White Rose Research Online

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk
Universities of Leeds, Sheffield and York

Deposited via The University of York.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/245184/

Version: Published Version

Article:
Yiu, Sean, Galactionova, Katya, Yuen, Steven et al. (2026) Bayesian Dynamic Borrowing to Enhance Evidence for New Therapies. Pharmacoeconomics. ISSN: 1179-2027

https://doi.org/10.1007/s40273-026-01656-7

Reuse
This article is distributed under the terms of the Creative Commons Attribution-NonCommercial (CC BY-NC) licence. This licence allows you to remix, tweak, and build upon this work non-commercially, and any new works must also acknowledge the authors and be non-commercial. You don't have to license any derivative works on the same terms. More information and the full terms of the licence here:
https://creativecommons.org/licenses/

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

UNIVERSITY OF LEEDS
University of Sheffield
UNIVERSITY of York

PharmacoEconomics
https://doi.org/10.1007/s40273-026-01656-7

PRACTICAL APPLICATION

Bayesian Dynamic Borrowing to Enhance Evidence for New Therapies

Sean Yiu[1] · Katya Galactionova[2] · Steven Yuen[3] · Chris Skedgel[4] · Isabelle Durand Zaleski[5] · Maarten Postma[6] · Mark Sculpher[7] · Keith R. Abrams[7,8]

Received: 6 January 2026 / Accepted: 2 August 2026
© The Author(s) 2026

Abstract

Estimating and interpreting treatment effects (TE) for rare or delayed clinical outcomes is often challenging. To address this, researchers may incorporate additional evidence sources, including historical trial data and concurrent information from intermediate outcomes. In this article, we present Bayesian dynamic borrowing (BDB) as a principled framework for integrating such data while maintaining control of bias and Type I error. Using hypothetical trials of a novel high-efficacy therapy for multiple sclerosis, we provide a step-by-step demonstration of how BDB can be used to combine an imprecise TE estimate for a final outcome with a prediction derived from historical data and information on a concurrent intermediate outcome. Our illustration includes calibration of BDB to meet desired Type I error and power properties, and sensitivity analyses to assess robustness to assumption violations. We also discuss key considerations for applying BDB in regulatory decision making and health technology assessment contexts.

1 Introduction

In rare or slowly progressing diseases, or in settings where placebo use is ethically constrained, conducting large multi-arm randomised controlled trials is often impractical [1–3]. Consequently, limited trial data can increase uncertainty in treatment effect (TE) estimates, potentially delaying regulatory decisions and restricting patient access to effective therapies [4]. To improve trial efficiency under such constraints, clinical programs may rely on surrogate or composite endpoints or adopt innovative study designs [5–7].

Bayesian approaches are increasingly used to strengthen TE estimation and interpretation [8–16], a trend reinforced

Key Points for Decision Makers

Leveraging supplementary data sources can enhance the estimation and interpretation of treatment effects of new therapies for rare or delayed outcomes.

Bayesian dynamic borrowing (BDB) is a principled approach to integrate supplementary data sources, while dynamically adjusting the degree of borrowing based on similarity to the primary data source.

BDB functions as a flexible Bayesian learning model that can be implemented with only aggregate-level data and without complex statistical algorithms.

BDB can significantly increase the precision of treatment effect estimators without introducing meaningful bias and Type I error inflation.

✉ Sean Yiu
sean_yiu@hotmail.com

1 Astellas Pharma Europe Ltd, Addlestone, UK
2 F. Hoffmann-La Roche Ltd, Basel, Switzerland
3 Genentech, Inc., South San Francisco, USA
4 Office of Health Economics, London, UK
5 Université de Paris, CRESS, INSERM, INRA, URCEco, AP-HP, Hôpital de l'Hôtel Dieu, 75004 Paris, France
6 University of Groningen, Groningen, Netherlands
7 University of York, York, UK
8 University of Warwick, Coventry, UK

Published online: 04 September 2026

by recent US Food and Drug Administration (FDA) draft guidance on Bayesian methods in clinical trials [17]. These approaches can support, for example, augmentation of concurrent control arms with nonconcurrent data in platform trials or sharing information on TEs across related populations in basket trials [17]. By incorporating prior information and adapting to accumulating evidence, Bayesian methods can enhance statistical power relative to traditional frequentist strategies [18, 19].

Among these, Bayesian dynamic borrowing (BDB) provides a structured way to integrate supplementary data while adjusting the amount of borrowing according to its consistency with the primary data source [20–22]. Information from historical trials, real-world data or other data sources can be synthesised into a prior distribution [23–27], which is then combined with the primary data source (e.g. from a new trial) to form the posterior. The 'dynamic' component reflects the method's ability to down-weight supplementary information when it conflicts with the primary data, thereby improving robustness [18, 19].

BDB has been applied in several clinical development scenarios, such as incorporating external adult data into paediatric studies [9–13] and enhancing subgroup TE estimation by borrowing from complementary populations [28]. It has also gained traction in health technology assessment (HTA), including for constructing external control arms in rare diseases [13–16]. While BDB is best applied with individual-level data, it can also be applied with only aggregate-level data and without complex statistical modelling, thereby facilitating reproducibility in evidence synthesis.

To support broader adoption, we outline step-by-step application of BDB using hypothetical multiple sclerosis (MS) trials of a novel therapy. Section 2 details the clinical setting and describes an approach for jointly incorporating historical trial data and concurrent information from an intermediate outcome to refine TE estimates for a rare or delayed final outcome. Section 3 applies this approach to the hypothetical MS application, including an assessment of performance in more general settings using cross-validation to assess bias and simulations to evaluate Type I error. Section 4 concludes with a discussion of implications for future research.

2 Overview of Proposed BDB Approach

2.1 Setting

We consider the setting where data on an intermediate and final outcome will be available from a new trial. However, estimating the TE on the final outcome is expected to be challenging because that outcome is likely to be rare or delayed. Furthermore, historical trials reporting TE estimates and corresponding standard errors (SEs) for the final and intermediate outcomes are available. Our aim is to combine evidence from the new and historical trials to improve estimation of the TE on the final outcome, while safeguarding against bias and inflated Type I error in the presence of prior data conflict.

2.2 Bayesian Dynamic Borrowing (BDB) Modelling Steps and Statistical Framework

Our proposal follows established methodology [18, 19] and is summarised in Fig. 1. In Step (1), we use historical trial data to develop a prediction model for the TE on the final outcome as the dependent variable and the TE on the intermediate outcome as the independent variable, for example by meta-regression. Evaluating this model at the new trial's estimated TE on the intermediate outcome yields a predicted TE on the final outcome. In Step (2), this prediction is encoded in a prior distribution, which is then updated using only the new trial's data on the final outcome to obtain a posterior distribution that synthesises all available evidence. Finally, in Step (3), the posterior distribution is used to draw inferences about the TE on the final outcome in the new trial.

2.2.1 Predictive Step

Implementing Step (1) requires identifying historical trials that report TEs on both the intermediate and final outcomes and determining which of these should be included when building the prediction model for the TE on the final outcome. Key considerations for trial selection include alignment with the new trial in terms of (a) estimand attributes (population, endpoint definitions, handling of intercurrent events), (b) measurement or assessment methods for the outcomes, (c) timing of data collection, and (d) broader contextual factors such as standard of care.

There is an inherent trade-off: strict inclusion criteria may reduce heterogeneity but leave too few trials to support a reliable prediction model, whereas more permissive criteria increase sample size at the expense of greater variability. A pragmatic approach could be to relax certain criteria where appropriate and account for between-trial differences directly within the prediction model.

It is also important to recognise that parameters from the prediction model may be misinterpreted. For example, regression coefficients from a meta-regression relating population-averaged TEs on the intermediate and final outcomes should not be interpreted as individual-level associations, as this can result in ecological bias.

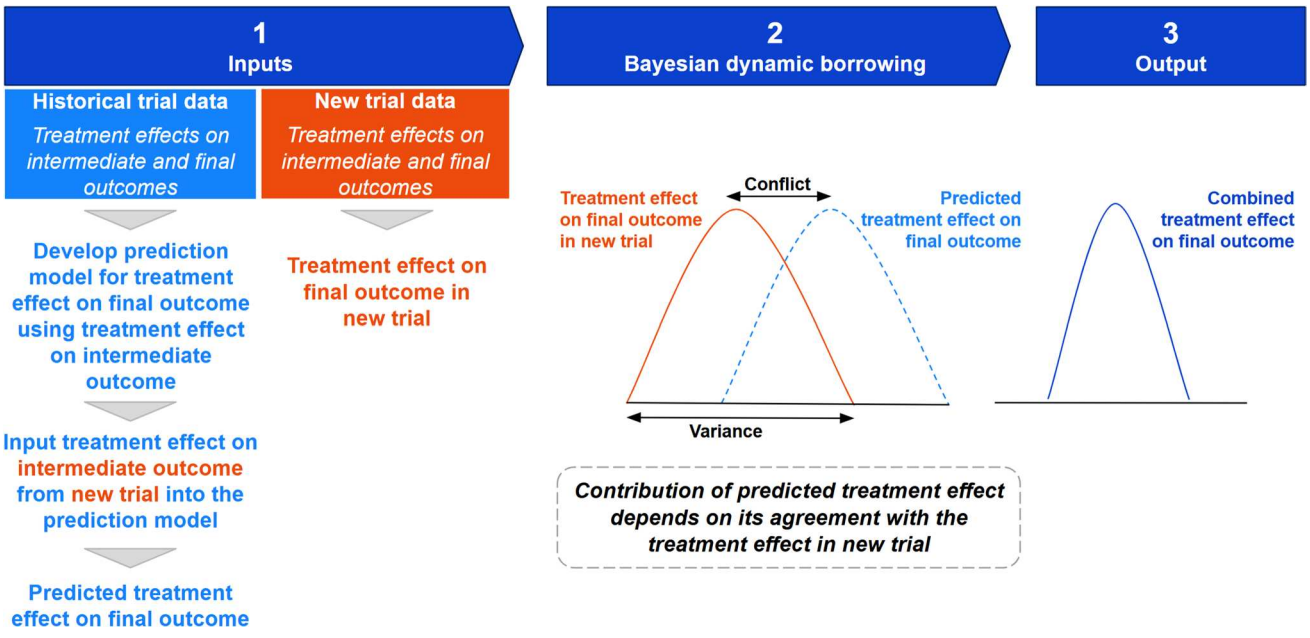

Fig. 1 Graphical description of the implementation of Bayesian dynamic borrowing

2.2.2 Constructing a Prior Distribution and Deriving the Posterior Distribution

A key feature of BDB is the construction of the prior distribution in Step (2). Under traditional Bayesian *static* borrowing, the prior is specified independently of the new trial data, meaning its influence on the posterior is fixed in advance (e.g. based on expert opinion) [20–22]. A limitation of this approach is that the posterior may rely too heavily on the prior. For example, when there is genuine confidence in the TE in the new trial, it is undesirable for a predicted TE to exert significant influence if it is not well aligned with the observed TE. 'Dynamic' borrowing addresses this issue by allowing the prior's influence to vary according to the agreement between the observed and predicted TEs [18, 19].

The prior and posterior distributions from our proposed BDB method is summarised in Fig. 2. Let $\widehat{\theta}_E$ be the TE estimator for the final outcome in the new trial on the linear predictor scale, assumed to follow a normal distribution with mean β_E and variance σ_E^2, where β_E is the true TE. To incorporate data from historical trials, we specify a conjugate prior distribution for β_E, assumed to be normally distributed with mean μ_P and variance σ_P^2. Following [29], the posterior distribution is:

$$\beta_E|\widehat{\theta}_E = \theta_E \sim N\left(\frac{\sigma_P^2}{\sigma_P^2+\sigma_E^2}\theta_E + \frac{\sigma_E^2}{\sigma_P^2+\sigma_E^2}\mu_P, \frac{\sigma_P^2\sigma_E^2}{\sigma_P^2+\sigma_E^2}\right) \quad (1)$$

From Eq. (1), the posterior mean is a weighted linear combination of the new trial TE estimate, θ_E, and the prior mean, μ_P, with greater weight on θ_E when the variance of $\widehat{\theta}_E$ is smaller than the prior distribution. Additionally, notice that the posterior variance is always smaller than the variance of $\widehat{\theta}_E$, especially when σ_P^2 is much smaller than σ_E^2.

To enhance robustness, we propose a mixture prior [30] combining an informative prior, $P_{IP}(\beta_E)$, and a sceptical prior, $P_{SP}(\beta_E)$:

$$\left\{1-\exp\left(-\Upsilon\frac{|\theta_E-\hat{y}|}{\sigma_E}\right)\right\}P_{\mathrm{SP}}(\beta_E), +\exp\left(-\Upsilon\frac{|\theta_E-\hat{y}|}{\sigma_E}\right)P_{\mathrm{IP}}(\beta_E),$$

$$P_{\mathrm{SP}}(\beta_E) \sim N(0,\sigma_S^2), P_{\mathrm{IP}}(\beta_E) \sim N(\hat{y}, \mathrm{Var}(\hat{y}) + \sigma_R^2)$$

Here, $\Upsilon > 0$ is a discount tuning parameter controlling the influence of $P_{\mathrm{IP}}(\beta_E)$, which also depends on the agreement between the observed TE, θ_E, and predicted TE, $\hat{y}$, after scaling by the SE, σ_E. The informative prior is centred at the predicted TE, $\hat{y}$, with variance reflecting both the uncertainty in parameter estimation, $\mathrm{Var}(\hat{y})$, and residual uncertainty σ_R^2 from the prediction model. The sceptical prior is centred at the null effect with variance σ_S^2 that needs to be specified.

To determine Υ, we propose to select a value that will control Type I error at a pre-specified level. This aligns with one of the approaches specified in the FDA guidance document [17], but others can also be considered. In practice, this approach involves conducting simulations under the joint null hypotheses of no TE on both the final and intermediate outcomes (see Sect. 3.2.1) and selecting Υ within the

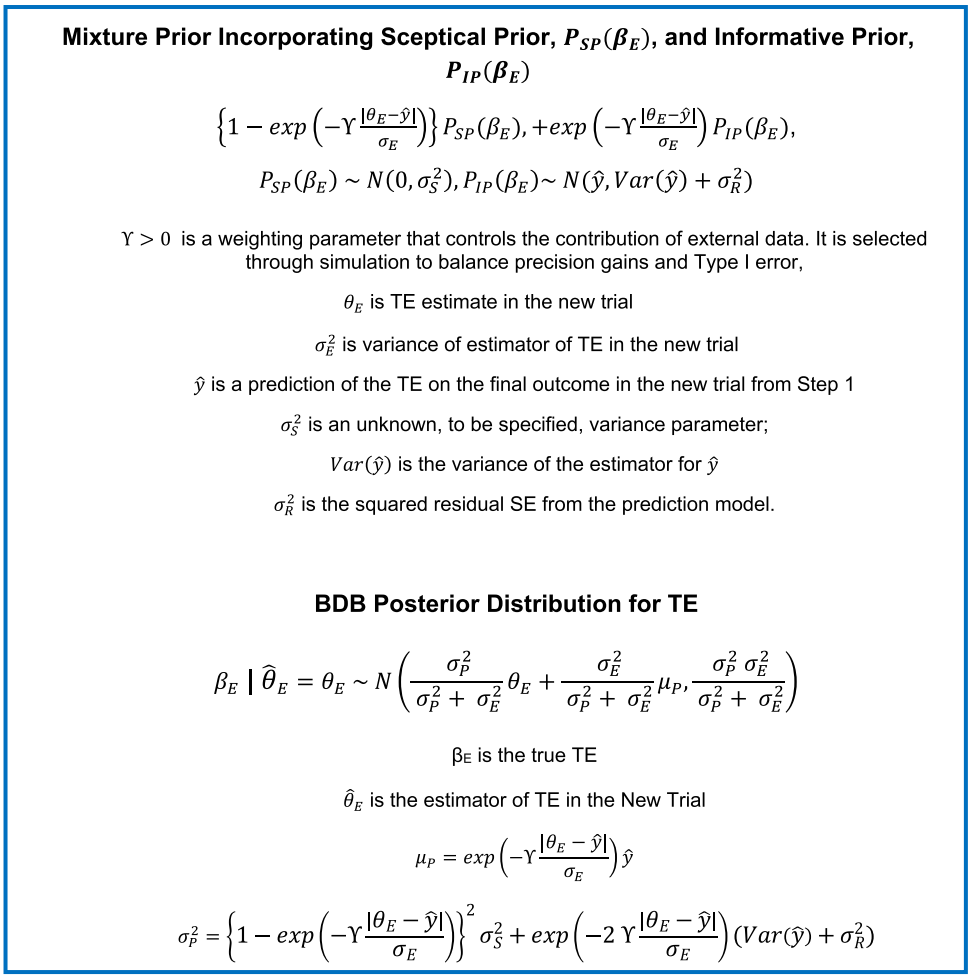

Fig. 2 Schematic describing the prior and posterior distributions for our proposed BDB method. *BDB* Bayesian dynamic borrowing, *SE* standard error, *TE* treatment effect

range of values where Type I error rate is acceptable. Power considerations may also inform the choice of Υ, as modest increases in Type I error might be acceptable if accompanied by large increases in power.

With the proposed mixture prior, the posterior distribution is obtained by substituting the values of μ_P and σ_P^2 into Eq. (1):

$$\mu_P = \exp\left(-\Upsilon\frac{|\theta_E - \hat{y}|}{\sigma_E}\right)\hat{y},$$

$$\sigma_P^2 = \left\{1 - \exp\left(-\Upsilon\frac{|\theta_E - \hat{y}|}{\sigma_E}\right)\right\}^2 \sigma_S^2 + \exp\left(-2\Upsilon\frac{|\theta_E - \hat{y}|}{\sigma_E}\right)\left(\mathrm{Var}(\hat{y}) + \sigma_R^2\right).$$

2.2.3 Sensitivity Analyses

When conducting simulations to select Υ based on Type I error and power, assumptions for the SE of the TE estimators on the intermediate and final outcomes in the new trial, as well as their correlation, must be specified. If these assumptions are implausible, BDB may not perform well and may need to be recalibrated. Therefore, to understand the robustness of BDB, it is essential to assess the performance of BDB in alternative settings where the initial assumptions do not hold. This approach is illustrated in Sect. 3.4.

3 Application of BDB to Hypothetical Trials in Multiple Sclerosis (MS)

Broader availability of high-efficacy therapies in MS has led to a steep decline in event rates for established disability outcomes, including relapses and confirmed disability progression measured by the Expanded Disability Status Scale (CDP-EDSS) [31, 32]. As a result, clinical programs [33–39] have shifted towards using composite confirmed disability progression (cCDP) as the primary efficacy outcome. cCDP combines conventional CDP-EDSS with confirmed worsening on the Timed Twenty-Five Foot Walk Test (CDP-T25FWT) and Nine-Hole Peg Test (CDP-9HPT) [40] (see electronic supplementary material [ESM], Figure S1).

In this application, we consider two trials for a hypothetical MS treatment: one in Relapsing MS (RMS) and one in Primary Progressive MS (PPMS). These trials are powered to detect a TE on cCDP but not on its individual components. However, because CDP-EDSS is the dominant driver of MS-related disability and costs, interpreting the effect on CDP-EDSS is of interest for decision makers [7]. To improve precision in estimating the TE on CDP-EDSS, we apply BDB, leveraging CDP-T25FWT as an intermediate outcome alongside historical data on both outcomes.

3.1 Step 1: Building a Prediction Model for the Treatment Effect (TE) on the Final Outcome Using the TE Estimate for an Intermediate Outcome

CDP-T25FWT and CDP-9HPT, included in cCDP, have been historically collected in MS clinical trials alongside CDP-EDSS. Kappos et al. [41] showed that the TE on CDP-T25FWT can provide reasonably accurate out-of-sample predictions of the TE on CDP-EDSS. The authors fitted an unweighted linear model to nine historical trials (ESM, Table S1) to obtain a predicted effect on the log hazard ratio (HR) scale (ESM, Table S2):

$$\hat{y} = -0.4173 + 0.5425\,\theta_T + 0.3377\,P,$$

where θ_T is an estimate of the log HR for CDP-T25FWT and P is an indicator variable such that $P = 1$ for PPMS or non-active secondary progressive MS trial populations, and $P = 0$ otherwise. An inverse weighted variance linear model was also fitted; however, its predictive performance was inferior to the unweighted linear model in cross validation.

The residual SE from the model was $\sigma_R = 0.1156$ and

$$\mathrm{Var}(\hat{y}) = \begin{pmatrix} 1 & \theta_T & \mathrm{P} \end{pmatrix} \begin{pmatrix} 0.0135 & 0.0267 & -0.0093 \\ 0.0267 & 0.0793 & -0.0145 \\ -0.0093 & -0.0145 & 0.0093 \end{pmatrix} \begin{pmatrix} 1 \\ \theta_T \\ P \end{pmatrix} + \mathrm{Var}\left(\hat{\theta}_T\right)\left(0.0793 + 0.5425^2\right).$$

The derivation of $\mathrm{Var}(\hat{y})$ is provided in the ESM. Note that $Var(\hat{y})$ incorporates the TE estimator variance for the intermediate outcome, $\mathrm{Var}(\hat{\theta}_T)$, as well as the covariance between the TE estimators for the intermediate and final outcome, which in this case is 0.5425; and the estimator variance of this covariance, which in this case is 0.0793.

3.2 Step 2: Prior and Posterior Distribution Construction Under BDB

The variance of the sceptical prior $P_{\mathrm{SP}}(\beta_E)$ is specified to be $\sigma_S^2 = 1$ as this implies a 16% chance of HR $< \exp(-1) = 0.37$ for CDP-EDSS, which is consistent with the observations in [42] where 12% (3/25) of trials demonstrated HR < 0.37.

Following the derivation of BDB in Sect. 2.2, the BDB-derived TE estimate and corresponding 95% credible interval (CRI) are obtained as follows:

$$\log(\text{HR estimate}) = \frac{\sigma_P^2}{\sigma_P^2 + \sigma_E^2}\theta_E + \frac{\sigma_E^2}{\sigma_P^2 + \sigma_E^2}\hat{y}\exp\left(-\frac{\Upsilon|\theta_E - \hat{y}|}{\sigma_E}\right)$$

$$95\%\text{ credible interval} = \log(\text{HR estimate}) \pm z_{0.975}\sqrt{\frac{\sigma_P^2\sigma_E^2}{\sigma_P^2 + \sigma_E^2}},$$

$$\sigma_P^2 = \left\{1 - \exp\left(-\frac{\Upsilon|\theta_E - \hat{y}|}{\sigma_E}\right)\right\}^2 + \exp\left(-\frac{2\Upsilon|\theta_E - \hat{y}|}{\sigma_E}\right)(\mathrm{Var}(\hat{y}) + 0.1156^2),$$

where $z_{1-a/2}$ is the $100(1-\alpha/2)th$ quantile of a standard normal distribution. Once the TE estimate and SE for the final outcome, θ_E and σ_E, and intermediate outcome, θ_T and σ_T, are available, $\hat{y}$ and $\mathrm{Var}(\hat{y})$ can be derived and the only remaining unknown parameter is Υ.

3.2.1 Calibrating Υ Based on Type I Error and Power Considerations

To evaluate Type I error and power, we conducted simulations under both the null hypothesis of no TE and the alternative hypothesis defined by the target product profile. This requires specifying the SEs of the TE estimator for the intermediate and final outcome, as well as their correlation.

For the hypothetical RMS and PPMS trials, the assumed target product profile and SEs for CDP-T25FWT were HR = 0.8 (SE for log HR = 0.114) for RMS and 0.6 (SE for log HR = 0.147) for PPMS, and for CDP-EDSS, HR = 0.6 (SE for log HR = 0.261) for RMS and 0.7 (SE for log HR = 0.182) for PPMS. Under these assumptions, the expected power to detect a TE on CDP-EDSS was 50% in both scenarios. The correlation between the TE estimator for the intermediate and final outcome was set to 0.8, informed by historical trials [41].

For a range of $\boldsymbol{\Upsilon}$ values, we simulated 100,000 pairs of TE estimates for CDP-T25FWT and CDP-EDSS from a bivariate normal distribution using the specified mean, variance and correlation parameters. Type I error and power were then computed as the proportion of simulations where the BDB-derived upper bound of the 95% CRI for the log HR was below zero. The R script used for these simulations is provided in the ESM.

Table 1 summarises the simulation results on Type I error and power across a range of $\boldsymbol{\Upsilon}$ values. As expected, smaller tuning parameter values, which allow greater borrowing

from the informative prior, lead to higher power for detecting TEs, but also increase Type I error. For example, achieving a 15-percentage-point gain in power with BDB over the conventional no-borrowing estimator requires setting Υ to 1 in the RMS setting and 0.8 in the PPMS setting. However, these choices raise the estimated Type I error rate from 2.5% to 6.3% in the RMS setting and 4.3% in the PPMS setting. For the remainder of this example, we assume that this trade-off is acceptable.

3.3 Output

Suppose that the TE estimates defined by the target product profile and expected SEs were observed in the new trial. That is, for CDP-EDSS the estimated TE is HR = 0.6 (95% confidence interval [CI] 0.36–1.0) for RMS and 0.7 (95% CI 0.49–1.0) for PPMS.

3.3.1 Input TE on Intermediate Outcome From a New Trial Into the Prediction Model

TE on CDP-T25FWT from the new trial and the qualifier for the MS subtype are plugged into the predictive meta-regression model above. For the hypothetical RMS and PPMS trials, $\theta_T = \log(0.8)$ and $\log(0.6)$, and $P = 0$ and 1, which results in $\hat{y} = \log(0.584)$ and $\log(0.7)$, and prediction errors of $\sqrt{\mathrm{Var}(\hat{y}) + \sigma_R^2} = 0.154$ and 0.184, respectively.

3.3.2 Estimate TE and 95% Credible Intervals using BDB

For the hypothetical RMS and PPMS trials, BDB results in TE estimates of 0.611 (95% CRI 0.453–0.794) and 0.700 (95% CRI 0.543–0.902), respectively. Therefore, BDB provides essentially the same TE estimates as those based only on the new trial data, but with much smaller estimates of uncertainty.

Table 1 Type I error and power results from applying BDB to 100,000 simulated TEs across a range of Υ values

Υ	RMS trial		PPMS trial	
	Type I error (%)	Power (%)	Type I error (%)	Power (%)
No borrowing	2.5	50	2.5	50
1.5	5.6	62.3	3.3	62.0
1.4	5.7	62.7	3.3	62.5
1.3	5.8	63.2	3.5	62.9
1.2	6.0	63.7	3.6	63.3
1.1	6.2	64.4	3.7	63.9
1.0	6.3	65.2	3.8	64.5
0.9	6.7	65.8	4.0	64.9
0.8	7.1	67.0	4.3	65.8

BDB Bayesian dynamic borrowing, *PPMS* primary progressive multiple sclerosis, *RMS* relapsing multiple sclerosis, *TE* treatment effect

3.4 Performance of BDB in Other Settings

3.4.1 Cross-Validation Evaluating BDB on Real Trial Data

To evaluate the performance of BDB on empirical trial data, we conducted 'leave one trial out' cross-validation. Specifically, for each historical trial used in the meta-regression, this consisted of regarding the historical trial (ESM, Table S1) as the new trial. Then, a prediction model was fitted to the data excluding the new trial. Finally, BDB was applied with the selected values of Υ to estimate the TE on CDP-EDSS in the new trial and the corresponding 95% CRI.

Table 2 summarises the results from the cross-validation exercise. Despite some differences between the individual trials, BDB produced TE estimates that were essentially identical to the observed estimates generated from only the new trial data. Predicted effects from meta-regression, on the other hand, varied from trial to trial, highlighting the advantage of dynamic (as opposed to static) Bayesian borrowing in mitigating the influence of overly optimistic/pessimistic predictions.

The impact of BDB on reducing uncertainty can be further appreciated by comparing the standard deviation (SD) of the posterior distribution to the SE. For trials where the observed and predicted effects were not in close agreement (i.e. AFFIRM [43], ASCEND [33], EXPAND [44], and OPERA I [31]) or where there was moderate agreement but the SE was small, indicating that data on CDP-EDSS was already mature (i.e. ORATORIO [45]), BDB led to only modest reductions in uncertainty (the SD of the posterior was 0.9–8.0% smaller than the SE). In contrast, when there was strong agreement (i.e. INFORMS [34] and PROMISE [46]) or moderate agreement, but the SE was not small (i.e. OLYMPUS [47] and OPERA II [31]), larger reductions in uncertainty were achieved with BDB (the SD of the posterior was 11.3–17.3% smaller than the SE).

To further contextualise the results, we translated the reduction in the BDB-derived SD of the TE into the number of events gained (Table 3) using the Schoenfeld approximation formula [48]. Particularly for INFORMS [34], OLYMPUS [47], OPERA II [31], ORATORIO [45] and PROMISE [46], the results suggest that BDB has the effect of greatly increasing the number of observed events, and hence the sample size of these studies.

Table 2 TE estimates on CDP-EDSS from the conventional Cox model without borrowing, meta-regression, and BDB

Trial	Cox model (no borrowing)		Meta-regression		BDB	
	Estimate (95% CI)	SE	Estimate	Prediction error	Estimate (95% CRI)	SD of posterior distribution
AFFIRM [43]	0.46 (0.31–0.68)	0.200	0.58	0.163	0.48 (0.32–0.67)	0.193
ASCEND [33]	1.06 (0.74–1.53)	0.185	0.86	0.134	1.05 (0.75–1.50)	0.177
EXPAND [44]	0.79 (0.65–0.95)	0.097	0.92	0.137	0.79 (0.65–0.95)	0.096
INFORMS [34]	0.88 (0.72–1.08)	0.103	0.90	0.156	0.89 (0.74–1.05)	0.091
OLYMPUS [47]	0.77 (0.55–1.09)	0.174	0.71	0.214	0.77 (0.57–1.04)	0.154
OPERA I [31]	0.57 (0.37–0.90)	0.227	0.46	0.209	0.59 (0.38–0.86)	0.212
OPERA II [31]	0.63 (0.42–0.92)	0.200	0.58	0.205	0.64 (0.45–0.89)	0.177
ORATORIO [45]	0.76 (0.59–0.98)	0.129	0.80	0.159	0.77 (0.60–0.96)	0.119
PROMISE [46]	0.87 (0.71–1.07)	0.105	0.88	0.145	0.88 (0.73–1.03)	0.087

BDB Bayesian dynamic borrowing, *CDP-EDSS* Confirmed Disability Progression-Expended Disability Status Scale, *CI* confidence interval, *CRI* credible interval, *SD* standard deviation, *SE* standard error, *TE* treatment effect

3.4.2 Simulation Experiments Evaluating BDB in Other Hypothetical Settings

We conducted simulation experiments to further assess the performance of BDB with the chosen values for Υ. The data-generating mechanism for the simulations were varied in two ways: (1) changing the correlation between the TE estimators for the intermediate and final outcomes from 0.8 to 0.6, 0.4 and 0.2, and (2) increasing the SEs for both TEs by 5% and 10%. These modifications yielded 12 distinct data generating scenarios.

Table 4 summarises results from applying BDB to 100,000 simulated pairs of TE estimates under each scenario. Across all settings, BDB produced essentially unbiased estimates, demonstrating strong robustness. However, Type I error rates varied. In particular, Type I error increased as the SEs of the TE estimators increased—especially in the RMS setting. This is expected because the originally assumed SEs in RMS trials were larger than those in PPMS trials; therefore, a 5–10% inflation results in a larger absolute increase for RMS trial settings.

Changes in correlation had opposite effects in the RMS and PPMS settings. In the RMS setting, lower correlation led to higher Type I error. When the TE on the intermediate outcome is zero, the predicted TE for the final outcome is moderate; thus, with high correlation, the informative prior receives low weight. As correlation decreases, the likelihood that the observed and predicted TEs on the final outcome align increases, which raises the probability that the informative prior exerts greater influence—thereby increasing Type I error. In contrast, in the PPMS setting, the predicted TE on the final outcome is close to zero when the TE on the intermediate outcome is zero. Hence, at high correlation, the informative prior is likely to dominate. As correlation decreases,

Table 3 The number of events gained by BDB

Trial	Trial randomisation ratio	Variance (SE^2)	Estimated number of events	BDB variance	Estimated number of events with BDB-derived variance	Number of events gained by applying BDB
AFFIRM [43]	2:1	0.040	113	0.037	121	8
ASCEND [33]	1:1	0.034	117	0.031	128	11
EXPAND [44]	2:1	0.009	478	0.009	489	11
INFORMS [34]	1:1	0.011	377	0.008	484	107
OLYMPUS [47]	2:1	0.030	149	0.024	190	41
OPERA I [31]	1:1	0.052	78	0.045	89	11
OPERA II [31]	1:1	0.040	100	0.031	127	27
ORATORIO [45]	2:1	0.017	270	0.014	318	48
PROMISE [46]	2:1	0.011	408	0.007	600	192

BDB Bayesian dynamic borrowing, *SE* standard error

Table 4 Simulation results assessing bias and Type I error under BDB

SE assumption	RMS trial		PPMS trial	
	Bias	Type I error (%)	Bias	Type I error (%)
Correlation between TE estimators for final and intermediate outcome = 0.8				
No change in SE	− 0.01	6.3	− 0.00	4.3
5% increase in SE	− 0.01	7.4	− 0.00	4.3
10% increase in SE	− 0.01	8.2	− 0.00	4.4
Correlation between TE estimators for final and intermediate outcome = 0.6				
No change in SE	− 0.01	7.0	− 0.00	4.0
5% increase in SE	− 0.01	7.9	− 0.00	4.1
10% increase in SE	− 0.01	8.5	− 0.00	4.2
Correlation between TE estimators for final and intermediate outcome = 0.4				
No change in SE	− 0.01	7.6	− 0.00	3.7
5% increase in SE	− 0.01	8.5	− 0.00	3.7
10% increase in SE	− 0.01	9.4	− 0.00	3.8
Correlation between TE estimators for final and intermediate outcome = 0.2				
No change in SE	− 0.00	8.0	− 0.00	3.4
5% increase in SE	− 0.01	8.8	− 0.00	3.4
10% increase in SE	− 0.01	9.5	− 0.00	3.5

BDB Bayesian dynamic borrowing, *PPMS* primary progressive multiple sclerosis, *RMS* relapsing multiple sclerosis, *SE* standard error, *TE* treatment effect

agreement between observed and predicted TEs becomes less likely, reducing the influence of the prior and therefore lowering Type I error.

Overall, BDB remains highly robust in the PPMS setting across positive correlations and SE inflations up to 10%. In the RMS setting, although BDB continues to provide essentially unbiased TE estimates, Type I error inflation can exceed a two-percentage-point increase when SEs increase by at least 5–10% and the correlation drops below 0.4–0.6.

4 Discussion

In an era of increasingly effective treatments, such as those for MS, drawing on the *totality of evidence* is essential for timely and reliable assessments of clinical efficacy. However, evidence varies in credibility, relevance, and quality; therefore, borrowing from additional evidence sources must be selective and grounded in methods that account for such differences. This need is heightened in disease areas where epidemiology and natural history have evolved while traditional legacy outcome measures remain unchanged. As trials increasingly adopt more sensitive intermediate outcomes, BDB provides a principled framework for integrating concurrent intermediate-outcome data with historical trial evidence to improve estimation of TEs on legacy outcomes, helping bridge the gap toward broader acceptance of newer measures.

In the application to real MS trials (Table 2), BDB improved the precision of TE estimators without materially changing the estimate of effect. Simulations further showed that BDB can limit bias and Type I error inflation under the null (Tables 1, 4). This performance arises from two key features: weighting the predicted TE according to its consistency with the observed TE and calibrating the relative influence of the observed and predicted TEs to maintain Type I error control. Together, these mechanisms underscore that BDB's effectiveness depends on the accuracy of the predicted TE and on careful tuning of the borrowing parameter.

Our application of BDB used aggregate rather than patient-level data, allowing synthesis across trials using only summary results. If, however, patient-level data were available, which might be the case for some, but not all trials [49], then such data can be used to harmonise endpoint definitions, align follow-up periods, and provide richer covariate information to build the prediction model. We also specified a single conjugate prior to obtain a closed-form posterior and simplify implementation. Therefore, our proposal needs to be extended if prior distributions are required for other parameters. In this setting, a more flexible model fitting approach, such as Markov Chain Monte Carlo (MCMC) [50], will likely be necessary.

In the context of HTA, estimated TEs for new therapies typically need to be incorporated into network meta-analyses to support indirect comparisons and establish a relative ranking of treatments. However, using BDB-derived TEs in this setting may be problematic, as it could introduce unfair comparisons unless TEs for other treatments in the network are also generated using similar methods. Even when the inclusion of BDB-derived estimates can be justified, additional challenges arise because historical trial data may be used twice—once in the direct evidence for older therapies and again to inform the BDB estimates for new therapies. This duplication undermines the assumption that TEs arise from independent random variables, complicating statistical inference. Further methodological work on how to appropriately incorporate BDB-derived TEs into network meta-analysis, and more generally to address the problem of double counting [51], would therefore be valuable.

On the other hand, where clinical benefit is the primary basis for HTA assessment, BDB may offer added value. In the application considered here, BDB complements a statistically powered composite outcome by clarifying the treatment effect on a clinically meaningful but underpowered component. This enhances the interpretability of the

composite outcome and reduces uncertainty around the direction and magnitude of effect on the component. With the primary efficacy claim established with the powered composite endpoint, a small increase in bias may be acceptable in exchange for greater certainty on secondary elements.

The recognition of the relevance of Bayesian methods for HTA [52–54] and the increasing openness of regulatory and HTA agencies to use of these methods [13, 17, 55–57] suggests that BDB is likely to be accepted by the health authorities as a supportive analytic strategy. This would likely be the case also for the new European Union (EU) Joint Clinical Assessment (JCA) whose evolving methodological guidance acknowledges the value of advanced evidence synthesis approaches in addressing clinical uncertainty [58]. The methodological rigor and possibility to implement BDB with aggregate-level data may further support its acceptability in this context.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s40273-026-01656-7.

Funding This study was funded by F. Hoffmann-La Roche Ltd. Employees of F. Hoffmann-La Roche Ltd contributed to the study design, development of the statistical approach, analysis, and interpretation of the simulated data. The expert panel members received no direct compensation for their authorship of this manuscript; however, they received consultancy fees from F. Hoffmann-La Roche Ltd on a related project. Editorial support, including copywriting, routing, and journal submission assistance, was provided by Nucleus Global, and funded by F. Hoffmann-La Roche Ltd.

Declarations

Conflicts of interest CS and MS are editorial board members of the journal but took no part in the review process or editorial decisions.

Author contributions S Yiu and KG contributed to the study design, development of the statistical approach, analysis, and interpretation of the simulated data. S Yiu, KG, S Yuen and CS wrote the first draft of the manuscript. All authors critically reviewed the manuscript and suggested or made critical edits to the scientific content of the manuscript.

Acknowledgements The authors would like to recognise early input on this work from the cCDP Steering Committee members and Roche contributors. The authors would like to thank Qing Wang, Kristina Weber, Sandro Gsteiger, Aino Launonen, and Claude Berge for helpful discussions that informed the analysis and manuscript. Sean Yiu initiated this work while he was employed by Roche Products Ltd, Welwyn Garden City, UK.

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

References

1. Feltmate K, Janiszewski PM, Gingerich S, et al. Delayed access to treatments for rare diseases: who's to blame. Respirology. 2015;20(3):361–9.
2. Whicher D, Philbin S, Aronson N. An overview of the impact of rare disease characteristics on research methodology. Orphanet J Rare Dis. 2018;13:14. https://doi.org/10.1186/s13023-017-0755-5.
3. Li Y, Izem R. Novel clinical trial design and analytic methods to tackle challenges in therapeutic development in rare diseases. Ann Transl Med. 2022;10(18):1034. https://doi.org/10.21037/atm-21-5496.PMID:36267797;PMCID:PMC9577738.
4. Duivenvoorden R, de Groot E, Stroes ES, Kastelein JJ. Surrogate markers in clinical trials—challenges and opportunities. Atherosclerosis. 2009;206(1):8–16.
5. Reilly MM, Herrmann DN, Pareyson D, Scherer SS, Finkel RS, Züchner S, et al. Trials for slowly progressive neurogenetic diseases need surrogate endpoints. Ann Neurol. 2023;93(5):906–10.
6. Marrie RA, Sormani MP, Apap Mangion S, Bovis F, Cheung WY, Cutter GR, et al. Improving the efficiency of clinical trials in multiple sclerosis. Mult Scler J. 2023;29(9):1136–48.
7. European Medicines Agency. Guideline on clinical investigation of medicinal products for the treatment of multiple sclerosis [Internet]. London: European Medicines Agency; 2015 [cited 2025 Mar 18]. http://www.ema.europa.eu/docs/en_GB/document_library/Scientific_guideline/2015/03/WC500185161.pdf. Accessed on April 18th 2026.
8. FDA draft guidance on Leveraging Existing Clinical Data for Extrapolation to Pediatric Uses of Medical Devices. https://www.fda.gov/media/91889/download. Accessed on April 18th 2026.
9. Goodman SN, Sladky JT. A Bayesian approach to randomized controlled trials in children utilizing information from adults: the case of Guillain–Barre. Clin Trials. 2005;2(4):305–10.
10. Greenhouse JB, Seltman H. Using prior distributions to synthesize historical evidence: comments on the Goodman-Sladky case study of IVIg in Guillain–Barré syndrome. Clin Trials. 2005;2(4):311–8.
11. FDA review of benlysta (belimumab) for intravenous infusion in children 5 to 17 years of age with SLE. https://www.fda.gov/media/127912/download. Accessed on April 18th 2026.
12. Travis J, Rothmann M, Thomson A. Perspectives on informative Bayesian methods in pediatrics. J Biopharm Stat. 2023;33(6):830–43.
13. Cockerham A, Gregory M, Iff J, Smith J, Wolfert D. Implications for rare diseases: NICE guidance and outcomes on implementing bayesian borrowing approaches for external control data. In: Presented at the Professional Society for Health Economics and Outcomes Research (ISPOR) Europe; November 15th, 2023; Copenhagen, Denmark. https://www.ispor.org/docs/default-source/euro2023/cockerhamnice-guidanceispor-eu-2023final-01nov2023upload132305-pdf.pdf?sfvrsn=5dbbb37f_0. Accessed on April 18th 2026.
14. Mackay E, Springford A, Heeg B, Arora P, Thorlund K. Combating sample scarcity: a novel bayesian approach to pediatric basket trials. In: Presented at the Professional Society for Health Economics and Outcomes Research (ISPOR) Europe; November 15th, 2023; Copenhagen, Denmark. https://www.ispor.org/docs/defau

lt-source/euro2023/isporeurope23mackaymsr82poster133896-pdf.pdf?sfvrsn=eee705e5_0. Accessed on April 18th 2026.
15. Potts J, Wang M. Application of a novel Bayesian dynamic borrowing approach using anonymized Finnish real-world data as an external control arm to augment an internal control arm of a randomized clinical trial. In: Presented at the Professional Society for Health Economics and Outcomes Research (ISPOR) Europe; November 15th, 2023; Copenhagen, Denmark. https://www.ispor.org/docs/default-source/intl2024/ispor24pottsmsr29poster138835-pdf.pdf?sfvrsn=46974e48_0. Accessed on April 18th 2026.
16. Struebing A, McKibbon C, Ruan H, Mackay E, Dennis N, Velummailum R, et al. Augmenting external control arms using Bayesian borrowing: a case study in first-line non-small cell lung cancer. J Comp Eff Res. 2024;13(5):e230175.
17. FDA draft guidance on the use of Bayesian methodology in clinical trials of drug and biological products. Draft guidance for industry use of bayesian methodology in clinical trials of drug and biological products. https://www.fda.gov/media/190505/download
18. Saint-Hilary G, Barboux V, Pannaux M, Gasparini M, Robert V, Mastrantonio G. Predictive probability of success using surrogate endpoints. Stat Med. 2019;38:1753–74.
19. Quan H, Xu Z, Luo J, Paux G, Cho M, Chen X. Utilization of treatment effect on a surrogate endpoint for planning a study to evaluate treatment effect on a final endpoint. Pharm Stat. 2023;22(4):633–49.
20. Viele K, Berry S, Neuenschwander B, Amzal B, Chen F, Enas N, et al. Use of historical control data for assessing treatment effects in clinical trials. Pharm Stat. 2014;13:41–54.
21. Kaplan D, Chen J, Yavuz S, Lyu W. Bayesian dynamic borrowing of historical information with applications to the analysis of large-scale assessments. Psychometrika. 2023;88(1):1–30.
22. Kaplan D, Chen J, Lyu W, et al. Bayesian historical borrowing with longitudinal large-scale assessments. Large-scale Assess Educ. 2023. https://doi.org/10.1186/s40536-022-00140-w.
23. Schmidli H, Gsteiger S, Roychoudhury S, O'Hagan A, Spiegelhalter D, Neuenschwander B. Robust meta-analytic-predictive priors in clinical trials with historical control information. Biometrics. 2014;70(4):1023–32.
24. Welton NJ, Phillippo DM, Owen R, Jones HE, Dias S, Bujkiewicz S, Ades AE, Abrams KR. Nice DSU report: CHTE2020 sources and synthesis of evidence; update to evidence synthesis methods. https://www.sheffield.ac.uk/media/34043/download?attachment. Accessed on April 18th 2026.
25. Tan SH, Abrams KR, Bujkiewicz S. Bayesian multiparameter evidence synthesis to inform decision making: a case study in metastatic hormone-refractory prostate cancer. Med Decis Mak. 2018;38(7):834–48.
26. Klugkist I, Volker TB. Bayesian evidence synthesis for informative hypotheses: an introduction. Psychol Methods. 2023;30(5):949–65.
27. Wheaton L, Papanikos A, Thomas A, Bujkiewicz S. Using Bayesian evidence synthesis methods to incorporate real-world evidence in surrogate endpoint evaluation. Med Decis Mak. 2023;43(5):539–52.
28. Best N, Price RG, Pouliquen IJ, Keene ON. Assessing efficacy in important subgroups in confirmatory trials: an example using Bayesian dynamic borrowing. Pharm Stat. 2021;20:551–62.
29. Spiegelhalter DJ. Incorporating Bayesian ideas into health-care evaluation. Stat Sci. 2004;19(1):156–74.
30. Röver C, Wandel S, Friede T. Model averaging for robust extrapolation in evidence synthesis. Stat Med. 2019;2019(38):674–94.
31. Hauser SL, Bar-Or A, Comi G, Giovannoni G, Hartung HP, Hemmer B, et al. Ocrelizumab versus interferon beta-1a in relapsing multiple sclerosis. N Engl J Med. 2017;376(3):221–34.
32. Gärtner J, Hauser SL, Bar-Or A, Montalban X, Cohen JA, Cross AH, et al. Efficacy and safety of ofatumumab in recently diagnosed, treatment-naive patients with multiple sclerosis: results from ASCLEPIOS I and II. Mult Scler. 2022;28(10):1562–75.
33. Biogen. Effect of natalizumab on disease progression in secondary progressive multiple sclerosis (ASCEND). ClinicalTrials.gov identifier: NCT01416181. Published August 10, 2011. Updated March 28, 2019. https://clinicaltrials.gov/study/NCT01416181. Accessed on April 18th 2026.
34. Novartis Pharmaceuticals. A study of fingolimod in primary progressive multiple sclerosis (INFORMS). ClinicalTrials.gov identifier: NCT00731692. Published August 7, 2008. Updated April 12, 2016. https://clinicaltrials.gov/study/NCT00731692. Accessed on April 18th 2026.
35. Sanofi. Efficacy and safety study of frexalimab (SAR441344) in adults with nonrelapsing secondary progressive multiple sclerosis. ClinicalTrials.gov identifier: NCT06141486. Published December 1, 2023. Updated January 30, 2024. https://clinicaltrials.gov/study/NCT06141486. Accessed on April 18th 2026.
36. Sanofi. A study to evaluate the efficacy and safety of tolebrutinib in adult participants with primary progressive multiple sclerosis (PERSEUS). ClinicalTrials.gov identifier: NCT04458051. Published July 8, 2020. Updated January 18, 2024. https://clinicaltrials.gov/study/NCT04458051. Accessed on April 18th 2026.
37. Hoffmann-La Roche. A study to evaluate the efficacy, safety and pharmacokinetics (pk) of a higher dose of ocrelizumab in adults with relapsing multiple sclerosis (RMS) (MUSETTE). ClinicalTrials.gov identifier: NCT04544436. Published September 9, 2020. Updated March 6, 2024. https://clinicaltrials.gov/study/NCT04544436. Accessed on April 18th 2026.
38. Hoffmann-La Roche. A study to evaluate the safety and efficacy of ocrelizumab in adults with primary progressive multiple sclerosis (GAVOTTE). ClinicalTrials.gov identifier: NCT04548999. Published September 10, 2020. Updated March 6, 2024. https://clinicaltrials.gov/study/NCT04548999. Accessed on April 18th 2026.
39. Hoffmann-La Roche. A study to evaluate the efficacy and safety of fenebrutinib compared with teriflunomide in participants with relapsing multiple sclerosis. ClinicalTrials.gov identifier: NCT04458051. Published July 8, 2020. Updated January 18, 2024. https://clinicaltrials.gov/study/NCT04458051. Accessed on April 18th 2026.
40. Cadavid D, Cohen JA, Freedman MS, Goldman MD, Hartung HP, Havrdova E, et al. The EDSS-Plus, an improved endpoint for disability progression in secondary progressive multiple sclerosis. Mult Scler. 2017;23(1):94–105.
41. Kappos L, Yiu S, Dahlke F, Coetzee T, Cutter GR, Yuen S, Bonati U, Lublin FD. Composite confirmed disabilityworsening/progression is a useful clinical endpoint for multiple sclerosis clinical trials. Neurology. 2025;104(10):e213558. https://doi.org/10.1212/WNL.0000000000213558. Epub 2025 Apr 21. PMID: 40258203; PMCID: PMC12136489.
42. Sormani MP, Bonzano L, Roccatagliata L, Mancardi GL, Uccelli A, Bruzzi P. Surrogate endpoints for EDSS worsening in multiple sclerosis. A meta-analytic approach. Neurology. 2010;75(4):302–9.
43. Polman CH, O'Connor PW, Havrdova E, et al. A randomized, placebo-controlled trial of natalizumab for relapsing multiple sclerosis. N Engl J Med. 2006;354(9):899–910.
44. Kappos L, Bar-Or A, Cree B, et al. Siponimod versus placebo in secondary progressive multiple sclerosis (EXPAND): a double-blind, randomised, phase 3 study. Lancet. 2018;391:1263–73.

45. Montalban X, Hauser SL, Kappos L, et al. Ocrelizumab versus placebo in primary progressive multiple sclerosis. N Engl J Med. 2017;376(3):209–20.
46. Wolinsky JS, Narayana PA, O'Connor P, Coyle PK, Ford C, Johnson K, et al. Glatiramer acetate in primary progressive multiple sclerosis: results of a multinational, multicenter, double-blind, placebo-controlled trial. Ann Neurol. 2007;61:14–24. https://doi.org/10.1002/ana.21079.
47. Hawker K, O'Connor P, Freedman MS, Calabresi PA, Antel J, Simon J, et al. Rituximab in patients with primary progressive multiple sclerosis: results of a randomized double-blind placebo-controlled multicenter trial‡. Ann Neurol. 2009;66:460–71. https://doi.org/10.1002/ana.21867.
48. Schoenfeld D. The asymptotic properties of nonparametric tests for comparing survival distributions. Biometrika. 1981;68:316–9.
49. Riley RD, Lambert PC, Staessen JA, Wang J, Gueyffier F, Thijs L, et al. Meta-analysis of continuous outcomes combining individual patient data and aggregate data. Stat Med. 2008;27(11):1870–93.
50. Brooks S, Gelman A, Jones G, Meng XL (Eds.). Handbook of Markov ChainMonte Carlo (1st ed.). Chapman and Hall/CRC. 2011 https://doi.org/10.1201/b10905
51. Hussein H, Nevill CR, Meffen A, Abrams KR, Bujkiewicz S, Sutton AJ, et al. Double-counting of populations in evidence synthesis in public health: a call for awareness and future methodological development. BMC Public Health. 2022;22(1):1827.
52. Cooper NJ, Sutton AJ, Abrams KR, Turner D, Wailoo A. Comprehensive decision analytical modelling in economic evaluation: a Bayesian approach. Health Econ. 2004;13(3):203–26.
53. Mackay EK, Springford A. Evaluating treatments in rare indications warrants a Bayesian approach. Front Pharmacol. 2023;14:1249611.
54. Bullement A, Stevenson MD, Baio G, Shields GE, Latimer NR. A systematic review of methods to incorporate external evidence into trial-based survival extrapolations for health technology assessment. Med Decis Mak. 2023;43(5):610–20.
55. European Medicines Agency. Reflection paper on statistical methodology for the comparative assessment of quality attributes in drug development. EMA/CHMP/138502/2017. London: European Medicines Agency; 2017. https://www.ema.europa.eu/en/documents/scientific-guideline/reflection-paper-statistical-methodology-comparative-assessment-quality-attributes-drug-development_en.pd. Accessed on April 18th 2026.
56. Haute Autorité de Santé. Choices in methods for economic evaluation. Saint-Denis La Plaine: HAS; 2020. https://www.has-sante.fr/jcms/p_3216041/en/methodological-guidance-2020-choices-in-methods-for-economic-evaluationH. Accessed on April 18th 2026.
57. Institut für Qualität und Wirtschaftlichkeit im Gesundheitswesen (IQWiG). General Methods: Version 7.0. Cologne: IQWiG; 2023. https://www.iqwig.de/methoden/general-methods_version-7-0.pdf. Accessed on April 18th 2026.
58. European Commission. Practical guideline for quantitative evidence synthesis: direct and indirect comparisons. Brussels: European Commission; 2023. https://health.ec.europa.eu/system/files/2023-03/hta_practical-guideline_direct-and-indirect-comparisons_en_0.pdf. Accessed on April 18th 2026.