




MAIN PAPER

Improving the assessment of the probability of success in late stage drug development

Lisa V. Hampson¹  | Björn Bornkamp¹  | Björn Holzhauer¹  |
Joseph Kahn² | Markus R. Lange¹ | Wen-Lin Luo² | Giovanni Della Cioppa³ |
Kelvin Stott⁴ | Steffen Ballerstedt¹

¹Analytics, Novartis Pharma AG,
Basel, Switzerland

²Analytics, Novartis Pharmaceuticals
Corporation, East Hanover,
New Jersey, USA

³Clinical R&D Consultants srl, Rome,
Italy

⁴Portfolio Analytics, Novartis Pharma AG,
Basel, Switzerland

Correspondence

Lisa V. Hampson, Analytics, Novartis
Pharma AG, Postfach 4002,
Basel, Switzerland.
Email: lisa.hampson@novartis.com

Abstract

There are several steps to confirming the safety and efficacy of a new medicine. A sequence of trials, each with its own objectives, is usually required. Quantitative risk metrics can be useful for informing decisions about whether a medicine should transition from one stage of development to the next. To obtain an estimate of the probability of regulatory approval, pharmaceutical companies may start with industry-wide success rates and then apply to these subjective adjustments to reflect program-specific information. However, this approach lacks transparency and fails to make full use of data from previous clinical trials. We describe a quantitative Bayesian approach for calculating the probability of success (PoS) at the end of phase II which incorporates internal clinical data from one or more phase IIb studies, industry-wide success rates, and expert opinion or external data if needed. Using an example, we illustrate how PoS can be calculated accounting for differences between the phase II data and future phase III trials, and discuss how the methods can be extended to accommodate accelerated drug development pathways.

KEYWORDS

Bayesian methods, expert elicitation, meta-analysis, quantitative decision making

1 | INTRODUCTION

Before a new medicine can be licensed, a sequence of trials, each with its own objectives, is required to confirm the medicine's efficacy and safety. Clinical research generally begins with small-scale phase I studies which focus on the tolerability and safety of the medicine. Phase II is then often subdivided into two distinct stages: phase IIa, intended to demonstrate proof-of-concept, and phase IIb, which is used to identify the dose and dosing schedule to be taken forward to phase III. The final stage of development involves performing confirmatory phase III trials. A program like this will be punctuated by a series of milestones at which the sponsor reviews all of the available evidence and decides whether or not to transition to the next stage of development. A key milestone occurs at the end phase IIb, since continuation requires investment in large-scale pivotal studies. Metrics such as a program's probability of success (PoS) are routinely used to quantify and communicate risk.

While PoS has been defined in a variety of ways in the literature,¹ here we define PoS as a marginal probability of success, which can be thought of as the expectation of the conditional probability of success calculated by averaging

over the prior distribution of unknown parameters such as key treatment effects.^{2,3} PoS will evolve as new information becomes available, such as interim results from an on-going pivotal study.⁴ Success can mean different things to different stakeholders. For example, at the trial-level, success is often taken to mean achieving statistical significance on the primary endpoint, in which case PoS simplifies to the Bayesian predictive power of the trial, also referred to as “assurance”.² In this article, we focus on calculating the probability of program success before initiating registration trials, which in general will occur at the end of phase II. We define program success as regulatory approval with effects on key endpoints sufficient to secure market access, that is, get a newly approved drug used and reimbursed. PoS can be used to inform trial design discussions, such as whether to include a futility interim analysis.⁵ It is also essential for calculating a program’s expected net present value (eNPV), defined approximately as $eNPV = PoS \times NPV(\text{Rewards}) - NPV(\text{Costs})$. The eNPV metric is important for informing investment decisions and has also been used as an objective function to optimize various aspects of program design including sample size and dose-selection strategies.^{6,7,8}

Several methods of increasing complexity have been proposed for calculating PoS. The first, simplest, approach summarizes industry data by aggregate success rates: so-called industry “benchmarks” are available for the probability of regulatory approval from a particular milestone, or the probability of successfully transitioning out of a development phase.^{9,10} Typically benchmarks are disaggregated across a limited set of covariates, such as therapeutic area or lifecycle class. The project team can then apply adjustments to the benchmark based on a subjective assessment of program-specific risks to arrive at a final, more tailored, PoS. While this approach is quick and simple, there are several drawbacks. Most importantly, subjective adjustments are open to heuristic biases^{11,12} and are likely to be applied inconsistently across programs. Recently, more advanced multivariable modeling and machine learning techniques have been applied to industry datasets to generate tailored industry benchmarks, which are estimates of the probability of approval adjusting for several (and sometimes tens) of program characteristics.^{13,14} However, when using this approach, the analyst must be careful to avoid leaking information by conditioning on prophetic variables which would not typically be known at the time an investment decision is made.

An alternative approach for calculating PoS is to present to a group of experts the relevant evidence and elicit from them a prior distribution for the treatment effect which is then used to drive the PoS evaluation.^{15,16} Evidence may comprise data from early phase trials, studies of the drug in related indications or trials of similar drugs, as well as more broadly relevant information such as the scientific rationale for the drug’s mechanism of action.

A third data-based approach to the PoS assessment is to take the phase II data at face value without adjusting for any selection that may have occurred at the end of phase II, and use a Bayesian approach to combine these with prior information. The resulting posterior for the treatment effect is then used to drive the PoS evaluation. The choice of prior will have an important impact on the PoS.^{17,3} One major limitation of this approach is that it can be applied only when there are no design differences between phases.

In this article, we extend existing methodology to propose a new quantitative approach for evaluating the PoS of a drug development program at the end of phase II. Referring to adverse events that would prompt abandonment of a development program despite positive efficacy data as safety showstopper events (SSEs), at the end of phase II there are three hurdles for success, as shown in Figure 1A. Specifically, we must: (a) meet statistical significance on the one or two efficacy endpoints needed for approval in all phase III trials without observing a SSE; (b) obtain regulatory approval; and (c) for all efficacy endpoints considered essential for market access (which is typically a larger set than the group of endpoints needed for approval), observe treatment effect estimates which are in excess of the minimum thresholds believed to be sufficient to secure access. We refer to these thresholds collectively as the “target product profile” (TPP) because this is the document which is used in many pharmaceutical companies to outline the desired profile of a new therapeutic, including efficacy and safety. We begin focusing on traditional development programs with data available from at least one phase IIb trial. Then, we can calculate the probability of taking all three of the hurdles in Figure 1(A) by combining several sources of information including the tailored industry benchmark, the phase IIb efficacy data, the design of the phase III studies and a qualitative assessment of any remaining risks. Frequently differences between phases will preclude a purely data-based PoS evaluation. In these cases, we propose leveraging expert opinion in order to relate the phase IIb data to the quantities of interest in phase III.

The remainder of this article proceeds as follows. In Section 2, we begin by giving an overview of the PoS calculation. Section 3 describes how we use program-specific efficacy data and industry benchmarks for reasons of attrition to estimate the probability of running a positive phase III program. In Section 4, we take a step back and discuss how prior distributions for parameters of the Bayesian meta-analytic model used to combine early phase efficacy data can be informed by tailored industry benchmarks. In Section 5, we propose a semi-quantitative approach to account for any remaining risks which are not accounted for in previous steps of the calculation, while in Section 6 we discuss how to

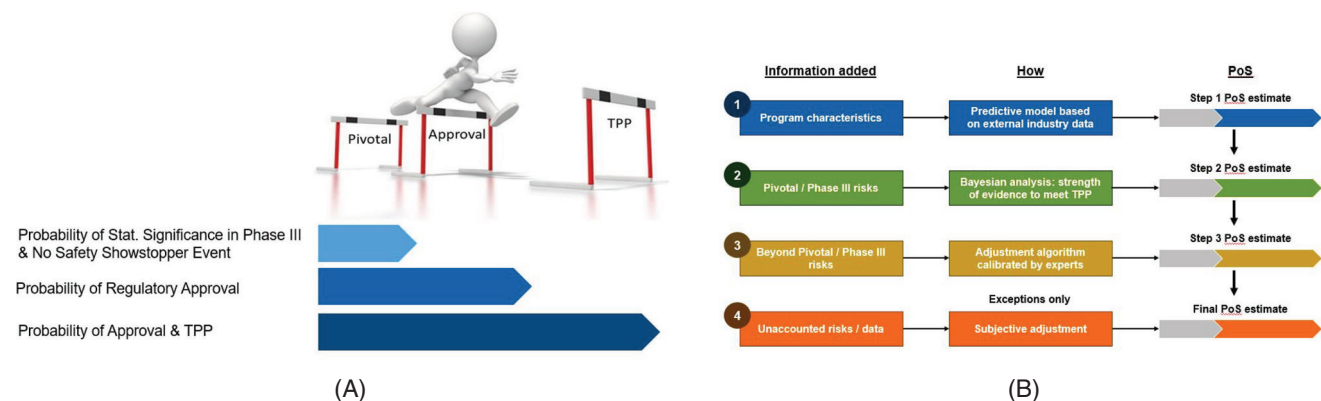


FIGURE 1 (A) Three hurdles for success at the end of phase II; (B) The proposed framework integrates all key evidence in three main steps. Step 4 would only be applied when key risks or data have been omitted from the first three steps

bridge across differences between phase IIb and phase III trials using expert opinion. We illustrate the proposed framework with an example in Section 7 and conclude by outlining further work in Section 8.

2 | ROAD MAP OF THE PROBABILITY OF SUCCESS CALCULATION

This section provides an overview of the PoS calculation; a schematic diagram can be found in Figure 1B. The evaluation begins by considering the risks associated with the phase III studies. The probability of a positive phase III program in which we demonstrate efficacy on key endpoints with no SSE is:

$$\mathbb{P}\{\text{Efficacy success in phase III on 1-2 key endpoints}\} \times \mathbb{P}\{\text{No SSE in phase III} | \text{Efficacy success in phase III on 1-2 key endpoints}\}. \tag{1}$$

Note that the probabilities in (1) do not condition on the unknown treatment effects. Efficacy success in phase III has two components. First, we must achieve statistical significance on the key endpoints in all phase III trials. Second, we must observe average treatment effect estimates for these endpoints which are at least in line with the TPP. While the first component is a minimum requirement for regulatory approval, the second is a prerequisite for market access. We begin by calculating the probability of efficacy success assuming data are available from $J \geq 1$ phase IIb studies which measured the primary endpoint(s) of the phase III trials in the target patient population and compared the selected dose and formulation of the novel drug to the phase III control. We describe in Section 3.1 the Bayesian meta-analytic model used to combine the phase IIb data allowing for between-study differences and the simulation approach used to calculate the probability of efficacy success in phase III. To limit complexity we restrict attention in this step to the one or two endpoints considered essential for approval. These will typically be the primary and key secondary endpoints of phase III.

Incorporating safety endpoints into the Bayesian meta-analytic model is challenging. This is because phase II trials are rarely of sufficient size and duration to identify or characterize those adverse events that would imperil phase III success. Furthermore, the correlation of efficacy and safety effects is likely to be poorly understood at the end of phase IIb. With this in mind, we describe in Section 3.2 a simpler, albeit coarser, approach for calculating the right-hand side term of (1) which is based on industry benchmarks rather than project-specific data.

Success on the key phase III efficacy endpoints is typically necessary but not sufficient for market access, and the TPP will usually include additional endpoints which are either measured as secondary endpoints in phase III or studied in dedicated phase IIIb trials. The probability of success, defined as regulatory approval with all endpoints needed for access meeting the TPP, is:

$$\mathbb{P}\{\text{Approval \& TPP} | \text{Efficacy success on 1-2 key endpoints \& no SSE in phase III}\} \times \mathbb{P}\{\text{Efficacy success on 1-2 key endpoints \& no SSE in phase III}\} \tag{2}$$

Section 5 describes the semi-quantitative approach taken to calculate the left hand side term of (2), referred to as the conditional PoS.

3 | LEVERAGING CLINICAL AND EXTERNAL DATA TO ASSESS THE CHANCE OF SUCCESS IN PHASE III

3.1 | Bayesian meta-analytic approach to calculating the probability of efficacy success in phase III

Suppose two endpoints will be key to efficacy success in phase III; the case for a single endpoint follows naturally. We refer to them as the primary endpoint P and secondary endpoint S, but the same approach can be applied if they are, for example, co-primary endpoints. We define $\theta_{2j} = (\theta_{P2j}, \theta_{S2j})$ as the study-specific treatment effects underpinning the j th phase IIb trial, for $j = 1, \dots, J$. Without loss of generality, we assume that the null effect consistent with no advantage versus control is 0 for each endpoint, and larger effects are consistent with greater efficacy. The TPP thresholds for endpoints P and S are δ_P and δ_S .

In what follows we will provide a high-level description on how to obtain a Markov Chain Monte Carlo (MCMC) sample from the predictive distribution for the phase III treatment effects. For details please refer to Appendix A.

In a first step a Bayesian normal-normal hierarchical model is used to synthesize information across the J phase IIb trials. For that purpose it is assumed that pairs of estimates originating from phase IIb at least approximately after suitable transformation follow a bivariate normal distribution (see Equation A1 in Appendix A). The correlation between pairs of estimates as well as further nuisance parameters will be assumed known and set to their estimates from phase II trials. We combine these data assuming study-specific treatment effects are exchangeable and samples from a bivariate normal random-effects distribution (see A2 in Appendix A). The prior for the mean of this distribution $\mu = (\mu_P, \mu_S)$ will be described in detail in Section 4. For the across study SDs τ_{P2} and τ_{S2} , we follow others^{18,19,20} to stipulate weakly informative half-normal priors. Adopting the nomenclature of Neuenschwander and Schmidli,²¹ the examples presented in this article will assume “small” between-trial heterogeneity in phase IIb for all key endpoints. For simplicity, the correlation ρ between the treatment effects in the bivariate normal random-effects distribution is treated as a fixed constant supplied by the analyst. Its specification could be based on a meta-regression of pairs of treatment effect estimates obtained from trials of drugs with a similar mechanism of action to the novel drug.

Based on this model one can create L MCMC samples from the posterior distribution for μ as $(\mu_P^{(1)}, \mu_S^{(1)})$, ..., $(\mu_P^{(L)}, \mu_S^{(L)})$. This sample can then be used to sample from the meta-analytic-predictive (MAP) prior for the study-specific treatment effects in the K planned phase III trials, denoted by $\theta_{3k} = (\theta_{P3k}, \theta_{S3k})$, for $k = 1, \dots, K$. If there are no differences between the target estimands of the phase IIb and phase III studies, the long-run averages of the study-specific effects in each phase should be identical. However, the degree of between-study heterogeneity is expected to be smaller in phase III than phase IIb because it is common for pivotal studies to run concurrently with one another and follow similar (if not identical) protocols. Let τ_{P3} and τ_{S3} denote the SDs of the phase III study-specific treatment effects on endpoints P and S, where again half-normal priors are used but now with medians corresponding to “very small” heterogeneity. Taking L independent samples from the prior distributions of τ_{P3} and τ_{S3} , and assuming that study-specific treatment effects in phases IIb and III are partially exchangeable, we can then sample,

$$\theta_{31}^{(\ell)}, \dots, \theta_{3K}^{(\ell)} | \mu^{(\ell)}, \rho, \tau_{P3}^{(\ell)}, \tau_{S3}^{(\ell)} \sim N \left(\begin{bmatrix} \mu_P^{(\ell)} \\ \mu_S^{(\ell)} \end{bmatrix}, \begin{bmatrix} \tau_{P3}^{(\ell)2} & \rho \tau_{P3}^{(\ell)} \tau_{S3}^{(\ell)} \\ \rho \tau_{P3}^{(\ell)} \tau_{S3}^{(\ell)} & \tau_{S3}^{(\ell)2} \end{bmatrix} \right) \text{ for } \ell = 1, \dots, L. \quad (3)$$

For each $\ell = 1, \dots, L$, given the study-specific treatment effects $\theta_{31}^{(\ell)}, \dots, \theta_{3K}^{(\ell)}$, we can simulate the outcome of the ℓ th phase III program assuming that treatment effect estimators follow the canonical joint distribution (A1) and setting information levels equal to their design values. Statistical significance in a trial is declared if the simulated treatment effect estimate exceeds the critical value of the planned hypothesis test. The totality of evidence from the phase III program is used to determine whether a TPP target has been met, and for simplicity we combine estimates via a fixed-effects meta-

analysis. Therefore, the TPP threshold for an endpoint in program ℓ is deemed to have been met if it is less than the weighted average of $\hat{\theta}_{31}^{\ell}, \dots, \hat{\theta}_{3K}^{\ell}$, weighting by the inverse variances.^{22,23} The predictive probability of efficacy success in phase III is given by

$$\frac{1}{L} \sum_{\ell=1}^L \mathbf{1}\{\text{Meet efficacy success criteria in } \ell\text{th phase III program} | \theta_{31}^{(\ell)}, \dots, \theta_{3K}^{(\ell)}\}.$$

We need to proceed slightly differently when a key efficacy endpoint is binary and the treatment effect is a difference in proportions. This is because the assumption of normality in parameter model (A2) could lead us to place probability mass on effects outside the interval $[-1, 1]$, particularly if response rates close to 0 or 1 are expected on either treatment arm. Appendix B describes how we handle this special case.

3.2 | Calculating the probability of no safety showstopper event in phase III

The probability of a positive phase III program in (1) depends on the conditional probability of no SSE in phase III given efficacy is demonstrated. We use industry benchmarks, rather than project-specific clinical data, to quantify the risk of a SSE. Several authors have reviewed the reasons for attrition in drug development and how these vary across phases.^{24,25,26} However, since only the primary reason for termination is typically reported in industry datasets, we cannot estimate the joint distribution of different causes for failure. In addition, failure attribution may not always be explicit: “strategic reasons” are commonly cited for termination²⁶ but we speculate this may be a coded version of poor efficacy or safety.

We simplify to assume that a program can only fail due to either inadequate efficacy or safety. Furthermore, we assume these two causes for failure are independent. Under the latter assumption, the conditional probability of no SSE in phase III given efficacy success in (1) simplifies to the unconditional probability of no SSE in phase III. This approximation is likely to be conservative because higher rates of serious adverse events on the novel drug would be expected to result in higher rates of study discontinuation or treatment switching which, depending on how these intercurrent events are handled, may dilute efficacy. Therefore, if we were told efficacy has been demonstrated, the risk of a SSE would decrease. We can express the unconditional probability of no SSE in phase III as:

$$\begin{aligned} \mathbb{P}\{\text{No SSE in phase III}\} &= 1 - \mathbb{P}\{\text{Fail in phase III}\} \times \mathbb{P}\{\text{SSE in phase III} | \text{Fail in phase III}\} \\ &= 1 - (1 - \mathbb{P}\{\text{Succeed in phase III}\}) \times \mathbb{P}\{\text{SSE in phase III} | \text{Fail in phase III}\}. \end{aligned} \quad (4)$$

We estimate $\mathbb{P}\{\text{Succeed in phase III}\}$ in (4) using a tailored industry benchmark which is obtained by evaluating a simple predictive model which was fitted to an industry dataset according to the approach described in Appendix C. The conditional risk $\mathbb{P}\{\text{SSE in phase } i | \text{Fail in phase } i\}$ in (4) is also estimated using an industry benchmark. To derive this, we used the Clarivate Global R&D performance metric program “CMR” (Centre for Medicines Research) database which provides aggregate summaries of transition rates and reasons for failure by phase. Assuming all failures that were not attributed to safety were due to lack of efficacy, restricting attention to programs entering a phase between 2012 and 2018 and excluding vaccines and biosimilars, we obtained estimates:

1. Non-oncology: $\mathbb{P}\{\text{SSE} | \text{Fail in phase II}\} = 0.101$, $\mathbb{P}\{\text{SSE} | \text{Fail in phase III}\} = 0.15$
2. Oncology: $\mathbb{P}\{\text{SSE} | \text{Fail in phase II}\} = 0.093$, $\mathbb{P}\{\text{SSE} | \text{Fail in phase III}\} = 0.01$.

Actually we found no failures in phase III for oncology programs were directly attributed to safety; we assign this event a small probability of 0.01 so as not to rule it out entirely. We stratify by disease area since different risk–benefit trade-offs will be acceptable in oncology and non-oncology programs. For simplicity, we assume the same mix of attrition reasons for programs in phase IIa and phase IIb, so that for both therapeutic areas $\mathbb{P}\{\text{SSE} | \text{Fail in phase IIa}\} = \mathbb{P}\{\text{SSE} | \text{Fail in phase IIb}\}$ which we set equal to our estimate of $\mathbb{P}\{\text{SSE} | \text{Fail in phase II}\}$.

4 | PRIOR DISTRIBUTION FOR THE AVERAGE TREATMENT EFFECT μ

4.1 | Motivation

We have yet to comment on what prior we will place on μ , the mean of the random-effects distribution linking study-specific treatment effects. Using an “off the shelf” weakly informative normal distribution would neglect the information we have from the industry benchmark, as well as the potential impact of selection bias on the phase II effect estimate(s). Figure 2A compares the probability density function (pdf) of the maximum likelihood estimate (MLE) of the treatment effect from a phase III trial with the conditional pdf of the MLE from phase II given statistical significance is achieved. Figure 2B shows how the magnitude of the selection bias in the phase II MLE from a statistically significant trial varies with the treatment effect and phase II sample size. The impact of the selection bias is highest when the phase II trial is poorly powered for its primary objective²⁷ and when the drug has little benefit versus control. We could try to account for this bias by explicitly modeling the phase III go/no-go criteria²⁸ although in practice this is not straightforward as investment decisions are influenced by multiple factors. Alternatively, based on a small review of the Pfizer portfolio, Kirby et al.²⁹ propose discounting the phase II estimate by 10%. However, applying a fixed discount factor ignores the influence of phase IIb sample size and drug efficacy effect on the selection bias. With these factors in mind, we try to ameliorate the impact of selection bias by using a prior for μ satisfying the following requirements:

1. It should incorporate some degree of skepticism.
 2. The degree of skepticism should reflect the historical success rates of similar projects at the same stage of development. Since μ measures the efficacy of the new drug, only the benchmark probability of efficacy success in phases II and III (given we start phase II) is relevant for informing the prior. The benchmark conditional probability of approval given we submit a new drug application (NDA) is not considered informative for μ since approval outcomes may be influenced by many factors beyond efficacy; see Section 5.
 3. The influence of the benchmark should decrease as the phase II sample size and/or as the phase IIb effect estimate increases.
- After motivating our choice of prior for μ , Section 4.2 provides more details on its specification.

4.2 | Defining the mixture prior for the average treatment effect

We specify a mixture prior for μ placing probability ω on a “null” component consistent with the hypothesis that the new treatment offers no clinically relevant advantage over control on either endpoint, and probability $(1 - \omega)$ on a “TPP” component consistent with the hypothesis that drug effects on both endpoints are close to the TPP thresholds. To capture these beliefs, we set:

$$f(\mu_P, \mu_S) = \omega f_1(\mu_P, \mu_S) + (1 - \omega) f_2(\mu_P, \mu_S). \quad (5)$$

For $c = 1, 2$, we define $f_c(\mu_P, \mu_S)$ as the pdf of a bivariate normal random variable with mean η_c and variance matrix Σ_c , where:

$$\eta_1 = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad \Sigma_1 = \begin{pmatrix} \sigma_{P1}^2 & \rho \sigma_{P1} \sigma_{S1} \\ \rho \sigma_{P1} \sigma_{S1} & \sigma_{S1}^2 \end{pmatrix}$$

and

$$\eta_2 = \begin{pmatrix} \delta_P \\ \delta_S \end{pmatrix} \quad \Sigma_2 = \begin{pmatrix} \sigma_{P2}^2 & \rho \sigma_{P2} \sigma_{S2} \\ \rho \sigma_{P2} \sigma_{S2} & \sigma_{S2}^2 \end{pmatrix}.$$

We assume that the correlation between μ_P and μ_S is the same as the linear correlation between study-specific effects in the random-effects distribution (A2). We find σ_{P1} as the solution to $\mathbb{P}\{\mu_P \geq \delta_P; \eta_1, \Sigma_1\} = 0.01$, and σ_{S1} is defined similarly: placing 1% probability in the upper tail is consistent with the interpretation of $f_1(\mu_P, \mu_S)$ as the “null” component of the

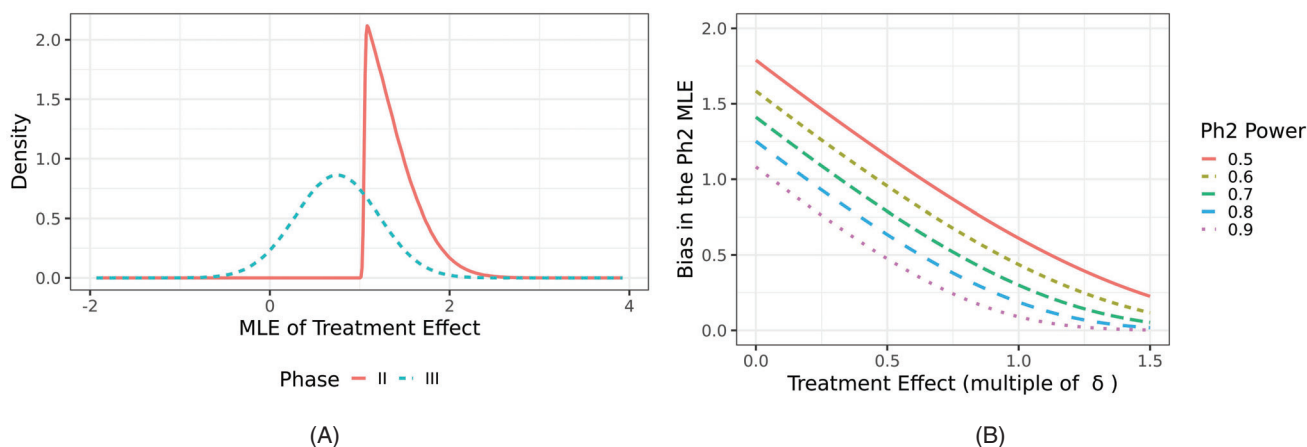


FIGURE 2 Results are for the case that the primary endpoint P is normally distributed with a known SD of 2, where the difference in average responses on the new drug versus control is identical across phases, that is, $\theta_{P21} = \theta_{P31} = \theta_P$. The phase IIb and phase III studies are designed to test $H_0 : \theta_P \leq 0$ versus $H_1 : \theta_P > 0$ with type I error rate $\alpha = 0.025$ at $\theta_P = 0$, and power is specified at $\theta_P = \delta_P$ setting $\delta_P = 1.5$. The phase III trial is designed to have power 0.9 at $\theta_P = \delta_P$. The figures plot: (A) Comparison of the unconditional pdf of $\hat{\theta}_{P31}$ with the conditional pdf of $\hat{\theta}_{P21}$ given we achieve statistical significance in phase IIb, when the phase IIb trial has power 0.8 at $\theta_P = \delta_P$ and in truth $\theta_P = 0.5\delta_P$. (B) Conditional bias in $\hat{\theta}_{P21}$ given statistical significance is achieved in phase IIb at level- α




FIGURE 3 Mixture prior for μ_P when $\delta_P = 10$ and $\omega = 0.5$. In this case, $\sigma_{P1} = -\delta_P/\Phi^{-1}(0.01)$ and $\sigma_{P2} = -\delta_P/\Phi^{-1}(0.01)$. Results assume there is a single endpoint P of interest

mixture. Meanwhile, we find σ_{P2} as the solution to $\mathbb{P}\{\mu_P \leq 0; \eta_2, \Sigma_2\} = 0.01$, consistent with the interpretation of $f_2(\mu_P, \mu_S)$ as the “TPP” component; σ_{S2} is defined similarly. In the examples we have considered, choosing the prior SD s in this way implies that as ω tends towards 0.5, marginal priors $f(\mu_P)$ and $f(\mu_S)$ roughly approximate uniform densities on the intervals $(0, \delta_P)$ and $(0, \delta_S)$, capturing our prior equipoise about whether or not the drug has a meaningful benefit. Figure 3 shows one such example with a single primary efficacy endpoint.

We calibrate prior (5) to the industry benchmark by searching for the value of ω such that the unconditional probability of efficacy success in a “standard” phase IIb and phase III program equals the corresponding tailored industry benchmark, the derivation of which is given in Appendix C. We characterize a “standard” development program as comprising one phase IIb study and either one or two phase III studies, depending on the disease area (one if oncology; two, otherwise). The unconditional probability of efficacy success is,

$$\int \mathbb{P}\{\text{Efficacy success in ‘standard’ phase IIb and III} | \mu_P, \mu_S\} f(\mu_P, \mu_S) d\mu.$$

For the purposes of prior calibration, we define efficacy success as observing a (one-sided) p -value < 0.05 in phase IIb for the endpoint associated with the smallest Fisher information for any given sample size; and observing p -values

< 0.025 on both endpoints P and S in all phase III studies. This is because based on our experience, success on one endpoint is typically considered sufficient in phase II. We assume a standard phase IIb (phase III) study is designed to have power 0.8 (0.9) to meet its objectives when treatment effects equal their TPP thresholds. When there is one key efficacy endpoint, a closed form expression for ω exists which is presented in Appendix D.

4.3 | Accommodating accelerated development programs which skip phases

So far we have restricted attention to traditional development pathways, where a phase III program is preceded by one or more phase IIb trials. However, accelerated pathways which skip phases are common in highly competitive research spaces, in conditions where there is high unmet medical need, or where there is an abundance of existing relevant evidence. For accelerated development programs, phase labels can also become somewhat arbitrary. For example, pivotal studies may be labeled phase II rather than phase III, although they are still intended to support registration. It is also common in the oncology space for phase Ib expansion-cohort studies to collect efficacy data in the target patient population and these trials play a similar role to that of phase IIa studies in other disease areas. If efficacy data are available from early phase Ib or phase IIa studies, it is straightforward to extend our approach to evaluate the PoS of these abbreviated programs, the principal methodological question being which industry benchmark should we use to calibrate the prior for μ ? If early-phase efficacy data come from a phase Ib or phase IIa study, we calibrate $f(\mu_P, \mu_S)$ in (5) so that the unconditional probability of efficacy success in a standard phase IIa, IIb and III program equals the corresponding industry benchmark. We assume a standard phase IIa program consists of a single study with (one-sided) type I error rate 0.1 and power 0.8; standard phase IIb and phase III programs remain as above.

Note that our focus is on calculating the PoS of pivotal trial(s) intended from the outset to support registration. Programs granted conditional approval based on overwhelmingly positive results from an early phase trial will typically fall outside the scope of the current work unless these studies were pre-specified as registrational.

4.4 | Illustrative example

We retrospectively evaluated the PoS of a Novartis program which at the time of the calculation had started phase III but not yet reported the results of the pivotal trial. To protect confidentiality, some details have been anonymized.

We began by recording the important program characteristics given in Table A2 (in Appendix C) associated with a project's probability of success in phases II and III. The drug (T) was a small molecule, orally administered, targeting an enzyme to treat a condition in the cardiovascular/metabolic/renal therapeutic area and was already approved for other indications. Drug T had not been granted a breakthrough designation for the indication in question. Prior to phase III, T had been studied in a single phase II trial, which we label as phase IIa and index as study $j = 1$. Table 1 lists the industry benchmarks given these characteristics obtained from the predictive models described in Appendix C. The primary endpoint P of the phase IIa trial was change from baseline at week 12 in a \log_2 -transformed continuous biomarker, which is normally distributed with SD 0.91. The objective was to demonstrate superiority of T versus control; larger reductions in the biomarker by week 12 reflect an advantage of T, and the TPP threshold was $\log_2(0.75) = -0.42$, interpreted as a 25% relative reduction in the geometric mean biomarker ratio to baseline at week 12.

TABLE 1 Tailored benchmark probabilities of overall, efficacy, and safety success by phase for our example. Success in the submission phase means obtaining regulatory approval

Phase	Prob. of success In phase	Prob. of efficacy Success in phase	Prob. of no SSE In phase
IIa	0.68	0.72	0.97
IIb	0.68	0.72	0.97
III	0.70	0.76	0.96
Submission	0.88	NA	NA

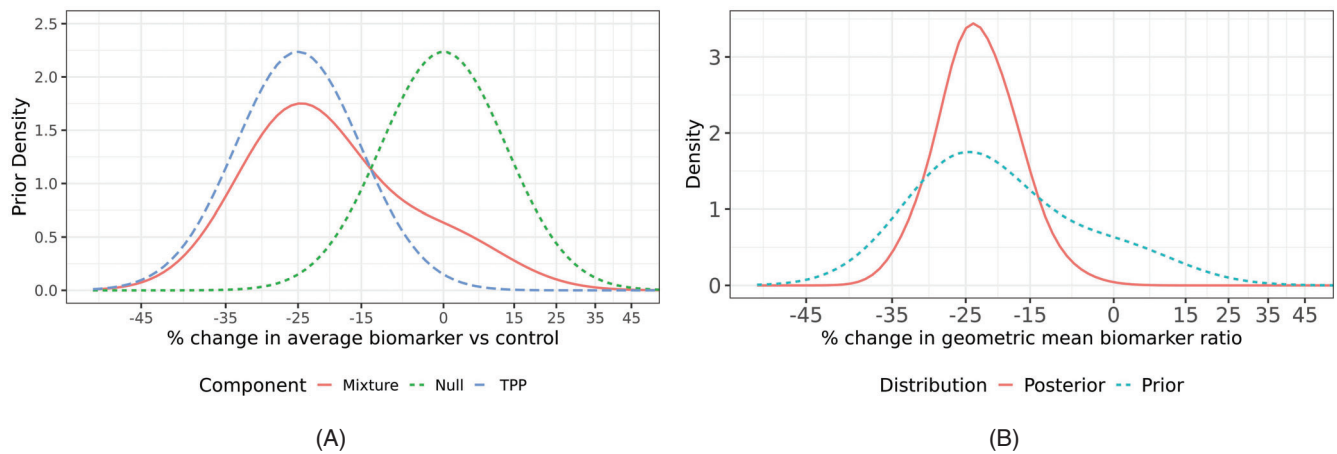


FIGURE 4 (A) Calibrated mixture prior for μ_p for the example of Section 4.4, with pdf $0.23 \times N(0, 0.03) + 0.77 \times N(-0.42, 0.03)$. A normal distribution is parameterized in terms of its mean and variance; (B) Comparing the calibrated mixture prior for μ_p with the posterior distribution obtained if we observe in phase IIa the point estimate $\hat{\theta}_{P21} = \log_2(0.77)$ with 95% CI: $\log_2(0.64)$ to $\log_2(0.92)$

Figure 4A shows the calibrated mixture prior for μ_p . From the phase IIa trial we observed $\hat{\theta}_{P21} = \log_2(0.77)$ with 95% CI $\log_2(0.64)$ to $\log_2(0.92)$, which is broadly consistent with the TPP. Figure 4B plots the posterior distribution for μ_p after fitting the Bayesian hierarchical model described in Section 3.1 to these data placing the mixture prior in Figure 4A on μ_p and stipulating $\tau_{P2} \sim HN(0.06^2)$, which has a median consistent with small between-trial heterogeneity in this context. All calculations were performed in R 3.6.1³⁰ using JAGS.³¹ To better understand the influence of the prior for μ_p in this context, Supplementary Figure D1 shows how the posterior distribution for μ_p would vary with the results of a “standard” phase IIa trial (following the standard design mentioned in Section 4.3) designed so that a t -test would have 80% power to declare superiority when $\theta_{P21} = \log_2(0.75)$ at (one-sided) significance level 0.1. We see that the posterior median is strongly attenuated towards the null if we observe a point estimate $\hat{\theta}_{P21} \ll -0.42$. The Bayesian estimate is compared with the approach of discounting $\hat{\theta}_{P21}$ by 10%.²⁹

5 | EVALUATING THE CONDITIONAL PROBABILITY OF APPROVAL AND ACCESS

We discuss how to assess $\mathbb{P}\{\text{Approval \& TPP} \mid \text{Efficacy success on 1-2 key endpoints \& no SSE in phase III}\}$, which we refer to as a program’s conditional PoS. This probability should capture risks known at the end of phase IIb which have not yet been accounted for in previous steps of the PoS evaluation described in Sections 3–4. These risks fall into five categories:

1. Regulatory alignment (R_1): Phase III design may not be aligned with regulatory expectations.
2. Unaccounted safety (R_2): Safety risks recently emerging from within the program (e.g., pre-clinical studies) and/or beyond the program (e.g., safety signals from clinical trials of a compound with the same mechanism of action) point towards an increased risk of a rare AE which, while unlikely to be detected in phase III, may raise concerns during submission. Such risks would not be captured in the SSE calculation.
3. Unaccounted TPP (R_3): Risk of not meeting the TPP on endpoints other than P and S which are necessary for approval and /or market access.
4. Quality and compliance (R_4): Known risks in quality and compliance that could jeopardize approval despite positive results on P and S. For example, poor internal audit outcome on phase II trial, issues with assay validation for key biomarker, phase III program to occur in areas with poor infrastructure.
5. Technical development (R_5): Known issues on formulation and/or device that could create uncertainties about dose selection or manufacturing.

Appendix C describes how we used industry data to fit a logistic model for the conditional probability of regulatory approval given NDA submission, adjusting for the lifecycle class of the drug and disease area. Evaluating this model for a program yields a tailored benchmark \hat{p}_{BS} . However, the logistic model does not capture the impact of $R_1 - R_5$. We propose a semi-quantitative approach to integrating these risks which begins by asking a team to score their program on a three-point risk scale (low; medium; high) for each risk R_1 to R_5 : the scorecard is included in Supplementary Materials A. A program's risk profile, defined as the configuration of the five low-med-high ratings, is then used to adjust $\hat{p}_{BS}/(1 - \hat{p}_{BS})$.

In order to translate a program's qualitative risk profile to a number that we can use to adjust the benchmark odds of approval given NDA submission, we need to understand the impact of $R_1 - R_5$ on a program's PoS. While there are no readily-available data on this, it does not mean there is no relevant evidence. To quantify what was known about the effect of $R_1 - R_5$ on PoS, we elicited the judgments of senior Novartis colleagues with experience of several submissions and market access negotiations. Each expert was asked to complete a survey listing 15 configurations of the low-med-high risk ratings: for each one, the expert was asked to state how many out of 100 hypothetical programs with the same risk profile would fail to gain approval and access despite having run a positive phase III program meeting statistical significance and the TPP on the 1–2 key efficacy endpoints without a SSE. Each survey was accompanied by a cover sheet providing background information which cited a crude historical regulatory approval rate of 90% after NDA submission. Since elicited conditional success rates were expected to be very high for some risk profiles, we preferred to ask experts for opinions on failure rates and deduce success rates from these.

Three different versions of the survey were circulated and are included in Supplementary Materials B. All experts were asked a common set of 11 questions to identify the main effects of $R_1 - R_5$. The remaining four questions were then tailored to explore one of the three pairwise interactions between R_1 , R_2 , and R_3 . In total, 46 experts spanning seven line functions were invited to participate in the survey, split between the three versions of the questionnaire (16:16:14). Experts were assigned to versions using purposeful sampling where appropriate, for example, to ensure experts in regulatory affairs received versions of the questionnaire relevant for understanding potential pairwise interactions between R_1 and R_2 , and R_1 and R_3 ; a similar strategy was applied to assign experts in safety and patient access to questionnaires.

In total, 31 of 46 experts responded. One completed survey was discarded due to a misunderstanding of the questions, meaning results are based on a denominator of 30. We model experts' individual opinions using a generalized linear mixed model, linking the average opinion on the conditional PoS to the main effects of $R_1 - R_5$ using a logit link function, and assuming a Gaussian random error term. We fit the model with a random expert intercept term in the linear predictor, treating all other terms as fixed effects. We represent $R_1 - R_5$ as categorical variables to avoid the need for assumptions about how the conditional odds of success change across risk levels. Let $\hat{e}p(r_1, \dots, r_5)$ denote the fitted conditional PoS of a program with risk profile ($R_1 = r_1, \dots, R_5 = r_5$) obtained from the mixed effects model. Figure 5 compares fitted values with elicited opinions; fitted values are also listed in Supplementary Materials C.

Recall that the tailored benchmark \hat{p}_{BS} incorporates information on a program's disease area and lifecycle class. Assuming the effects on the conditional PoS of $R_1 - R_5$, lifecycle class and disease area are additive on the logit scale, we can leverage $\hat{e}p(r_1, \dots, r_5)$ to derive a multiplicative adjustment to $\hat{p}_{BS}/(1 - \hat{p}_{BS})$. We denote this adjustment factor by $C(r_1, \dots, r_5)$. For ease of presentation, we will henceforth drop the risk arguments to $\hat{e}p$ and C .

As expert opinions were elicited with a crude benchmark conditional PoS of 0.9 in mind, the adjustment factor C must satisfy:

$$\frac{\hat{e}p}{1 - \hat{e}p} = C \frac{0.9}{1 - 0.9}, \quad (6)$$

Applying this adjustment factor to $\hat{p}_{BS}/(1 - \hat{p}_{BS})$, our estimate of the conditional odds of success reflecting information on $R_1 - R_5$, disease area and lifecycle class is given by,

$$\frac{\mathbb{P}\{\text{Approval \& TPP} \mid \text{Positive phase III on key endpoints}\}}{1 - \mathbb{P}\{\text{Approval \& TPP} \mid \text{Positive phase III on key endpoints}\}} = C \times \frac{\hat{p}_{BS}}{(1 - \hat{p}_{BS})}. \quad (7)$$

Substituting in our expression for C from (6) into (7), we obtain,




FIGURE 5 Comparing the fit of a linear mixed effects model adjusting only for the main effects of risk factors $R_1 - R_5$ (red points) with experts' stated opinions summarized by the mean ± 1 SD. Here, we code the risk levels as A (low risk); B (medium risk); C (high risk). For example, the risk profile ($R_1 = \text{low}, R_2 = \text{med}, R_3 = \text{low}, R_4 = \text{low}, R_5 = \text{med}$) appears as ABAAB

$$\mathbb{P}\{\text{Approval \& TPP} \mid \text{Positive phase III on key endpoints}\} = \frac{0.1 \times \hat{e}p \times \hat{p}_{BS}}{0.9(1 - \hat{e}p) + \hat{p}_{BS}(\hat{e}p - 0.9)}. \tag{8}$$

The final PoS estimate is then obtained as the product of $\mathbb{P}\{\text{Positive phase III program}\}$ in (1) and the conditional PoS in (8).

6 | ACCOMMODATING DIFFERENCES BETWEEN PHASE IIB AND PHASE III

So far, we have restricted attention to the relatively simple scenario that similar treatment effects are measured in phase IIB and phase III. However, differences between early phase and pivotal trials are common. Examples of differences include: pushing out the time-point at which the primary endpoint is measured; switching from measuring a biomarker to a clinical outcome; broadening-out the patient population; or refining the drug formulation in a manner which impacts on efficacy. In disease areas where the treatment landscape is rapidly evolving, we may also find the control used in phase IIB has been replaced as standard of care by the time phase III studies launch. Failure to examine the impact of these differences on the effect of treatment will make it difficult to interpret just how predictive statistical significance in phase IIB is of success in phase III. An FDA report³² highlighting 22 case-studies of phase II and III trials with divergent results included projects where positive phase II results on a short-term endpoint turned out to be inconsistent with the lack of long-term benefit in phase III.

When short- and long-term endpoints are chosen consistently across an indication, one could perform a Bayesian meta-regression of data from trials reporting pairs of effect estimates on these two endpoints³³; the association between treatment effects can then be used to bridge from the phase IIB data to derive a MAP prior for the long-term treatment effect in phase III. Alternatively, a network meta-analytic approach could be used to bridge across phases when there are differences between control arms. Both meta-analytic approaches rely on the availability of relevant historical data. However, these data are often unavailable, rendering a purely data-driven PoS evaluation impossible. This does not necessarily imply however that we are in complete ignorance about the relationship between the quantities of interest in phases II and III.

Expert elicitation is a scientific approach to quantifying knowledge about unknown parameters^{11,16} which can be adopted in this situation. There are several examples of elicited prior distributions being used to inform the design and analysis of clinical trials and drug development decisions.^{34,35,15} SHELF (the Sheffield Elicitation Framework)³⁶ is a

package of templates, software and methods intended to facilitate a systematic approach to prior elicitation minimizing the scope for heuristics and bias. We propose using the SHELF extension method³⁷ to elicit the functional relationship between effects on different endpoints including experts' uncertainty.

To illustrate how this would proceed, suppose different efficacy endpoints are studied in phases II and III and this is the only difference; applications to other scenarios follow directly. For simplicity, suppose the phase III program comprises one trial indexed by $k = 1$ and let θ_{P31} denote the effect of treatment on endpoint P in this study. Furthermore, let θ_{P2*} denote the treatment effect on the phase II primary endpoint if this were to be measured in the new phase III trial. In this setting, we can use the phase IIb data to derive a MAP prior for θ_{P2*} and then follow the SHELF extension method to elicit experts' conditional judgments on θ_{P31} given θ_{P2*} . We summarize the experts' beliefs by asking them to consider what a rational impartial observer (RIO) would believe after listening to their discussions. Then, by repeatedly sampling first from the MAP prior for θ_{P2*} and then from RIO's conditional prior distribution for $\theta_{P31} | \theta_{P2*}$, we obtain a set of Monte Carlo samples from the marginal MAP prior for θ_{P31} , which can be used to simulate the phase III trial. It is straightforward to extend this process to the case when the phase III program comprises K trials. More details on the SHELF extension method are given elsewhere.³⁷ In Section 7, we describe how we applied this approach to evaluate the probability of success for the example described in Section 4.4 where there was a change in endpoint in phase III.

7 | ILLUSTRATIVE EXAMPLE CONTINUED

We revisit the example of Section 4.4 assessing the PoS of a cardiovascular drug in lifecycle management. A single phase III trial, which we index by $k = 1$, was planned. While the primary endpoint of the phase IIa trial was a biomarker, the phase III trial would compare drug T with control on the basis of two long-term clinical outcomes: the primary endpoint (P) was the number of occurrences of a composite recurrent event endpoint, while the key secondary endpoint (S) was the time to cause-specific mortality. Let θ_{P31} , a log rate-ratio, and θ_{S31} , a log hazard ratio (HR), represent treatment effects on endpoints P and S in the phase III study. Negative effects $\theta_{P31} < 0$ and $\theta_{S31} < 0$ are consistent with a benefit of T versus control. The phase III trial was to be analyzed testing the primary null hypothesis $H_0: \theta_{P31} \geq 0$ at (one-sided) type I error rate 0.025, and the TPP threshold was a 15% reduction in the annual rate of the composite recurrent event. No significance testing was planned for the key secondary endpoint; instead, demonstrating a positive trend ($HR < 1$) would be sufficient. The team considered success on both endpoints to be essential.

We analyzed the phase IIa data on the biomarker in Section 4.4. We related the phase IIa data to the biomarker treatment effect in a new phase III study by assuming $\theta_{P2*} \sim N(\mu_P, \tau_{P2*}^2)$, where the prior distribution for μ_P is the posterior distribution given the phase IIa data summarized in Figure 4B and $\tau_{P2*} \sim \text{HN}(0.03^2)$ is centered at very small between-study heterogeneity. Figure 6 shows the MAP prior for θ_{P2*} , which places a high predictive probability of 0.998 on the event that drug T would have a beneficial effect on the biomarker in the phase III study.

We convened an elicitation workshop to quantify what was currently understood about the association between the effect of T (vs. C) on the biomarker and endpoints P and S. Four experts from Novartis were invited, three from the program team (2 statisticians, 1 clinician) and one independent clinician with knowledge of the disease area. The elicitation process largely followed the SHELF extension method.^{36,37} However, there were some deviations from that procedure as the elicitation workshop was run as an internal pilot of the elicitation process, meaning the team had the opportunity to test a modified version of the approach. Most notably, when eliciting conditional judgments we used three (rather than five) conditioning points and used the roulette method to elicit full conditional distributions at each one to avoid the need for any assumptions on the conditional variances. Furthermore, due to restrictions on time, we used mathematical rather than behavioral aggregation to obtain consensus priors. Based on our learnings from this example, we plan to adopt the SHELF extension method as described in Holzhauser et al.³⁷ for forthcoming workshops, but for transparency we describe what was actually done in the pilot meeting.

Prior to the workshop, we circulated an evidence dossier summarizing the key data as well as their limitations. During the meeting, we first elicited experts' conditional judgments on the rate ratio for the primary composite recurrent event, $\exp(\theta_{P31})$, given the biomarker treatment effect. We then elicited conditional judgments on the HR for the secondary endpoint, $\exp(\theta_{S31})$, given the biomarker treatment effect. This strategy assumes beliefs about treatment effects on P and S are conditionally independent given θ_{P2*} . We used the roulette method³⁸ to elicit from each expert a sequence of three conditional priors for $\exp(\theta_{P31})$ and $\exp(\theta_{S31})$. Experts were asked to condition their judgments on:

1. $\theta_{P2*} = -0.47$, interpreted as 28% relative reduction of geometric means between baseline and week 12.




FIGURE 6 MAP prior for θ_{P2^*} , the biomarker treatment effect in the new phase III trial, given $\hat{\theta}_{P21} = \log_2(0.77)$ with 95% CI: $\log_2(0.64)$ to $\log_2(0.92)$

2. $\theta_{P2^*} = -0.40$, corresponding to a 24% reduction
3. $\theta_{P2^*} = -0.30$, corresponding to a 19% reduction.

These conditioning values correspond to the 22, 45, and 74th percentiles of the MAP prior for θ_{P2^*} . We implemented the roulette method by asking an expert to allocate a total of 25 chips, each representing a probability of 4%, to bins covering their plausible range for the treatment effect given a particular value of θ_{P2^*} . To determine conditional priors for θ_{P31} and θ_{S31} given θ_{P2^*} , we took log-transformations of the quantiles elicited for $\exp(\theta_{P31})$ and $\exp(\theta_{S31})$. For example, if an expert stated that $\mathbb{P}\{\exp(\theta_{P31}) \leq q\} = p$ we took this to imply she/he believed $\mathbb{P}\{\theta_{P31} \leq \log(q)\} = p$. We then fitted parametric distributions to an expert's conditional opinions on θ_{P31} and θ_{S31} using the SHELF package in R.³⁹ Due to time constraints, we used mathematical, rather than behavioral, aggregation to derive “consensus” conditional priors, assigning equal weights to each expert. Supplementary Figures D2(a-c) and D3(a-c) compare individual and pooled conditional priors.

To determine a marginal prior for θ_{P31} , we began by calculating the 10th, 50th, and 90th percentiles of each of the three pooled conditional prior distributions. Let $F_p(a)$ denote the p th percentile of the pooled conditional prior for θ_{P31} given $\theta_{P2^*} = a$. For each $p = 10, 50, 90$, for simplicity we assumed a piecewise linear relationship connecting $F_p(\theta_{P2^*})$ and θ_{P2^*} , meaning we could interpolate to deduce $F_p(\theta_{P2^*})$ for any $\theta_{P2^*} \in [-0.47, -0.30]$. We extrapolated beyond the range of observation by extending the straight line connecting $F_p(-0.47)$ and $F_p(-0.40)$ to the left, and extending the straight line connecting $F_p(-0.40)$ and $F_p(-0.30)$ to the right. We then sampled from the marginal prior for θ_{P31} by following four steps:

1. Sample $\theta_{P2^*}^{(1)}, \dots, \theta_{P2^*}^{(L)}$ from the MAP prior for θ_{P2^*} .
2. Using linear interpolation, calculate $F_p(\theta_{P2^*}^{(1)})$, for $p = 10, 50, 90$. Find the best fitting statistical distribution for these percentiles, and sample $\theta_{P31}^{(1)}$ from this.
3. Repeat Step 2 to generate L samples from the marginal prior distribution of θ_{P31} .

A similar process was used to generate L samples from the marginal prior for θ_{S31} . We set $L = 40,000$. Kernel estimates of the marginal prior densities are shown in Figures 7(a,b), while the joint distribution of $(\theta_{P31}, \theta_{S31})$ is shown in Figure 8. Using a common set of samples for θ_{P2^*} in Step 1 for both endpoints induces a Spearman correlation of 0.4 between pairs of trial-specific treatment effects. For each pair of samples, we simulated one phase III trial and recorded whether we achieved the efficacy success criteria on endpoints P and S, and overall.

From Table 1, we see that for this program the probability of no SSE in phase III is 0.96, while the benchmark probability of regulatory approval after a positive phase III program (given the disease area and lifecycle class) is 0.88. The project team also completed the risk scorecard described in Section 5: they scored the program low risk on all factors except “Unaccounted TPP risks”, which they considered medium risk. Incorporating this information, we calculated that the conditional probability of obtaining approval and meeting the TPP on all endpoints needed for access given a positive Phase III program (succeeding on P and S without a SSE) is 0.80.

On the basis of the simulations of the phase III trial and information on beyond phase III risks, we estimated that the probability of:




FIGURE 7 (A) Pooled marginal prior distribution for the log rate ratio for endpoint P. 10th, 50th and 90th percentiles of this marginal prior were -0.44 , -0.23 , -0.07 , respectively. (B) Pooled marginal prior for the log-HR for endpoint S. 10th, 50th, and 90th percentiles of this marginal prior were -0.30 , -0.14 , 0.02 , respectively




FIGURE 8 Joint MAP prior for study-specific treatment effects on endpoints P and S in the phase III trial

1. Statistical significance on P in the phase III trial was 0.57.
2. Meeting the above criterion and meeting the TPP for P and a positive trend on S was 0.50.
3. Meeting the above criterion and seeing no SSE was 0.48.
4. Meeting the above criterion and obtaining approval and meeting the TPP on all remaining endpoints was 0.38.

In conclusion, the PoS of T before entering pivotal trials was retrospectively estimated at 38%.

8 | DISCUSSION

In this article, we have presented a comprehensive approach for calculating the PoS of a program at the end of phase II. Since the end of 2020, this framework has been formally adopted within Novartis to evaluate PoS before the launch of pivotal studies. While the new approach may require more time and effort than approaches based on simple industry benchmarks say, when we balance the substantial costs associated with large scale phase III trials against the benefits of the new framework, we believe this effort to be fully justified. Immediate improvements include the fact that the framework focuses on a more pertinent definition of success than statistical significance in a single trial. It also ensures all project-internal and -external information can be leveraged, using expert opinion to bridge gaps in the phase II data when these are not fully consistent with the phase III design, and assess risks beyond key phase III outcomes. Furthermore, the new approach increases consistency and transparency, allowing the identification of “pain points” for a project which can then be prospectively mitigated.

To confirm the framework increases the predictive accuracy of PoS estimates, we will need to follow-up a broad range of programs for their phase III (and possibly phase IIIb) outcomes and regulatory approval decisions. While these long-term data are not yet available, we do have several pilot programs where the phase III data are now available, meaning we can compare our prediction of the phase III results with the observed trial outcomes (for example, see Ref. 37). While these comparisons only capture the accuracy of Steps 1 and 2 of the framework, these steps are often the key drivers of the PoS evaluation; therefore these early evaluations lead us to believe the approach produces more accurate PoS estimates that can help teams to assess the adequacy of the phase III design and evaluate whether TPP targets are too ambitious. If applied early, prior to phase IIb, the process can even help teams to rethink their phase II design, for example, by considering whether the knowledge generated by measuring the phase III endpoint in phase II would offset the time and cost required to do so.

Despite the flexibility of proposed approach, not all development programs will fit perfectly into the PoS framework and further adaptations beyond those discussed in Section 4.3 will be needed. For example, lifecycle management programs which skip straight to phase III can be accommodated if it is feasible to follow Section 6 and use expert opinion to bridge phase III data from the approved indication to the effect of treatment in the new indication. Phase III data from the approved indication could be combined using a meta-analytic approach stipulating a weakly informative rather than a calibrated mixture prior for the mean of the random effects distribution. This is because selection bias is likely to be less of a concern for these data due to the size of the previous phase III studies.

In some programs we have encountered, no relevant clinical data are available at the time of the PoS assessment. In these cases, we propose calculating PoS based on the calibrated prior for the efficacy treatment effects described in Section 4. Another challenge is that for some highly innovative medicines (e.g., novel gene therapies), existing industry benchmarks may be deemed to be irrelevant. Further work is needed to identify how best to proceed in this scenario, although one idea would be to elicit expert opinion directly on the treatment effect parameter¹⁵ and use this (uncalibrated) prior to drive the PoS assessment.

One scenario that the proposed framework may not handle particularly well is the case where the new drug is truly transformational, with a treatment effect far in excess of the TPP threshold. In this setting, the calibrated mixture prior described in Section 4.2 may substantially shrink the phase II point estimate of the treatment effect towards the null. The exact degree of shrinkage will depend on both how far the point estimate overshoots the TPP target and on the industry benchmark. To handle these rare cases (without excessively shrinking the phase II data), we could modify the proposed calibrated prior to allow for surprises. For example, using a weighted average of the calibrated prior and a weakly informative prior centered at the null. Or we could define the calibrated mixture prior as a mixture of *t*-distributions rather than Gaussians. A systematic exploration of these options will be considered in future work.

Subgroup selection is common at the end of phase II and if unaccounted for, may introduce additional selection bias into the phase II effect estimate. While there are a number of approaches to correct for subgroup selection bias from a given set of subgroups (see Thomas and Bornkamp⁴⁰ for a review or Guo and He⁴¹ for recent developments), the subgroup selection process may not always be totally quantitative or driven entirely by the data observed in the study. A pragmatic approach in line with the proposed overall procedure here would be to downweight the TPP component of the phase II prior according to how plausible the selected subgroup is (e.g., if the subgroup is considered to be among the top three hypotheses before start of phase II, a multiplier of 1/3 would be applied to the TPP component). This approach will be investigated in future applications.

ACKNOWLEDGMENTS

The authors would like to thank Günther Müller-Velten, Jim Gong, Claudio Gimplewicz, Victor Shi, Wolfgang Kothny, Pritibha Singh, and Michael Wittpoth for helpful discussions during the development of this work. We would also like to thank Professor Anthony O'Hagan who facilitated the Bayesian expert elicitation meeting described in Section 7.

DATA AVAILABILITY STATEMENT

Individual patient data from the phase IIa study described in Section 4.4 are not publicly available; however, only aggregate statistics are required for the PoS calculation, which are listed in the caption to Figure 6.

ORCID

Lisa V. Hampson  <https://orcid.org/0000-0003-3074-9170>

Björn Bornkamp  <https://orcid.org/0000-0002-6294-8185>

Björn Holzhauser  <https://orcid.org/0000-0002-9209-757X>

REFERENCES

1. Kunzmann K, Grayling MJ, Lee KM, Robertson DS, Rufibach K, Wason JMS. A review of Bayesian perspectives on sample size derivation for confirmatory trials. *Am Stat*. 2021;75(4):424-432.
2. O'Hagan A, Stevens JW, Campbell MJ. Assurance in clinical trial design. *Pharm Stat*. 2005;4:187-201.
3. Ruifbach K, Burger HU, Abt M. Bayesian predictive power: choice of prior and some recommendations for its use as probability of success in drug development. *Pharm Stat*. 2016;15(5):438-446.
4. Rufibach K, Jordan P, Abt M. Sequentially updating the likelihood of success of a phase 3 pivotal time-to-event trial based on interim analyses or external information. *J Biopharm Stat*. 2016;26(2):191-201.
5. Crisp A, Miller S, Thompson D, Best N. Practical experiences of adopting assurance as a quantitative framework to support decision making in drug development. *Pharm Stat*. 2018;17(4):317-328.
6. Patel N, Bolognese J, Chuang-Stein C, Hewitt D, Gammaitoni A, Pinheiro J. Designing phase 2 trials based on program-level considerations: a case for neuropathic pain. *Ther Innov Regul Sci*. 2012;46(4):439-454.
7. Marchenko O, Miller J, Parke T, Perevozskaya I, Qian J, Wang Y. Improving oncology clinical programs by use of innovative designs and comparing them via simulations. *Ther Innov Regul Sci*. 2013;47(5):602-612.
8. Antonijevic Z, Kimber M, Manner D, Burman C-F, Pinheiro J, Bergenheim K. Optimizing drug development programs: type 2 diabetes case study. *Ther Innov Regul Sci*. 2013;47(3):363-374.
9. Hay M, Thomas DW, Craighead JL, Economides C, Rosenthal J. Clinical development success rates for investigational drugs. *Nat Biotechnol*. 2014;32:40-51.
10. Wong CH, Siah KW, Lo AW. Estimation of clinical trial success rates and related parameters. *Biostatistics*. 2019;20:273-286.
11. O'Hagan A, Buck CE, Daneshkhah A, et al. *Uncertain Judgements: Eliciting Experts' Probabilities*. Wiley & Sons; 2006.
12. Kahneman D. *Thinking, Fast and Slow*. Penguin; 2011.
13. Lo AW, Siah KW, Wong CH. Machine learning with statistical imputation for predicting drug approvals. *Harvard Data Sci Rev*. 2019; 1(1). Available from: <https://hdsr.mitpress.mit.edu/pub/ct67j043>
14. Feijoo F, Palopoli M, Bernstein J, Siddiqui S, Albright TE. Key indicators of phase transition for clinical trials through machine learning. *Drug Discov Today*. 2020;25(2):414-421.
15. Dallow N, Best N, Montague TH. Better decision making in drug development through adoption of formal prior elicitation. *Pharm Stat*. 2018;17(4):301-316.
16. O'Hagan A. Expert knowledge elicitation: subjective but scientific. *Am Stat*. 2019;73(1):69-81.
17. Carroll KJ. Decision making from phase II to phase III and the probability of success: reassured by "assurance"? *J Biopharm Stat*. 2013; 23(5):1188-1200.
18. Spiegelhalter DJ, Abrams KR, Myles JP. *Bayesian Approaches to Clinical Trials and Healthcare Evaluation*. Wiley & Sons; 2004.
19. Neuenschwander B, Capkun-Niggli G, Branson M, Spiegelhalter DJ. Summarizing historical information on controls in clinical trials. *Clin Trials*. 2010;7(1):5-18.
20. Friede T, Röver C, Wandel S, Neuenschwander B. Meta-analyses of few small studies in orphan diseases. *Res Synth Methods*. 2017;8(1): 79-91.
21. Neuenschwander B, Schmidli H. Use of Historical Data. *Bayesian Methods in Pharmaceutical Research*. 1st ed. Chapman & Hall/CRC; 2020:1-27.
22. Whitehead A. *Meta-Analysis of Controlled Clinical Trials*. John Wiley & Sons; 2002.
23. Borenstein M, Hedges LV, Higgins JPT, Rothstein HR. *Introduction to Meta-Analysis*. Wiley; 2009.
24. Kola I, Landis J. Can the pharmaceutical industry reduce attrition rates? *Nat Rev Drug Discov*. 2004;3:711-715.
25. Waring MJ, Arrowsmith J, Leach AR, et al. An analysis of the attrition of drug candidates from four major pharmaceutical companies. *Nat Rev Drug Discov*. 2015;14:475-486.
26. Harrison RK. Phase II and phase III failures: 2013-2015. *Nat Rev Drug Discov*. 2016;15:817-818.
27. Gelman A, Carlin J. Beyond power calculations: assessing type S (sign) and type M (magnitude) errors. *Perspec Psychol Sci*. 2014;9(6): 641-651.
28. Bayarri MJ, Berger J. Robust Bayesian analysis of selection models. *Ann Stat*. 1998;26(2):645-659.
29. Kirby S, Burke J, Chuang-Stein C, Sin C. Discounting phase 2 results when planning phase 3 clinical trials. *Pharm Stat*. 2012;11(5): 373-385.
30. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing; 2019.
31. M Plummer. JAGS Version 4.3.0 User Manual, 2017.
32. Food and Drug Administration. 22 case studies where phase 2 and phase 3 trials had divergent results. U.S. Department of Health and Human Services, 2017.
33. Saint-Hilary G, Barboux V, Pannaux M, Gasparini M, Robert V, Mastrantonio G. Predictive probability of success using surrogate endpoints. *Stat Med*. 2019;38(10):1753-1774.
34. Hampson LV, Whitehead J, Eleftheriou D, Brogan P. Bayesian methods for the design and interpretation of clinical trials in very rare diseases. *Stat Med*. 2014;33(24):4186-4201.
35. Ramanan AV, Hampson LV, Lythgoe H, et al. Defining consensus opinion to develop randomised controlled trials in rare diseases using Bayesian design: an example of a proposed trial of adalimumab versus pamidronate for children with CNO/CRMO. *PLoS One*. 2019;14:1-16.

36. Oakley JE, O'Hagan A. *SHELF: The Sheffield Elicitation Framework (Version 4)*. Sheffield, UK: School of Mathematics and Statistics, University of Sheffield; 2019. <http://tonyohagan.co.uk/shelf/>. Accessed August 30, 2021.
37. Holzhauser B, Hampson LV, Gosling JP, et al. Eliciting judgements about dependent quantities of interest: the SHELF extension and copula methods illustrated using an asthma case study. arXiv:2102.02852, 2021.
38. Oakley JE, Daneshkhah A, O'Hagan A. Nonparametric prior elicitation using the roulette method, 2020. Available at: <http://www.tonyohagan.co.uk/academic/pdf/elic-roulette.pdf>. Accessed August 30, 2021.
39. Oakley JE. SHELF: Tools to Support the Sheffield Elicitation Framework. R package version 1.6.0, 2019.
40. Thomas M, Bornkamp B. Comparing approaches to treatment effect estimation for subgroups in clinical trials. *Stat Biopharm Res.* 2017; 9(2):160-171.
41. Guo X, He X. Inference on selected subgroups in clinical trials. *J Am Stat Assoc.* 2021;116(535):1498-1506.
42. Jennison C, Turnbull BW. Group-sequential analysis incorporating covariate information. *J Am Stat Assoc.* 1997;92(405):1330-1341.
43. Scharfstein DO, Tsiatis AA, Robins JM. Semiparametric efficiency and its implication on the design and analysis of group-sequential studies. *J Am Stat Assoc.* 1997;92(405):1342-1350.
44. Alhussain ZA, Oakley JE. Assurance for clinical trial design with normally distributed outcomes: eliciting uncertainty about variances. *Pharm Stat.* 2020;19(6):827-839.
45. Neuenschwander B, Roychoudhury S, Schmidli H. On the use of co-data in clinical trials. *Stat Biopharm Res.* 2016;8(3):345-354.

SUPPORTING INFORMATION

Additional supporting information may be found in the online version of the article at the publisher's website.

How to cite this article: Hampson LV, Bornkamp B, Holzhauser B, et al. Improving the assessment of the probability of success in late stage drug development. *Pharmaceutical Statistics.* 2022;21(2):439-459. doi:10.1002/pst.2179

APPENDIX A: Appendices

A. Bayesian meta-analytic model for phase IIb data

Suppose the j th phase IIb trial provides an estimate $\hat{\theta}_{2j}$ of θ_{2j} . In many cases, such as when estimates are obtained from fitting a generalized linear model using maximum likelihood estimation or a Cox proportional hazards model using maximum partial likelihood,^{42,43} $\hat{\theta}_{2j}$ will follow, at least approximately after suitable transformation, a bivariate normal distribution:

$$\begin{pmatrix} \hat{\theta}_{P2j} \\ \hat{\theta}_{S2j} \end{pmatrix} | \theta_{2j} \sim N \left(\begin{pmatrix} \theta_{P2j} \\ \theta_{S2j} \end{pmatrix}, \begin{pmatrix} \mathcal{I}_{P2j}^{-1} & \kappa \sqrt{\mathcal{I}_{P2j}^{-1} \mathcal{I}_{S2j}^{-1}} \\ \kappa \sqrt{\mathcal{I}_{P2j}^{-1} \mathcal{I}_{S2j}^{-1}} & \mathcal{I}_{S2j}^{-1} \end{pmatrix} \right), \quad (\text{A1})$$

where κ is the correlation of the effect estimates, and \mathcal{I}_{P2j} and \mathcal{I}_{S2j} are the Fisher information levels for θ_{P2j} and θ_{S2j} . In the case of normally distributed responses, balanced randomization and common variances, κ will be the within-patient correlation of responses on endpoints P and S. We treat κ as known and set it equal to the estimate from phase IIb. For many types of data, information levels will depend on one or more “nuisance” parameters. For example, for normal data information levels will depend on the response variance while for binary data (under the null hypothesis of no treatment effect) they will depend on the common response probability. One approach would be to stipulate prior distributions for all unknown nuisance parameters and incorporate this uncertainty into the PoS calculation.⁴⁴ However, for simplicity, we prefer to set information levels equal to the values obtained assuming nuisance parameters are equal to their estimates based on the phase IIb data.

We assume that study-specific treatment effects in phase IIb are exchangeable, so that

$$\theta_{21}, \dots, \theta_{2J} | \mu, \tau_{P2}, \tau_{S2}, \rho \sim N \left(\begin{bmatrix} \mu_P \\ \mu_S \end{bmatrix}, \begin{bmatrix} \tau_{P2}^2 & \rho \tau_{P2} \tau_{S2} \\ \rho \tau_{P2} \tau_{S2} & \tau_{S2}^2 \end{bmatrix} \right). \quad (\text{A2})$$

We interpret τ_{P2} and τ_{S2} as the *SDs* of the phase IIb study-specific effects on endpoints P and S, while ρ is the within-study correlation of treatment effects on these two endpoints. For simplicity, we treat ρ as a fixed constant supplied by the analyst; its specification could be based on a meta-regression of pairs of treatment effect estimates obtained from trials of drugs with a similar mechanism of action to the novel drug. The Bayesian meta-analytic model for the phase IIb data is completed by stipulating priors for the average treatment effect vector $\mu = (\mu_P, \mu_S)$, and τ_{P2} and τ_{S2} . Discussion of the prior for μ is given in detail in Section 4. We follow others^{18,19,20} to stipulate weakly informative half-normal priors for the heterogeneity parameters with $\tau_{P2} \sim HN(z_{P2}^2)$ and $\tau_{S2} \sim HN(z_{S2}^2)$, where $HN(z^2)$ is the distribution of $|X|$ if $X \sim N(0, z^2)$.¹⁸ Neuenschwander and Schmidli (Table 3 of²¹) characterize different degrees of heterogeneity (large, substantial, moderate, small) in terms of multiples of the “unit-information standard deviation”, which in this context is the standard error of the effect estimate based on two patient responses (one on each arm) or a single event. We expanded this categorization to introduce a “very small” level of heterogeneity, as shown in Supplementary Materials E. We choose z_{P2} (z_{S2}) to ensure that the prior median estimate of τ_{P2} (τ_{S2}) is equal to the multiple of the unit-information *SD* corresponding to the stated degree of between-study heterogeneity. Adopting the nomenclature of Neuenschwander and Schmidli,²¹ the examples presented in this article will assume “small” between-trial heterogeneity in phase IIb for all key endpoints. The hyperparameters z_{P2} and z_{S2} will take different values if either endpoints P and S follow different distributions or, more generally, if different levels of heterogeneity are attributed to each. Given the phase IIb data $\hat{\theta}_{21}, \dots, \hat{\theta}_{2J}$, we fit the model defined in (A1)–(A2) using Markov chain Monte Carlo.

B. Handling a binary endpoint where the treatment effect summary is a risk difference

For the reasons outlined in Section 3.1, when synthesizing the phase IIb data we need to proceed slightly differently when a key efficacy endpoint is binary and the treatment effect is a difference in proportions. To outline how we proceed in this special case, suppose a single endpoint P is of interest and estimates of the response probabilities on the

TABLE A1 Measured program characteristics available in the industry benchmark dataset. Missing values on the variables Molecule, Target, and RoA were imputed using random sampling with replacement

Program Feature	Type of variable	Levels
Disease Area	Categorical	Allergy/respiratory Autoimmune/immunology/dermatology/rheumatology Cardiovascular/metabolic/renal Endocrine Hematology (malignant & benign) Infectious diseases Neurology Oncology (solid tumors) Ophthalmology Psychiatry Others
Molecule	Categorical	Small molecule, protein-antibody, protein-other, other
Target	Coded as 3 Dummy variables	Receptor, enzyme, other
Route of administration	Coded as 5 Dummy variables	Oral, intramuscular, intravenous, subcutaneous, Topical, other
Size of sponsor	Binary	Yes = sponsor is in top 20 R&D spend
Lifecycle class	Categorical	New molecular entity, lifecycle management, biosimilar
Breakthrough designation	Binary	Yes/no
Special protocol assessment	Binary	Yes/no

new drug and control are available from each phase IIb study. Rather than perform a Bayesian meta-analysis of the risk difference estimates, the analyst is instead asked to provide the sample size and number of responders per arm and study. We then use these data to run two analyses. First, we derive estimates of the study-specific log-odds ratios and combine these using a Bayesian meta-analysis based on a normal-normal hierarchical model.⁴⁵ Second, we perform a Bayesian meta-analysis of the total number of responders on control in each phase IIb study, assuming that these data follow a binomial distribution and the study-specific log-odds of response on control are samples from a normal random-effects distribution. Let p_{3k}^T , p_{3k}^C , and η_{3k} denote the study-specific probabilities of response on the new drug and control, and the log-odds ratio in the k th phase III trial. From the two analyses described above, we can obtain samples $\eta_{3k}^{(1)}, \dots, \eta_{3k}^{(L)}$ and $p_{3k}^{C(1)}, \dots, p_{3k}^{C(L)}$ from the MAP priors for η_{3k} and p_{3k}^C , respectively. The ℓ th pair of samples $(\eta_{3k}^{(\ell)}, p_{3k}^{C(\ell)})$ is transformed to obtain $(p_{3k}^{T(\ell)}, p_{3k}^{C(\ell)})$ which is used to simulate the outcome of the k th phase III trial.

C. Deriving tailored industry benchmarks

We describe below how we derived tailored benchmarks for the probability of success in phase IIa, IIb, III and submission. Tailored benchmarks are obtained from predictive models fitted to industry data. The commercial dataset we had access to contained records on 7956 programs reporting clinical trial results between 2007 and 2018: 4652 programs started phase II; 1846 started phase III; and a NDA was submitted for 1308 programs. The dataset did not distinguish between phase IIa and phase IIb trials. However, under the assumption that risks are discharged equally across stages IIa and IIb, the benchmark probability of success in phase IIb is given by the square root of the phase II benchmark. It was sufficient to use the industry data to fit logistic models for:

- $\mathbb{P}\{\text{Success in phase II}\}$: conditional probability of success in phase II given we start phase II.
- $\mathbb{P}\{\text{Success in phase III}\}$: conditional probability of success in phase III given we start phase III.
- $\mathbb{P}\{\text{Success in submission}\}$: conditional probability of regulatory approval given we submit a NDA.

TABLE A2 Program characteristics used to derive tailored benchmarks for the success probability within a development phase given a program starts that stage. Where a covariate is coded as dummy variables, the selected dummy variables are listed in parentheses

Phase	Program characteristics
II	Disease, lifecycle, molecule, target (receptor, enzyme, other), RoA (IV)
III	Disease, lifecycle, molecule, RoA (SQ, IM, other), sponsor, breakthrough
Submission	Disease, lifecycle

Table A1 shows the program characteristics available in the database and how these were coded. For both models (a) and (b), forward variable selection was used to identify important predictors of success from the following options: disease area; lifecycle class; drug molecule class; drug target; route of administration; size of sponsor; breakthrough designation; and special protocol assessment status. The last two regulatory characteristics were only considered for inclusion in model (b) because these designations can be granted at any time prior to the start of phase III and may be influenced by the phase II data. Table A2 lists the predictors that were actually selected for inclusion in each model. We needed to take a slightly different approach to fit model (c) since only 164 of the 1308 submitted programs failed to obtain regulatory approval, and this small number of events limited model complexity. We identified a limited set of predictors from discussions with drug development experts, and fitted a logistic model adjusting only for disease area and lifecycle class. Fitted values of parameters in logistic models (a-c) can be found in Supplementary Materials F.

Tailored benchmarks for the probability of success in a phase are used to calculate the probability of not seeing a SSE in phase III in Equation (4). They are also used to calculate tailored benchmarks for the probability of efficacy success in phase IIb and phase III, which themselves are needed to calibrate the prior for μ in Section 4.2. The probability of efficacy success in phase i , for $i \in \{\text{IIb}, \text{III}\}$, is given by:

$$\mathbb{P}\{\text{Efficacy success in phase } i\} = 1 - (1 - \mathbb{P}\{\text{Success in phase } i\})\mathbb{P}\{\text{Fail on Efficacy in phase } i \mid \text{Fail in phase } i\}, \quad (\text{A3})$$

where $\mathbb{P}\{\text{Fail on Efficacy in phase } i \mid \text{Fail in phase } i\}$ is 1 minus the conditional probability of failing due to a SSE in phase i under the assumption that failures are due to poor efficacy or poor safety are mutually exclusive, and the latter conditional risk is estimated using the aggregate statistics presented in Section 3.2. We take the square root of the benchmark chance of efficacy success in phase II as the phase IIb benchmark, assuming that risks are discharged equally across stages IIa and IIb.

D. Calibrating the mixture prior when there is a single efficacy endpoint

Suppose a new drug is being developed in a therapeutic area outside oncology, so that the standard phase II and phase III program comprises:

1. A single phase II trial designed to test $H_0 : \theta_{P2} \leq 0$ versus $H_1 : \theta_{P2} > 0$ with (one-sided) type I error rate α_2 at $\theta_{P2} = 0$ and power $1 - \beta_2$ at $\theta_{P2} = \delta_P$.
2. Two phase III trials designed to test $H_0 : \theta_{P3} \leq 0$ versus $H_1 : \theta_{P3} > 0$ with (one-sided) type I error rate α_3 at $\theta_{P3} = 0$ and power $1 - \beta_3$ at $\theta_{P3} = \delta_P$.

For the purposes of prior calibration, we assume there is no between-study heterogeneity, so that all trials are underpinned by a common treatment effect, denoted by μ_P , with prior

$$f(\mu_P) = \frac{\omega}{\sigma_{P1}} \phi\left(\frac{\mu_P}{\sigma_{P1}}\right) + \frac{(1-\omega)}{\sigma_{P2}} \phi\left(\frac{\mu_P - \delta_P}{\sigma_{P2}}\right),$$

where $\phi(\cdot)$ is the pdf of a standard normal random variate, $\sigma_{P1} = -\delta_P / \Phi^{-1}(0.01)$ and $\sigma_{P2} = -\delta_P / \Phi^{-1}(0.01)$. For $i = 2, 3$, let

$$c_i = \Phi^{-1}(1 - \alpha_i) \quad \text{and} \quad \mathcal{I}_i = \frac{\{\Phi^{-1}(1 - \alpha_i) + \Phi^{-1}(1 - \beta_i)\}^2}{\delta_P^2}$$

Then, we demonstrate efficacy in Phase II if and only if $Z_2 \geq c_2$, where $Z_2 | \mu_P \sim N(\mu_P \sqrt{\mathcal{I}_2}, 1)$. Similarly, we demonstrate efficacy in the k th study of the Phase III program if and only if $Z_{3k} \geq c_3$, where $Z_{3k} | \mu_P \sim N(\mu_P \sqrt{\mathcal{I}_3}, 1)$.

Letting η_2 (η_3) denote the benchmark probability of efficacy success within Phase II (III), ω is given by:

$$\omega = \frac{(\eta_2 \eta_3 - B)}{A - B},$$

where

$$A = \int_{-\infty}^{\infty} \Phi(\mu_P \sqrt{\mathcal{I}_2} - c_2) \left\{ \Phi(\mu_P \sqrt{\mathcal{I}_3} - c_3) \right\}^2 \frac{1}{\sigma_{P1}} \phi\left(\frac{\mu_P}{\sigma_{P1}}\right) d\mu_P \quad (\text{A4})$$

$$B = \int_{-\infty}^{\infty} \Phi(\mu_P \sqrt{\mathcal{I}_2} - c_2) \left\{ \Phi(\mu_P \sqrt{\mathcal{I}_3} - c_3) \right\}^2 \frac{1}{\sigma_{P2}} \phi\left(\frac{\mu_P - \delta_P}{\sigma_{P2}}\right) d\mu_P \quad (\text{A5})$$

We interpret A as the unconditional probability of demonstrating efficacy in phase II and phase III given $\mu_P \sim N(0, \sigma_{P1}^2)$, and B as the unconditional probability of efficacy success given $\mu_P \sim N(\delta_P, \sigma_{P2}^2)$. The single-fold integrals in Equations (A4) and (A5) can be evaluated numerically, for example using the integrate function in R.³⁰ The expression for ω when the new drug is an oncology therapy, so that the standard Phase III program comprises one trial, follows directly.