may be quantified using the *gain in ESS*, which is based on the relative width of 95% intervals *with* and *without* consideration of the (informative) prior. First, consider the relative width q, the ratio of interval widths (or standard errors) ($q = \frac{\text{width using informative prior}}{\text{width using vague prior}}$). Assuming that standard errors are proportional to the inverse of the square root of the sample size, the gain in ESS then is given by $q^{-2} - 1$. For example, if the informative prior yields an interval that is only half as wide ($q = 0.5$), this would otherwise have required a quadrupled sample size, or a $q^{-2} - 1 = 300\%$ increase. If the interval is 90% as wide ($q = 0.9$), this corresponds to an approximate $q^{-2} - 1 = 23\%$ increase.[9]

A.3. Heterogeneity priors

Figures A1 and A2 show several prior densities and cumulative distribution functions (CDFs) discussed in Section 3.1.2 that are all scaled to a common median (of 0.34, the median of an HN(0.5) distribution).

A.4. Variance of the MAP prior

The prior predictive distribution $p(\theta_2|y_1, s_1)$ is a normal *scale mixture* with (constant) mean y_1 and (conditional) variance $s_1^2 + 2\tau^2$, where τ is distributed according to the specified heterogeneity prior. The mixture distribution's marginal variance results as $\mathrm{Var}(\theta_2|y_1, s_1) = s_1^2 + 2\,\mathrm{E}[\tau^2]$ and depends on the (prior) expectation of the *squared* heterogeneity $\mathrm{E}[\tau^2]$ (see (2.6)). Table A1 summarizes expectations for τ^2 for a range of common heterogeneity priors. Note also the related Table B1 in the appendix of Röver et al. (2021)[12] giving some additional details for the prior distributions shown here.

A.5. Power prior exponent's distribution

For any *fixed* heterogeneity value τ, the (conditional) MAP prior is equivalent to a power prior with exponent $a_0 = \left(2\frac{\tau^2}{s_1^2} + 1\right)^{-1}$ (see also Section 2.5).[9] Through this functional relationship, any prior

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

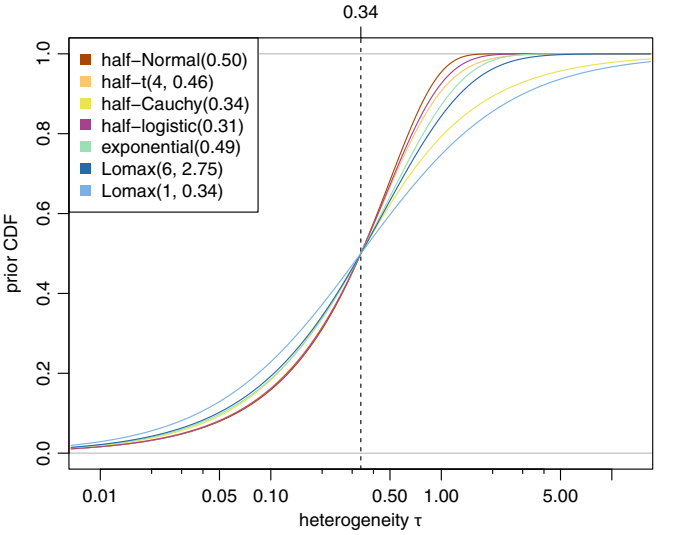


Figure A2. *Illustration of the several heterogeneity priors compared in Section 3.1.2 in terms of their cumulative distribution functions (CDFs). All priors are scaled such that they have a common median of 0.34 (dashed line).*

Table A1. *Expected values of τ^2 based on various common (prior) distributions for τ, depending on their scale parameter* s*. An asterisk (∗) indicates that there is no simple analytical expression. The half-Cauchy prior would be an additional option, but does not have a finite expectation (it is in fact also a special case of the half-Student-*t *prior, with $\nu = 1$ degree of freedom).*

	Standard deviation τ				Variance τ^2	
Distribution	Median	95% Quantile	Expectation		Expectation	
$\tau \sim$ half-Normal(s)	$\approx 0.674s$	$\approx 1.96s$	$\sqrt{\frac{2}{\pi}}\, s \approx 0.798s$		s^2	
$\tau \sim$ half-Student-$t_\nu(s)$	$*$	$*$	$2\sqrt{\frac{\nu}{\pi}}\frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})(\nu-1)}s$	$(\nu > 1)$	$\frac{\nu}{\nu-2}s^2$	$(\nu > 2)$
$\tau \sim$ half-logistic(s)	$\log(3)\, s \approx 1.10s$	$\approx 3.66s$	$\log(4)\, s \approx 1.39s$		$\frac{\pi^2}{3}s^2 \approx 3.29s^2$	
$\tau \sim$ Exp(s)	$\log(2)\, s \approx 0.693s$	$\approx 3.00s$	s		$2s^2$	
$\tau \sim$ Lomax(α, s)	$(2^{\frac{1}{\alpha}} - 1)\, s$	$(20^{\frac{1}{\alpha}} - 1)\, s$	$\frac{1}{\alpha-1}s$	$(\alpha > 1)$	$2\frac{\Gamma(\alpha-2)}{\Gamma(\alpha)}s^2$	$(\alpha > 2)$
$\tau \sim$ Uniform$[0, s]$	$\frac{s}{2}$	$0.95s$	$\frac{s}{2}$		$\frac{1}{3}s^2$	

density $p_\star(\tau)$ for the heterogeneity implies a corresponding prior for the exponent a_0 with probability density function

$$p(a_0) \;=\; \frac{s_1}{2\sqrt{2}}\,\frac{\sqrt{\frac{a_0}{1-a_0}}}{a^2}\; p_\star\!\left(s_1\sqrt{\frac{1-a_0}{2\,a_0}}\right). \tag{A.8}$$

Figure A3 illustrates such densities for an example value of $s_1 = 0.451$ (as in the Alport example from Section 3.1). The prior densities for a_0 are shown for half-normal priors with scales 0.25, 0.50, and 1.00 (corresponding to the cases also illustrated in Figure 2).

A value of $a_0 = 1$ for the exponent corresponds to *full borrowing*, while smaller values imply increasing degrees of discounting of prior information. While the half-Normal(0.25) prior places a substantial share of prior probability near $a_0 = 1$, larger prior scale parameters correspond to less a-priori expected borrowing, eventually resulting in bimodal priors for a_0.

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

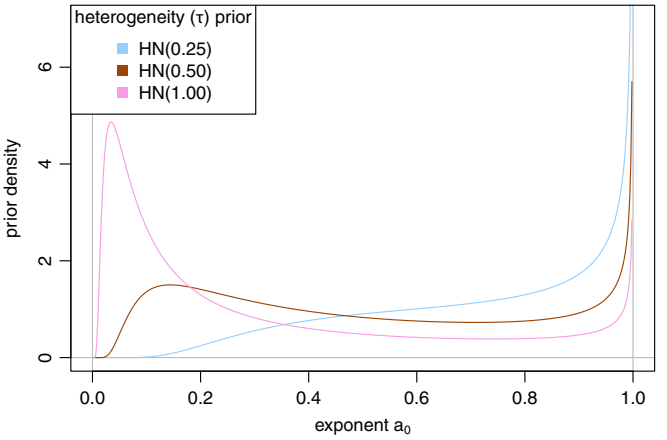


Figure A3. *Illustration of prior distributions for the power prior exponent a_0 corresponding to certain prior distributions assumed for the heterogeneity τ (and $s_1 = 0.451$).*

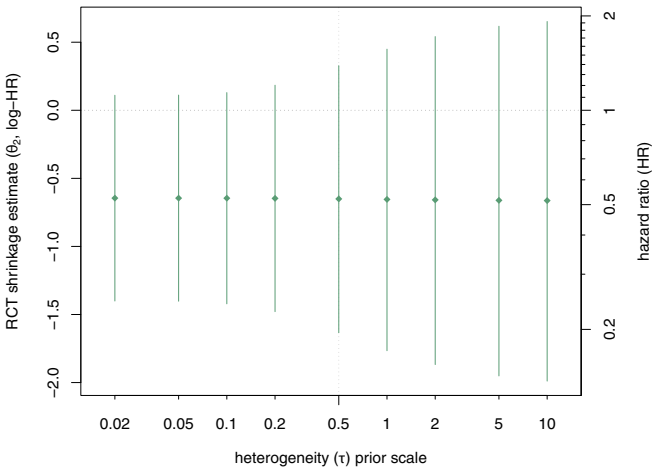


Figure A4. *Illustration of the effect of varying the (half-normal) heterogeneity prior's scale on the resulting RCT shrinkage estimate from Section 3.1.*

Note that (as elaborated in Section 2.5) the mapping between τ and a_0 always depends on the "source" study's standard error (s_1). A prior for τ may be motivated independent of s_1, while specification of a value (or distribution) for a_0 may have odd consequences when varying the source study's size or precision (s_1).

A.6. Sensitivity analysis (Alport example)

In the Alport example application from Section 3.1, varying the half-Normal heterogeneity prior's scale parameter affects the precision of the corresponding MAP prior, and with that, the eventual amount of borrowing from the observational data. Figure A4 shows how the resulting shrinkage estimate for the RCT is affected. A smaller prior scale leads to more borrowing and hence a shorter shrinkage interval eventually approaching the common-effect estimate. A larger prior scale leads to less borrowing, with the shrinkage interval eventually approaching the interval based on the RCT data alone. Since a log-HR of zero always remains included, in this example, there is no "tipping point" for the heterogeneity's prior scale parameter.

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

A.7. Example R *code*

A.7.1. Alport example

```
# data from:
#   Gross et al. (2020)
#   https://doi.org/10.1016/j.kint.2019.12.015
#   Figure 2(b)

GrossEtAl2020 <- cbind.data.frame("data"     = c("observational", "RCT"),
                                  "patients" = c(70, 20),
                                  "events"   = c(29, 8),
                                  "hr"       = c(0.53, 0.51),
                                  "lower"    = c(0.22, 0.12),
                                  "upper"    = c(1.29, 2.20))
GrossEtAl2020$loghr <- log(GrossEtAl2020$hr)
GrossEtAl2020$se    <- (log(GrossEtAl2020$upper)-log(GrossEtAl2020$lower))/
                       (2*qnorm(0.975))

# show data (see Table 1):
GrossEtAl2020

# convert effect measures to "escalc" object:
library("metafor")
es <- escalc(measure="GEN",
             yi=loghr, sei=se, ni=patients,
             slab=data, data=GrossEtAl2020)


############################################################
# perform joint meta-analysis of both estimates (MAC approach):
library("bayesmeta")
bm <- bayesmeta(es,
                tau.prior=function(t){dhalfnormal(t, scale=0.50)})

# show data as well as overall and shrinkage estimates:
forestplot(bm, predict=FALSE,
           xlab="log-HR",
           txt_gp = fpTxtGp(ticks = gpar(cex=1), xlab = gpar(cex=1)))

# check out shrinkage estimate(s):
bm$theta  # (log-HRs)
exp(bm$theta[c("median","95% lower", "95% upper"),"RCT"])  # (HR)


############################################################
# meta analysis of observational data only (MAP approach):
map <- bayesmeta(es[1,],
                 tau.prior=function(t){dhalfnormal(t, scale=0.50)})

# check out predictive distribution (MAP prior):
map$summary[,"theta"]  # (log-HR)


############################################################
# illustrate prior, likelihood, posterior / shrinkage estimate
```

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

```
# (see Figure 1):
colvec <- c("map"    = "red3",
            "likeli" = "blue",
            "shrink" = "green3")
x <- seq(-4, 2.5, length=200)
plot(c(-3.5, 2.0), c(0.0, 0.85), type="n",
     main="Alport example (Gross et al., 2020)",
     xlab="log-HR", ylab="")
# MAP prior (predictive density):
lines(x, map$dposterior(theta=x, predict=TRUE), type="l", col=colvec[1])
# RCT likelihood:
lines(x, dnorm(x, mean=es$yi[2], sd=sqrt(es$vi[2])), type="l", col=colvec[2])
# posterior (=shrinkage estimate):
lines(x, bm$dposterior(theta=x, individual="RCT"), type="l", col=colvec[3])
abline(h=0, col="grey20")
lines(c(0,0), c(0, 1), col="grey20", lty=2)
legend("topleft", c("MAP prior (observational)", "likelihood (RCT)",
                    "shrinkage (RCT)"),
       col=colvec, pch=15)
```

A.7.2. Heart failure example

```
# TOPCAT study (Pitt et al., 2024) quotes a HR of:
# 0.89 (0.77, 1.04)

# log-HR:
log(0.89)  # = -0.117

# standard error:
(log(1.04) - log(0.77)) / (2*qnorm(0.975))  # = 0.077

# derive MAP prior:
library("bayesmeta")
bm <- bayesmeta(y=-0.117, sigma=0.077, label="Topcat",
                tau.prior=function(t){dhalfnormal(t, scale=0.25)})

# (MAP-) prior predictive interval:
round(rbind("log-HR"=bm$summary[c(2,5:6), "theta"],
            "HR"=exp(bm$summary[c(2,5:6), "theta"])), 3)
# prior probability of HR<0:
bm$pposterior(theta=0, predict=TRUE)

##########################################
# illustrate TOPCAT likelihood
# and associated MAP prior (see Figure 6):
x <- seq(-1.8, 1.8, length=200)
plot(c(-1.3, 1.3), c(0.0, 5.18), type="n",
     xlab="log-HR", ylab="")

# TOPCAT likelihood:
lines(x, dnorm(x, mean=-0.117, sd=0.077), lty=2)
# MAP prior:
lines(x, bm$dposterior(theta=x, predict=TRUE), col="red3", lwd=1.5)
# axes, legend:
abline(h=0, v=0, col="grey60")
```

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press

```
legend("topright", lty=c(2,1), col=c("black", "red3"), lwd=c(1,1.5),
       c("TOPCAT likelihood", "MAP prior"))

# compute TOPCAT's unit information standard deviation (UISD):
# (standard error was 0.077, based on 3445 patients)
uisd(n=3445, sigma=0.077)  # = 4.5

# compute MAP prior's effective sample size (ESS_ELIR):
ess(bm, uisd=4.5)  # = 399 patients
```

Cite this article: Röver, C, Friede T. Meta-analytic-predictive priors based on a single study. *Research Synthesis Methods*. 2026;17: 797–815. https://doi.org/10.1017/rsm.2026.10081

https://doi.org/10.1017/rsm.2026.10081 Published online by Cambridge University Press