UNIVERSITY of York

Discussion Papers in Economics

No. 20/01

Value-based clinical trials: selecting trial lengths and recruitment rates in different regulatory contexts

Andres Alban, Stephen E. Chick, Martin Forster

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

Value-based clinical trials: selecting trial lengths and recruitment rates in different regulatory contexts

Andres Alban, Stephen E. Chick
Technology and Operations Management, INSEAD, Fontainebleau, FRANCE 77300, andres.alban@insead.edu, stephen.chick@insead.edu

Martin Forster
Department of Economics and Related Studies, University of York, Heslington, York YO42 2GA, UK, mf8@york.ac.uk

Health systems are placing increasing emphasis on improving the design and operation of clinical trials, with a view to increasing the rate of innovation and adoption of health technologies in a 'value-based' world. We present a value-based, Bayesian decision-theoretic model of a two-armed clinical trial and health technology adoption decision in which the recruitment rate and duration of the recruitment period are optimised. We account for a wide range of regulatory and practical contexts, addressing questions of how health is valued (considering discounting, the horizon of an adoption decision, and the endogenisation of outcomes for patients in the trial), and the state of clinical practice prior to commencement of the trial (we consider both exploratory trials for pharmaceutical research and pragmatic trials which compare technologies currently in use). We apply the model using research and treatment cost data from an existing trial and health technology assessment and challenge traditional perceptions concerning the efficiency, length and knowledge that may be gained from clinical research when trial teams are charged with delivering 'value' efficiently.

Key words: clinical trials; health technology assessment; cost-effectiveness; health economics; Bayesian statistics; value of information; sample size.

The need to establish 'value for money' in health care systems is becoming increasingly important, with service providers facing the dual challenges of rising demand for technologies and growing pressure on their budgets. A particular concern surrounds the evaluation of the clinical efficacy, effectiveness and cost-effectiveness of health technologies. With a large body of evidence suggesting that there is a productivity crisis in biopharmaceutical R&D (Paul et al. 2014, DiMasi et al. 2016), and with an estimated US$100 billion of public funds invested in medical research worldwide (Chakma et al. 2014), there is a growing focus on trying to improve what is often termed the 'efficiency' of clinical trial designs.

In this paper we present a value-based, Bayesian decision-theoretic model of a two-armed clinical trial and health technology adoption decision, with the goal of improving the efficiency of the health technology innovation process. This topic is timely and significant. In the United Kingdom, a National Institute for Health Research (NIHR)/Medical Research Council partnership has launched the 'Efficacy and Mechanism Evaluation' programme. This seeks to fund studies with 'novel methodological designs that deliver results more efficiently, reduce the study timeline, and

maximise the knowledge gained' (NIHR 2018). The NIHR is also funding projects which explicitly focus on innovation in trial design, via its 'Annual Efficient Studies' funding calls (NIHR 2020). In the United States, the National Institutes for Health (NIH) have launched a drive to improve the quality and 'efficiency, accountability and transparency' in clinical research (Hudson et al. 2016).

The notion of efficiency in clinical research is not always clearly or consistently defined. One common view is that greater efficiency implies that fewer patients are recruited to the trial, given predefined type I and type II error rates for identifying a smallest clinically relevant difference. Another perspective sees greater efficiency resulting from the way way trials operate, by improving retention rates, making it easier to recruit sites, and so on. In this paper, we see the idea of efficiency as being a way of designing a trial so as to maximise the value obtained for the patient population who may benefit from treatment with the health technologies under consideration, accounting for the costs of carrying out that trial, switching technologies and the costs of the technologies themselves. This proposed notion is analogous to the well-known definition from supply chains related to inefficiencies caused by double-marginalization. If a manufacturer and retailer collaborate, the supply chain as a whole can retain more value than if each optimises separately. Here, the unknown cost-effectiveness of treatments being compared takes the role of uncertainty, the trial manager is the manufacturer, the technology adopter is the retailer, and collaboration is the alignment of trial design decisions with technology adoption decisions.

Our Bayesian framework maximises the overall value of the clinical trial through the optimal choices of the trial's recruitment period, as well as its recruitment rate. It accounts for value accruing to the full patient population who may benefit from the technologies under consideration: those included in the trial, those available for inclusion in the trial but not recruited, those who require treatment while endpoints for health benefit and outcome are followed up, and those in the post-trial population. We account for trials in which there is mixed practice prior to commencement, and we consider different regulatory contexts according to whether costs and rewards are discounted and in the definition of the post trial 'patient horizon': one in which the number of patients affected by the technology adoption decision is a fixed value, the other in which it is a function of the trial's duration. We provide structural results which help characterise solutions, give comparative statics results for key performance measures for the trial with respect to disease prevalence and other parameters, and asymptotic results which shed insight about trial size. We consider both one-shot and sequential versions of the model, and apply it to a recent pragmatic trial carried out in the United Kingdom's National Health Service (Rangan et al. 2015, Handoll et al. 2015). Although Bayesian in nature, the model provides frequentist power curves, in line with FDA (2019) guidance regarding the communication of complex and innovative trial designs.

In comparison with exploration-exploitation methods in other applications, several specific features of the problem make it interesting. Firstly, the rate of recruitment of patients in a trial is capped by the disease incidence and recruitment capacity may be costly and nonlinear. Hence the problem is richer than one which involves simply selecting an optimal number of samples. Secondly, the regulatory context influences the valuation of benefits accruing to the patient population and, in turn, the expected value of information that is to be obtained from patients recruited to the trial. Thirdly, there may be significant delays between the time a trial participant is treated and the time that health outcomes and treatment costs are observed. By addressing these matters, the model assists diverse groups of health care decision makers in addressing a range of questions: (1) Clinical trial managers: What is the optimal recruitment rate for the trial? How many recruitment sites should be opened? How long should the trial run? (2) Funders of trials: is the recruitment rate and number of patients in the trial appropriate, given disease prevalence and potential health benefits? If several trials are proposed, which have the greater expected health benefit? (3) Public sector policy makers: How do the aforementioned regulatory issues affect the value of optimal trial designs? Is a trial even worth running?

We review background literature in section 1, set up the model in section 2, and present structural results for the optimal solution, comparative statics, and asymptotic properties in section 3. Section 4 illustrates how this framework can be applied in a practical trial setting. Section 5 shows how the framework can be extended to allow the recruitment rate and sampling duration to be adapted sequentially as outcomes are observed. Section 6 discusses our main results and presents directions for future work. Appendices provide supplementary information.

1. Background literature

Claxton and Posnett (1996) criticise the classical approach to clinical trial design because it ignores economic principles, such as the value of information and its cost of acquisition. They propose, as an alternative, a decision-theoretic approach using rules from cost-effectiveness analysis. Draper (2013) advocates for the use of a Bayesian decision-theoretic approach which uses a utility function comprising clinically relevant outcomes, such as Quality Adjusted Life Years (QALYs). In line with these initiatives, a range of Bayesian decision-theoretic models have been proposed as alternatives to the classical approach (Claxton et al. 2000, Gittins and Pezeshk 2000, Willan and Pinto 2005, Eckermann and Willan 2007, Griffin et al. 2010). These are based on a comparison of the cost of carrying out research with the value that the additional research generates, using so-called 'value of information' calculations (Raiffa and Schlaifer 1961, Brennan et al. 2007, Strong et al. 2015).

The majority of these decision-theoretic models concern one-shot trials. More recently, interest has grown in approaches that are adaptive in nature (Pertile et al. 2014, Ahuja and Birge 2016,

Villar et al. 2015a,b, Williamson et al. 2017, Chick et al. 2017, 2018). For example, Chick et al. (2017) solve a Bayesian decision-theoretic model of a two arm clinical trial with delay in observing the outcome for cost-effectiveness. Villar et al. (2015b) consider adaptive allocation to treatments within a multiple-arm setting using the Gittins Index. The related Bayesian ranking and selection literature has proposed various combinations of discounted or undiscounted rewards and online and offline learning (Branke et al. 2007, Frazier et al. 2008, Ryzhov et al. 2012, Russo 2020). None of these works have considered optimising the rate of recruitment, although this may be useful in applications besides clinical trials, such as when the optimal treatment choice from an A/B test has a reward which is highly time-sensitive.

Although it is natural for a Bayesian approach to maximise expected value, or equivalently minimise expected regret, it would also be possible to consider frequentist approaches to expected regret (e.g., see Chick and Wu 2005 for frequentist regret in ranking and selection, or the rich literature on asymptotic regret in bandit problems (Bubeck and Cesa-Bianchi 2012)). It is also possible to have Bayesian beta-bernoulli models in clinical trial design for sequential allocation in 0-1 trials (Villar et al. 2015a, Williamson et al. 2017). We choose a Bayesian, value-based framework to be consistent with the UK's National Institute for Health and Care Excellence (NICE) guidance for uncertainty quantification for probabilistic sensitivity analysis for health technology assessments (NICE 2012, Section 7).

There exists a range of other approaches to value in clinical trial design. Some consider changing the balance of allocation to treatment arms as a function of the past history of allocations. Examples include Berry and Eick (1995) and Villar et al. (2015b). Others maintain balanced allocation, but allow for the trial to stop at any stage of the process as a function of the accumulating evidence. Examples include Berry and Ho (1988), Chick et al. (2017) and Jennison and Turnbull (1989). Although our principal focus is on one-shot designs, we also provide extensions for adaptive trials,

The societal perspective to measuring value that is adopted in this work contrasts with the Bayesian decision-theoretic contributions of Gittins and Pezeshk (2000) and Willan (2008), which consider a trial's optimal choice of sample size from the industry perspective. In these studies, the terminal reward of the trial is a function of the probability that the technology is approved by a regulator and the market share that it may gain. Jobjörnsson et al. (2016) consider the optimal sample size and pricing decision for a new pharmaceutical product, given uncertainty over an insurer's willingness to pay and a prior distribution for efficacy. While contracting for incentive alignment is an interesting question, our model focuses on the social welfare approach to value, noting that contracting for public-private procurement is outside of scope.

Implementing a value-based trial in practice requires collecting cost and QALY data (or some other health outcome which can be converted to money). While many clinical trials do not have

QALYs as a primary endpoint, much less data on treatment costs, accounting for QALYs in clinical trials has seen increasing attention (e.g., Angus et al. 2001, Ferguson et al. 2013, NICE 2014). Costs and QALYs are also part of the rich tradition in operations research, outside of clinical trial design, in resource planning for health interventions (Long et al. 2008, is one of many examples).

To the best of our knowledge, a model addressing scenarios with mixed clinical practice has never been proposed in the literature. This is surprising, given that they are commonly encountered in clinical research (the ProFHER trial (Handoll et al. 2015), Caesarean section versus natural birth (Betrán et al. 2016), plates versus wires for certain fractures (Costa et al. 2014), etc.). Generalising the model in this direction has wide-ranging implications for the way in-trial benefits are valued, the way switching costs and patient horizons are handled when a technology adoption decision is made, and the choice set that is available to the trial team. Further, the above contributions provide limited analysis of the optimal value and trial length under differing mechanisms for how the post-trial patient population is defined, a matter that we address in detail. Finally, we provide both analytical comparative statics and numerical sensitivity analyses for some of the key parameters in the model, lending insight into where we can claim definitive results for directional changes, as well as applications which illustrate their absolute size.

The operations literature is also interested in the management of research and development (R&D) projects, including for the pharmaceutical pipeline. Jacob and Kwak (2003) present a real option approach to valuing such projects in response to changes in the health care economy and scientific advances. Girotra et al. (2007) explore the value of portfolio management with Phase III drug trials. These works call for operational flexibility and management of the pipeline process. A control perspective on R&D investment decisions has been taken by a number of authors who address uncertainties in costs and durations of clinical trial processes (Lucas Jr 1971, Schwartz 2004). Kouvelis et al. (2017) further link data from trials to a theoretical model of recruitment rate optimisation decisions. Our work differs from these streams, in that we focus on uncertainties related to the efficacy (and hence health benefits) of the technology adoption decision, rather than on the uncertainties in the cost streams or patient accrual, and our focus is on optimising a rate of recruitment and trial length.

2. Mathematical model of a value-based clinical trial

We present a Bayesian decision-theoretic model of a clinical trial comparing two health technologies on cost-effectiveness grounds. The objective is to maximise the monetary value of health benefits generated for the target population, minus the financial costs of carrying out trial and any costs incurred in technology adoption. Below we discuss choice of decision variables, outcome measure, objective function and the regulatory jurisdictions addressed by the model. A table of principal notation is presented in Appendix A.

2.1. Trial design and decision variables

A clinical trial compares two health technologies – N ('new') and S ('standard') – on cost-effectiveness grounds. We consider trials in which no patients are being treated with technology N prior to commissioning (exploratory trials), as well as trials in which there exists mixed practice (pragmatic trials like ProFHER). Define $p_{\mathrm{N}} \in [0, 1/2]$ as the fraction of the eligible population which is treated with technology N at the trial's inception, so that the fraction $1 - p_{\mathrm{N}}$ is treated with technology S. Without loss of generality, assume that a greater proportion of patients receive S (it is always possible to rename the two technologies to satisfy this constraint).

The trial design comprises three decision variables: the adoption decision rule, $\mathcal{D}$, the recruitment period duration, T, and the recruitment rate, r, which is assumed constant. $\mathcal{D} \in \{\mathrm{S}, \mathrm{N}, \mathrm{M}\}$ is a function that maps the state of information at the end of a trial to a treatment allocation decision for all patients to be treated upon the trial's conclusion. If $\mathcal{D} = \mathrm{S}$, all patients are treated with the standard technology; if $\mathcal{D} = \mathrm{N}$, all patients are treated with the new technology; $\mathcal{D} = \mathrm{M}$ implies that patients are treated according to the mix that was in place prior to the trial's inception.

The duration of the recruitment period, T, is constrained to lie in the interval $[0, T_{\max}]$, where factors such as regulatory or funding requirements might constrain $T_{\max}$. The recruitment rate, r, is constrained to lie in $[0, r_{\max}]$. The maximum permissible recruitment rate, $r_{\max}$, is determined by factors such as the availability of recruitment sites and the incidence rate of the condition, $\zeta \in \mathbb{R}_{>0}$. A decision to make $T = 0$ or $r = 0$ means that the trial does not run and the adoption decision is made using information available at the start of the trial alone.

Given these assumptions, the sample size of the trial is equal to rT. In contrast to the majority of the literature, which focuses on sample size alone as the decision variable, we are able to study the impact of both choice of trial length and recruitment rate on the total value generated by the trial. We assume that patients are randomised to the two arms of the trial in pairs, and define $Q = Tr/2 \in [0, Q_{\max}]$ as the number of recruited pairs, where $Q_{\max} \leq T_{\max} r_{\max}/2$.

If the trial runs, it is assumed that it incurs costs, $c_{\mathrm{cap}}(r)$, which reflect set-up, recruitment site and post-trial costs (such as those incurred in writing up and disseminating results). For example, one model is $c_{\mathrm{cap}}(r) = c_{\mathrm{fixed}} + c_r \lceil r/x \rceil$ where c_{fixed} is the fixed cost, $c_r \geq 0$, and x patients per unit time can be recruited per site. A model which accounts for an increasing cost of marginal capacity is $c_{\mathrm{cap}}(r) = c_{\mathrm{fixed}} + c_r r^3$.

2.2. Value-based criterion for technology adoption decision

The value-based nature of the trial requires that we compare health outcomes and costs in a common currency. We use the difference between the net monetary benefit of using treatment N over S – the 'incremental net monetary benefit' (INMB) (Gold et al. 1996):

$$W = \lambda(E_{\mathrm{N}} - E_{\mathrm{S}}) - (C_{\mathrm{N}} - C_{\mathrm{S}}), \tag{1}$$

Figure 1 **A fixed horizon trial. Each area represents the number of patients in the following classes: (1) recruited to the trial; (2) not recruited to the trial during the recruitment period; (3) treated with current practice during waiting period; (4) benefiting from the adoption decision.**

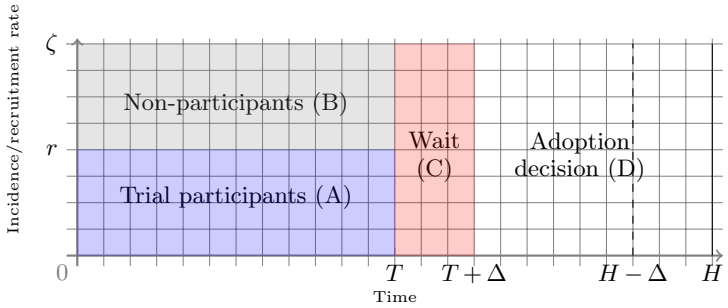

where $E_j \in \mathbb{R}$ and $C_j \in \mathbb{R}$ measure expected effectiveness (health or quality of life improvement) and cost of technology $j \in \{\mathrm{N}, \mathrm{S}\}$, respectively. The parameter $\lambda \in \mathbb{R}_{>0}$ is the monetary valuation of one unit of effectiveness, defined by the regulatory body responsible for the study population.[1] The value of W is not known a priori. Its prior distribution is assumed to be $W \sim \mathcal{N}(\mu_0, \sigma_0^2)$, as might be informed by phase II trials or a pilot study. Let $n_0 = \sigma_X^2/\sigma_0^2$ be the effective sample size of the prior distribution. Observations of INMB arrive with a fixed delay of $\Delta \geq 0$ units of time after treatment allocation and all outcomes must be observed before the adoption decision is made.

2.3. Objective function: expected net gain

The expected net gain of the trial is the difference between the expected values of: (1) running the trial and treating post-trial patients with the technology recommended by its results and (2) continuing to treat patients according to the practice in place before the trial commenced. Figure 1 shows how the expected net gain of the target population may be divided into four constituent parts. The horizontal axis plots time, covering the horizon, $[0, H]$, over which the trial operates and its results are to be used for treating patients, for some $H > \Delta$. The vertical axis plots the incidence and recruitment rates of patients. Areas labelled are: (A) patients recruited to the trial; (B) patients not recruited to the trial during the recruitment period; (C) patients in the 'waiting period'; (D) post-adoption patients treated with the recommended technology.

2.3.1. Patients recruited to the trial. $Tr/2$ pairs of patients are randomised during the trial. The trial incurs a recruitment cost, c, for each patient enrolled. For patient pair $i \in 1, \ldots, Q_{\max}$, we observe a noisy observation of INMB, defined as X_i. We assume that $X_i \mid W \overset{\text{i.i.d.}}{\sim} \mathcal{N}(W, \sigma_X^2)$, where σ_X^2 is the known sampling variance. The expected net gain for patients recruited to the trial is $(Tr/2)(1-2p_{\mathrm{N}})\mathbb{E}[W] - cTr$. The proportion receiving a different treatment under balanced randomization, compared to what they would have received had the trial not taken place, is $1-2p_{\mathrm{N}}$.

[1] The UK's National Institute for Health and Care Excellence may value one quality-adjusted life year at £20,000 (NICE 2013).

2.3.2. Patients not participating in the trial during the recruitment period. During the recruitment period, $(\zeta - r)T$ patients are not enrolled in the trial. We assume that these patients continue being treated with the practice in place before the trial commenced. Hence they do not affect the expected net gain (they incur no additional cost and their outcomes are the same as they would have been had the trial not taken place).

2.3.3. Waiting period. The adoption decision is made at time $T + \Delta$. During the period $[T, T + \Delta]$, patients are treated according to the practice that was in place prior to the start of the trial and so these patients do not contribute to the expected net gain.

2.3.4. Post-adoption patients. Define $I_{\mathcal{D}}$ as the total cost of technology adoption. We assume that $I_{\mathrm{M}} = 0$ because neither technology is adopted, and $I_{\mathrm{N}}, I_{\mathrm{S}} \geq 0$. Let P be the number of post-adoption patients, assumed to be known at the start of the trial. If $\mathcal{D} = \mathrm{N}$, the expected net gain for these patients is $P(1 - p_{\mathrm{N}})\mathbb{E}[W] - I_{\mathrm{N}}$, where $P(1 - p_{\mathrm{N}})$ is the number of patients who, absent the trial, would be treated with technology S. If $\mathcal{D} = \mathrm{S}$, the expected net gain is $-Pp_{\mathrm{N}}\mathbb{E}[W] - I_{\mathrm{S}}$. If $\mathcal{D} = \mathrm{M}$, the expected net gain is zero because the trial did not change practice. Thus the expected net gain for this portion of the population is $\mathbb{E}[\mathbf{1}_{\mathcal{D}=\mathrm{N}}(P(1-p_{\mathrm{N}})W - I_{\mathrm{N}}) + \mathbf{1}_{\mathcal{D}=\mathrm{S}}(-Pp_{\mathrm{N}}W - I_{\mathrm{S}}) \,|\, \mathcal{D}, T, r]$, where $\mathbf{1}_F$ is the indicator function, equal to one if F is true and zero otherwise. We condition on $\mathcal{D}$, T, and r to clarify that the expectation depends on $\mathcal{D}$ after the outcomes from $rT/2$ pairwise allocations have been observed.

2.3.5. Expected net gain. Define the expected net gain of a trial design by $V(T, r, \mathcal{D})$. If the trial does not run, the post-adoption population is the entire population, and the expected net gain is:

$$V(T, 0, \mathcal{D}) = V(0, r, \mathcal{D}) = \mathbb{E}[\mathbf{1}_{\mathcal{D}=\mathrm{N}}(P(1-p_{\mathrm{N}})W - I_{\mathrm{N}}) + \mathbf{1}_{\mathcal{D}=\mathrm{S}}(-Pp_{\mathrm{N}}W - I_{\mathrm{S}}) \,|\, \mathcal{D}, T, r]. \tag{2}$$

If the trial recruits at least one pair of patients, the expected net gain is the sum of the expected net gains of the enrolled and post-adoption patients, minus the trial costs:

$$\begin{aligned} V(T, r, \mathcal{D}) = &- \underbrace{(c_{\mathrm{cap}}(r) + cTr)}_{\text{trial cost}} + \underbrace{\delta_{\mathrm{on}}(Tr/2)(1 - 2p_{\mathrm{N}})\mathbb{E}[W]}_{\text{trial participants}} \\ &+ \underbrace{\mathbb{E}[\mathbf{1}_{\mathcal{D}=\mathrm{N}}(P(1-p_{\mathrm{N}})W - I_{\mathrm{N}}) + \mathbf{1}_{\mathcal{D}=\mathrm{S}}(-Pp_{\mathrm{N}}W - I_{\mathrm{S}}) \,|\, \mathcal{D}, T, r]}_{\text{post-adoption}}. \end{aligned} \tag{3}$$

Use of the indicator variable, δ_{on}, permits us to model 'online learning' ($\delta_{\mathrm{on}} = 1$, so that benefits to trial participants are counted in the calculation of expected net gain) as well as 'offline learning' ($\delta_{\mathrm{on}} = 0$, so that benefits are not counted).

2.4. The regulatory context

The above model already handles one important matter reflecting the regulatory context in which the trial is conducted, namely online or offline learning. Online learning may be particularly relevant in trials for orphan diseases. In addition to online learning, we extend the model to handle two further matters concerned with the regulatory context: the size of the population that is affected by the technology adoption decision and whether rewards are discounted or not.

2.4.1. Post-adoption population. We allow the number of post-adoption patients to depend on the duration of the trial: $P(T)$, which we assume is not increasing in T. We focus our analysis further on two possible cases motivated by potential scenarios in a health technology development:

1. There exists a *fixed pool* of patients, $P(T) = P$.
2. There exists a *fixed horizon*, $H \geq T_{\max} + \Delta$, defined prior to the start of the trial and covering both the trial and adoption horizons, so that $P(T) = \zeta(H - T - \mathbf{1}_{T>0}\Delta)$.

Case 1 is motivated by regulatory regimes which grant exclusive marketisation rights for a health technology for a defined period post-authorisation (see, e.g., FDA 2015), but we use it in a more general sense to refer to situations in which there exists a fixed patient population to be treated post-adoption (for example, when a regulator approves a health technology for a defined period, prior to reviewing additional evidence). Case 2 is motivated by situations in which a patent protection agreement operates from time 0 for a fixed time horizon H, during which the trial may be run and the adoption decision implemented.[2]

2.4.2. Discounted expected net gain. The foregoing assumes that costs and benefits during the trial are not discounted, an assumption which may be realistic for some jurisdictions but not others. For example, discounting is not used in the design of 'traditional' clinical trials which use statistical criteria to determine the sample size, whereas the UK's NICE (2013) has recommended an annual discount factor of 3.5% for health technology assessment decisions. Under continuous time discounting at the rate $\rho > 0$, define the discounted recruitment period duration as $\tilde{T}_\rho(T)$ and the discounted post-adoption number of patients as $P_\rho(T)$, under the assumption that the $P(T)$ patients arrive at a constant rate, ζ, over a duration of time $P(T)/\zeta$:

$$\tilde{T}_\rho(T) = \int_0^T e^{-\rho s} ds = \rho^{-1}(1 - e^{-\rho T}); \; P_\rho(T) = \int_0^{P(T)/\zeta} \zeta e^{-\rho s} ds = (\zeta/\rho)(1 - e^{-\rho P(T)/\zeta}), \qquad (4)$$

and $\tilde{T}_\rho(T) = T$ and $P_\rho(T) = P(T)$ if $\rho = 0$.

[2] Patent protection may be applied for at any time during the development process of a new drug (FDA 2015). Our use of the term is more restrictive, in that the protection is assumed to apply over the interval $[0, H]$.

We can handle discounting by letting $P(T) = P_\rho(T)$ in (2) and (3):

$$V(T,0,\mathcal{D}) = V(0,r,\mathcal{D}) = \mathbb{E}\left[\mathbf{1}_{\mathcal{D}=\mathrm{N}}(P_\rho(0)(1-p_\mathrm{N})W - I_\mathrm{N}) + \mathbf{1}_{\mathcal{D}=\mathrm{S}}(-P_\rho(0)p_\mathrm{N}W - I_\mathrm{S}) \mid \mathcal{D}, T, r\right], \quad (5a)$$

$$\begin{aligned} V(T,r,\mathcal{D}) = &-(c_\mathrm{cap}(r) + c\tilde{T}_\rho(T)r) + \delta_\mathrm{on}(\tilde{T}_\rho(T)r/2)(1-2p_\mathrm{N})\mathbb{E}\left[W\right] \\ &+ e^{-\rho(T+\Delta)}\mathbb{E}\left[\mathbf{1}_{\mathcal{D}=\mathrm{N}}((1-p_\mathrm{N})P_\rho(T)W - I_\mathrm{N}) + \mathbf{1}_{\mathcal{D}=\mathrm{S}}(-p_\mathrm{N}P_\rho(T)W - I_\mathrm{S}) \mid \mathcal{D}, T, r\right], \end{aligned} \quad (5b)$$

where (5b) applies when the trial runs ($rT > 0$).

2.5. One-shot optimal trial design problem

Our focal optimal value-based trial design problem is

$$\begin{aligned} V^* = &\max_{T,r,\mathcal{D}} V(T,r,\mathcal{D}) \\ &\text{s.t. } T \in [0, T_\mathrm{max}],\ r \in [0, r_\mathrm{max}], \end{aligned} \quad (6)$$

where $V(T,r,\mathcal{D})$ is defined by (5a) when the trial does not run and is otherwise defined by (5b). The expectations which determine $V(T,r,\mathcal{D})$ are with respect to the prior distribution for W in section 2.2 (which can be informed by prior pilot studies or expert opinion) and the likelihood for the samples in section 2.3.1 which inform $\mathcal{D}$. We refer to this problem as a 'one-shot' trial because it fixes the trial parameters T, r at the start and does not vary them as the trial progresses. The model is general, in the sense that it accounts for all regulatory context options in section 2.4.

3. Analysis of the one-shot optimal trial design problem

We first show that the optimal adoption decision rule is easily found. Then, we consider the structural properties of T^* and r^*. Some mild assumptions guarantee the existence of a solution. We then provide comparative statics and explore the asymptotics of the optimal solution.

3.1. Optimal adoption decision rule

Define Z_{Tr} as the posterior mean of W given that realisations of incremental net monetary benefit for $rT/2$ pairwise allocations will be observed:

$$Z_{Tr} \equiv \mathbb{E}\left[W \mid X_1, \ldots, X_{rT/2}\right].$$

Define $\sigma_Z^2 = \sigma_X^2(rT/2)/[n_0(n_0 + rT/2)]$. Then it can be shown that (DeGroot 1970)

$$\begin{aligned} Z_{Tr} = \frac{n_0\mu_0 + \sum_{i=1}^{rT/2} X_i}{n_0 + rT/2} \sim \mathcal{N}\left(\mu_0, \sigma_Z^2\right), \\ \text{and } W \mid Z_{Tr} \sim \mathcal{N}(Z_{Tr}, \sigma_X^2/(n_0 + rT/2)). \end{aligned} \quad (7)$$

Thus, Z_{Tr} is a sufficient statistic for all the information obtained in the trial, and $\mathcal{D}^*$ is a function of Z_{Tr} instead of the sequence X_i of all observations.

Figure 2 **Predictive distribution for the posterior mean and regions where it is optimal to select N, S, or M.**

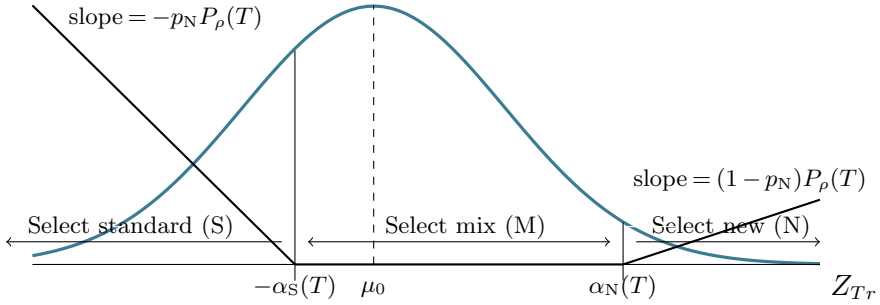

The optimal adoption decision rule maximises the benefits for the $P_\rho(T)$ post-adoption patients:

$$\mathcal{D}^*(Z_{Tr}) = \underset{\mathcal{D}\in\{\mathrm{M,N,S}\}}{\arg\max}\, \{0\,,\, P_\rho(T)(1-p_{\mathrm{N}})Z_{Tr} - I_{\mathrm{N}}\,,\, -P_\rho(T)p_{\mathrm{N}}Z_{Tr} - I_{\mathrm{S}}\,\},$$

where the terms inside the arg max operator are the post-adoption expected net gains for each of M, N and S, respectively.

Let $\alpha_{\mathrm{N}}(T) = I_{\mathrm{N}}/((1-p_{\mathrm{N}})P_\rho(T))$ and $\alpha_{\mathrm{S}}(T) = I_{\mathrm{S}}/(p_{\mathrm{N}}P_\rho(T))$ be the expected per patient switching costs given adoption of technology N or S, respectively. To avoid undefined expressions, we set

$$\alpha_{\mathrm{S}}(T) = \begin{cases} 0, & I_{\mathrm{S}} = 0 \\ \infty, & p_{\mathrm{N}} = 0 \text{ or } P_\rho(T) = 0; \end{cases} \qquad \alpha_{\mathrm{N}}(T) = \begin{cases} 0, & I_{\mathrm{N}} = 0 \\ \infty, & P_\rho(T) = 0. \end{cases}$$

The optimal adoption rule divides the open interval for posterior beliefs into three regions, delineated by $\alpha_{\mathrm{N}}(T)$ and $\alpha_{\mathrm{S}}(T)$: if $Z_{Tr} > \alpha_{\mathrm{N}}(T)$, it is optimal to adopt N; if $Z_{Tr} < -\alpha_{\mathrm{S}}(T)$, it is optimal to adopt S, otherwise it is optimal to continue with the current mix. We refer to $\alpha_{\mathrm{N}}(T)$ and $\alpha_{\mathrm{S}}(T)$ as 'indifference points', because one is indifferent between N and M when $Z_{Tr} = \alpha_{\mathrm{N}}(T)$, and one is indifferent between S and M when $Z_{Tr} = -\alpha_{\mathrm{S}}(T)$. Figure 2 shows a distribution for Z_{Tr} and the rewards for the optimal adoption decision for a prior mean lying between the indifference points. The slopes of the linear reward functions are given by $-p_{\mathrm{N}}P_\rho(T)$ when S is adopted and $(1-p_{\mathrm{N}})P_\rho(T)$ when N is adopted.

3.2. Simplified objective function

Given the optimal adoption decision, we now define, with some abuse of notation, the expected net gain as a function of T and r: $V(T,r) = V(T,r,\mathcal{D}^*)$. Thus, (5b) simplifies to the following expression by conditioning on Z_{Tr}, setting $x^+ = \max\{0,x\}$, and using the tower property of expectations:

$$\begin{aligned} V(T,r) = &-\left(c_{\mathrm{cap}}(r) + c\tilde{T}_\rho(T)r\right) + \delta_{\mathrm{on}}(\tilde{T}_\rho(T)r/2)(1-2p_{\mathrm{N}})\mathbb{E}\left[\,Z_{Tr}\,\right] \\ &+ e^{-\rho(T+\Delta)}\mathbb{E}\left[\,(P_\rho(T)(1-p_{\mathrm{N}})Z_{Tr} - I_{\mathrm{N}})^+ + (-P_\rho(T)p_{\mathrm{N}}Z_{Tr} - I_{\mathrm{S}})^+\,\right]. \end{aligned} \tag{8}$$

Define the normal linear loss function $\Psi(z) \equiv \mathbb{E}\left[(Z-z)^+\right] = \phi(z) - z(1-\Phi(z))$, where Z is a standard normal random variable with cumulative distribution function Φ and probability density function ϕ. When $Tr > 0$, (8), and therefore (5b), can be computed as follows:

$$\begin{aligned} V(T,r) = &-\left(c_{\text{cap}}(r) + c\tilde{T}_\rho(T)r\right) + \delta_{\text{on}}(\tilde{T}_\rho(T)r/2)(1-2p_{\text{N}})\mu_0 \\ &+ e^{-\rho(T+\Delta)}P_\rho(T)\sigma_Z\left[(1-p_{\text{N}})\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right) + p_{\text{N}}\Psi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right]. \end{aligned} \tag{9}$$

The probabilities of adopting N, S and M are, respectively, $1-\Phi((\alpha_{\text{N}}(T)-\mu_0)/\sigma_Z)$, $1-\Phi((\alpha_{\text{S}}(T)+\mu_0)/\sigma_Z)$, and $\Phi((\alpha_{\text{N}}(T)-\mu_0)/\sigma_Z) + \Phi((\alpha_{\text{S}}(T)+\mu_0)/\sigma_Z) - 1$. When $rT = 0$,

$$V(T,r) = \max\{0,\ (1-p_{\text{N}})P_\rho(0)\mu_0 - I_{\text{N}},\ -p_{\text{N}}P_\rho(0)\mu_0 - I_{\text{S}}\} \tag{10}$$

In the remainder of this section, we focus the analysis to solving a simplified version of (6):

$$\begin{aligned} V^* = &\max_{T,r} V(T,r) \\ &\text{s.t. } T \in [0, T_{\max}],\ r \in [0, r_{\max}], \end{aligned} \tag{11}$$

3.3. Structural properties of the optimal trial design

The following proposition proves the existence of an optimal fixed sample size trial design under two reasonable assumptions. Firstly, we assume that $c_{\text{cap}}(r)$ is non-decreasing, that is, additional recruitment capacity is costly, and lower semi-continuous, a mathematical condition for the existence of a solution. This assumption is not restrictive because it accepts any continuous, increasing function. Secondly, we assume that $P(T)$ is non-increasing, that is, the post-trial population does not increase with a longer trial. To guarantee the existence of a solution, we assume that $P(T)$ is bounded and upper semi-continuous. The two special cases (fixed patient pool and fixed horizon) that were introduced in section 2.4.1 both satisfy these conditions, as does the example cost of capacity model in section 2.1. The proof of the following proposition is presented in Appendix C.1.

PROPOSITION 1. *If $c_{cap}(r)$ is non-decreasing and lower semi-continuous, and $P(T)$ is non-increasing, bounded, and upper semi-continuous, then an optimal solution (T^*, r^*) to (11) exists.*

A closed-form solution to (11) is not available, but first order conditions may be obtained (see Appendix B). The function $V(T,r)$ is not guaranteed to have a unique local optimum, so the global optimum is found by starting a common optimisation algorithm at several random points.[3]

[3] In our numerical experiments, we have not found more than two local optima for the product Tr.

3.4. A taxonomy of value-based designs and their solutions

Table 1 presents a taxonomy of commonly encountered fixed sample size clinical trials which may be solved using our model. We classify the trials as cases I–IV, according to whether the recruitment site cost function, c_{cap}, is constant and whether the trial incorporates discounting and a fixed patient pool. For each case, Table 1 records the version of (11) which should be solved. The following propositions show that, for cases I–III but not case IV, the problem reduces to the optimal choice of a single variable: T or r or their product. These cases are useful for comparative statics results in section 3.5.

3.4.1. Case I: constant costs, undiscounted rewards and fixed patient pool. Choose the decision variables T and r so their product optimises the number of pairwise allocations, Q^*. This gives some flexibility in selecting T as long as Q^* is optimised.

PROPOSITION 2. *If $c_{cap}(r) = c_{cap}$, $\rho = 0$, $P(T) = P$, then all members of the set $\mathcal{S} = \{(T, r) \in [0, T_{\max}] \times [0, r_{\max}] : rT/2 = Q^*\}$, for some Q^* that represents the optimal number of pairwise allocations, are solutions of (11).*

As a corollary of Prop. 2, the marginal benefit of an additional unit of recruitment for Case I is exactly 0: an increase in r is accompanied by a proportional decrease in T to retain the same optimal Q^*. This is because a postponement of rewards is costless from the perspective of both marginal costs of recruitment and the benefit to the adopting population.

3.4.2. Case II: constant costs and discounted rewards and/or variable patient pool. For such trials, an increase in the recruitment rate accrues more benefits to patients by permitting an earlier adoption decision, without incurring additional cost. It is optimal to recruit as fast as possible and to optimise over T.

PROPOSITION 3. *If $c_{cap}(r) = c_{cap}$, there is an optimal solution (T^*, r^*) to (11) with $r^* = r_{\max}$:*

$$V^* = \max_{T \in [0, T_{\max}]} V(T, r_{\max}). \tag{12}$$

3.4.3. Case III: nonconstant costs, undiscounted rewards and fixed patient pool. The presence of a fixed patient pool and undiscounted rewards means there is no penalty for recruiting later rather than earlier. Without loss of generality, it is optimal to run the trial for as long as possible and optimise over the recruitment rate.

PROPOSITION 4. *If $P(T) = P$, $\rho = 0$, there is an optimal solution (T^*, r^*) to (11) with $T^* = T_{\max}$.*

$$V^* = \max_{r \in [0, r_{\max}]} V(T_{\max}, r). \tag{13}$$

Table 1 **A taxonomy of value-based designs for commonly encountered fixed sample size clinical trials.**

	Undiscounted fixed patient pool	Otherwise
Constant setup costs	**I**. Optimise sample size (section 3.4.1, Proposition 2)	**II**. Fix $r^* = r_{\max}$, optimise over T (section 3.4.2, Proposition 3)
Otherwise	**III**. Fix $T^* = T_{\max}$, optimise over r (section 3.4.3, Proposition 4)	**IV**. Optimise over both r and T (section 3.4.4)

3.4.4. Case IV: nonconstant costs and discounted rewards and/or fixed patient pool. It necessary to optimise over both r and T to solve (11).

3.5. Comparative statics.

We now discuss comparative statics to assess how the optimal expected net gain V^* and the optimal decision variables T^*, r^* and Q^*, depend on key parameters P, H, n_0 and p_{N}. These results are summarised in Table 2.

The sensitivity of V^* to changes in a parameter b (one of P, H, n_0 and p_{N}) may be obtained if we assume that the optimal values of the decision variables lie in the interior of the domain, satisfying the relevant first order necessary conditions. By the envelope theorem:

$$\frac{\mathrm{d}V^*}{\mathrm{d}b} = \frac{\partial V(a^*)}{\partial b},$$

where $a^* = Q^*$ for case I of section 3.4, T^* for case II, r^* for case III and (T^*, r^*) for case IV.

Comparative statics of the optimal values of Q^*, T^* and r^* with respect to b may be obtained by applying the implicit function theorem to the relevant first order necessary conditions evaluated at the optimal values of the decision variables. For example, for a Case IV problem:

$$\frac{\mathrm{d}T^*}{\mathrm{d}b} = \frac{\begin{vmatrix} -V_{Tb} & V_{Tr} \\ -V_{rb} & V_{rr} \end{vmatrix}}{|\mathbf{H}|}, \tag{14}$$

where the denominator of (14) is the determinant of the Hessian for the problem, $\mathbf{H}$, and is strictly negative at an interior solution (a similar expression applies for the partial derivative of r^* with respect to b). Equation (14) simplifies to $\mathrm{d}a^*/\mathrm{d}b = -[\partial^2 V(a^*)/\partial a \partial b][\partial^2 V(a^*)/\partial a^2]^{-1}$ for Cases I–III. The denominator of each expression is strictly negative at an interior solution, so the problem of signing the derivative of interest reduces to one of signing the numerator. In Appendix C.3, we illustrate the algebra that leads to the results presented in this section.

Results for $\mathrm{d}V^*/\mathrm{d}b$ are the same across cases I–IV: the maximised expected net gain of the trial is strictly increasing in P and H and is strictly decreasing in n_0. It can be positive, zero or negative for p_{N}, according to values taken by other parameters of the model. This latter result is not surprising: an increase in p_{N} increases the size of the population that benefits from adopting S but, at the same time, decreases the size of the population that benefits from adopting N.

Table 2 Comparative statics results for the value function and decision variables, as a function of cases I-IV in section 3.4.

Parameter b	Value function $\mathrm{d}V^*/\mathrm{d}b$ Cases I–IV	Optimal choice variables $\mathrm{d}a^*/\mathrm{d}b$ Cases I–III	Case IV
P	≥ 0	$\lessgtr 0$[1]	$\lessgtr 0$
H	≥ 0	$\lessgtr 0$[1]	$\lessgtr 0$
n_0	≤ 0	$\lessgtr 0$[2]	$\lessgtr 0$
p_{N}	$\lessgtr 0$	$\lessgtr 0$[3]	$\lessgtr 0$

[1]Strictly positive when $\alpha_{\mathrm{N}}(T) = \alpha_{\mathrm{S}}(T) = 0$.
[2]Strictly negative when $n_0 > rT^*/4$.
[3]$\partial a^*/\partial p_{\mathrm{N}} = 0$ when $I_{\mathrm{N}} = I_{\mathrm{S}} = \mu_0 = 0$, or $I_{\mathrm{N}} = I_{\mathrm{S}} = 0$ with offline learning and fixed patient pool.

It is not possible to unambiguously sign the expressions for the optimal values of the decision variables in any of the four cases. These results highlight some counterintuitive behaviour of the decision variables for cases I–III in response to changes in the parameters of interest for which we now provide intuition.

Firstly, $\partial a^*/\partial P$ (fixed patient pool) and $\partial a^*/\partial H$ (fixed horizon) may be positive or negative because of two forces operating in opposite directions. Increasing the population affected by the adoption decision forces a^* upwards because each observation is more valuable. However, increasing the population decreases $\alpha_{\mathrm{N}}(T)$ and $\alpha_{\mathrm{S}}(T)$ – the expected per patient switching costs – which makes the adoption decision less costly, forcing a^* downwards. When $\alpha_{\mathrm{N}}(T) = \alpha_{\mathrm{S}}(T) = 0$, the latter force disappears and a^* is increasing in P and H.

Secondly, $\partial a^*/\partial n_0$ is not always negative. When $n_0 \geq r^*T^*/4$, a^* is decreasing with n_0 because a larger effective sample size for the prior distribution requires fewer additional observations to achieve the same degree of precision to make an adoption decision. For $n_0 < r^*T^*/4$, an additional force may change the sign of $\partial a^*/\partial n_0$. Intuitively, an increase in n_0 means that we have more confidence in our prior beliefs. Thus, to produce any changes in the adoption decision requires more observations that contradict our prior beliefs. Mathematically, we show that $\partial^2\sigma_Z/\partial T\partial n_0$ can be negative in the region where $n_0 < r^*T^*/4$.

Finally, although $\partial a^*/\partial p_{\mathrm{N}}$ has no definitive sign, we point out two special cases in which $\partial a^*/\partial p_{\mathrm{N}} = 0$ for any p_{N}. The first is when $I_{\mathrm{N}} = I_{\mathrm{S}} = \mu_0 = 0$; the second is when $I_{\mathrm{N}} = I_{\mathrm{S}} = 0$ with offline learning and fixed patient pool. Notice that the effect of p_{N} on the optimal trial design is highly dependent on I_{N} and I_{S}. We explore this interaction with numerical results in section 4.2.1.

3.6. Asymptotically large $P(T)$

We conclude this section by presenting results as the number of post-trial patients who benefit from the adoption decision, $P(T)$, approaches infinity. Results are of methodological relevance because

they permit us better to understand the effect of some parameters on the solution; they are of practical relevance because trials often target a large population and the approximations can be very accurate, as our application in section 4 shows.

3.6.1. Undiscounted rewards. We study the limiting behavior of V^* and Q^* as $P \to \infty$ for a fixed patient pool and as $H \to \infty$ for a fixed horizon. We study these limits while ignoring the constraints $T \le T_{\max}$ and $r \le r_{\max}$ unless they are necessary for the existence of a solution to (11). Prop. 5 presents this asymptotic behaviour for the case of undiscounted rewards (cases II and IV are necessarily fixed horizon). It assumes that $c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 > 0$; otherwise, the limits often do not exist. It further assumes that $c_{\text{cap}}(r) = c_{\text{fixed}} + c_r r$ where $c_r = 0$ for cases I and II with constant setup costs and where $c_r > 0$ for cases III and IV with variable costs.

PROPOSITION 5. *Assume* $c - \delta_{\text{on}}(1-2p_N)\mu_0/2 > 0$, $\rho = 0$, *and* $c_{cap}(r) = c_{fixed} + c_r r$. *Then, for cases I* ($c_r = 0$) *and III* ($c_r > 0$)

$$\lim_{P\to\infty} \frac{Q^*}{\sqrt{P}} = \left(\frac{\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{4(c - \delta_{\text{on}}(1-2p_N)\mu_0/2)} \right)^{1/2} ; \lim_{P\to\infty} \frac{V^*}{P} = (1-p_N)\sigma_0\Psi(-\mu_0/\sigma_0) + p_N\sigma_0\Psi(\mu_0/\sigma_0).$$

For cases II ($c_r = 0$) *and IV* ($c_r > 0$)

$$\lim_{H\to\infty} \frac{V^*}{\zeta(H-\Delta)} = (1-p_N)\sigma_0\Psi(-\mu_0/\sigma_0) + p_N\sigma_0\Psi(\mu_0/\sigma_0).$$

For case II, where $c_r = 0$,

$$\lim_{H\to\infty} \frac{Q^*}{\sqrt{\zeta(H-\Delta)}} = \left(\frac{\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{4(c - \delta_{\text{on}}(1-2p_N)\mu_0/2) + 4\zeta\sigma_0(\Psi(\mu_0/\sigma_0) + (1-p_N)\mu_0/\sigma_0)/r_{\max}} \right)^{1/2}.$$

For case IV, where $c_r > 0$,

$$\lim_{H\to\infty} \frac{Q^*}{\sqrt{\zeta(H-\Delta)}} = \left(\frac{\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{4(c - \delta_{\text{on}}(1-2p_N)\mu_0/2)} \right)^{1/2}.$$

Appendix C.4 proves Prop. 5. The proofs introduce additional results on the asymptotics of T^* and r^* that are particularly interesting for case IV as both H and incidence ζ grow without bound.

For a finite value of $P(T)$, asymptotic approximations to the optimal value of the decision variable and maximised value function may be obtained by multiplying the right-hand side constants by the relevant denominators in the left-hand side of Prop. 5. Denote with $(\hat{T}, \hat{r})$ the asymptotic approximations. The approximations are accurate when $P(\hat{T})\sigma_0$, which we may think of as the standard deviation of the prior distribution for expected benefits for the post-trial population, is much larger than μ_0, I_{N}, I_{S}, and $c\hat{r}\hat{T}$, and $\hat{r}\hat{T}$ is much larger than n_0.

Prop. 5 provides three additional insights. First, the optimal sample size increases in the limit of large $P(T)$ as the square root of P (fixed patient pool) and the square root H (fixed horizon).

Second, both the fixed patient pool and fixed horizon models attain the same expected net gain in the limit. Finally, the switching costs I_{N} and I_{S} do not appear in these propositions and can be ignored if $P(T)$ is large enough.

3.6.2. Discounted rewards. We now turn to the case of a positive discount rate, $\rho > 0$. This is only applicable to cases II and IV (because cases I and III require undiscounted rewards). We denote with subscripts the adoption decision type. For instance, $V_{\mathrm{H}}(T,r)$ refers to the expected net gain with fixed horizon, while $V_{\mathrm{P}}(T,r)$ refers to the expected net gain with fixed patient pool. Unlike the undiscounted rewards asymptotic results, V^*, T^*, r^*, and Q^* do not diverge in the limit of large $P(T)$. Therefore, the results here take a completely different form. Results are proven in Appendix C.5.

It is easy to check that $\lim_{P\to\infty} P_\rho = \zeta/\rho$ with fixed patient pool, and $\lim_{H\to\infty} P_\rho(T) = \zeta/\rho$ with fixed horizon. Define similarly $\alpha'_{\mathrm{N}} = \lim_{P\to\infty}\alpha_{\mathrm{N}} = \lim_{H\to\infty}\alpha_{\mathrm{N}}(T) = I_{\mathrm{N}}\rho/((1-p_{\mathrm{N}})\zeta)$, and $\alpha'_{\mathrm{S}} = \lim_{P\to\infty}\alpha_{\mathrm{S}} = \lim_{H\to\infty}\alpha_{\mathrm{S}}(T) = I_{\mathrm{S}}\rho/(p_{\mathrm{N}}\zeta)$. Prop. 6 states that $V(T,r)$ converges to the following expression with both fixed patient pool and fixed horizon, which is obtained by substituting $P_\rho(T)$ with ζ/ρ, $\alpha_{\mathrm{N}}(T)$ with α'_{N}, and $\alpha_{\mathrm{S}}(T)$ with α'_{S} in (9) and (10):

$$V_\infty(T,r) = \begin{cases} -(c_{\mathrm{cap}}(r) + cr\tilde{T}_\rho(T)) + \delta_{\mathrm{on}}(r\tilde{T}_\rho(T)/2)(1-2p_{\mathrm{N}})\mu_0 \\ \qquad + \frac{\zeta e^{-\rho(T+\Delta)}\sigma_Z}{\rho}\left[(1-p_{\mathrm{N}})\Psi\left(\frac{\alpha'_{\mathrm{N}}-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\Psi\left(\frac{\alpha'_{\mathrm{S}}+\mu_0}{\sigma_Z}\right)\right], & \text{if } rT > 0 \\ \max\{0,\ (1-p_{\mathrm{N}})\zeta\mu_0/\rho - I_{\mathrm{N}},\ -p_{\mathrm{N}}\zeta\mu_0/\rho - I_{\mathrm{S}}\}, & \text{if } rT = 0. \end{cases} \tag{15}$$

Proposition 6. *If $\rho > 0$, then $V_P(T,r)$ (as $P\to\infty$) and $V_H(T,r)$ (as $H\to\infty$) converge uniformly to $V_\infty(T,r)$ on the compact domain $\{(T,r) : 0 \le T \le T_{\max}, 0 \le r \le r_{\max}\}$ and*

$$\lim_{P\to\infty} V_P^* = \lim_{H\to\infty} V_H^* = \max_{\substack{0\le T\le T_{\max}\\ 0\le r\le r_{\max}}} V_\infty(T,r).$$

In addition, if $(T_\infty, r_\infty) = \arg\max_{T,r} V_\infty(T,r)$ is unique, then $\lim_{P\to\infty} T_P^* = \lim_{H\to\infty} T_H^* = T_\infty$, $\lim_{P\to\infty} r_P^* = \lim_{H\to\infty} r_H^* = r_\infty$, and $\lim_{P\to\infty} Q_P^* = \lim_{H\to\infty} Q_H^* = r_\infty T_\infty/2$.

In summary, for finite but large $P(T)$, we can approximate (11) with

$$\begin{aligned} &\max_{T,r} V_\infty(T,r) \\ &\text{s.t. } T \in [0, T_{\max}],\ r \in [0, r_{\max}]. \end{aligned} \tag{16}$$

The function $V_\infty(T,r)$ is equivalent to the discounted fixed patient pool model with $P_\rho = \zeta/\rho$. Thus, all the previous results related to discounted fixed patient pool are also valid when solving (16). Both the fixed patient pool and fixed horizon models converge to the same function and maximisers. However, we find in numerical examples in section 4.2.1 that fixed patient pool converges faster than fixed horizon.

4. Application to the ProFHER pragmatic trial

We apply our model to data from the ProFHER pragmatic trial (Rangan et al. 2015, Handoll et al. 2015, Corbacho et al. 2016) using a series of numerical experiments. The application is illustrative and is not intended to advocate for a given technology decision. The ProFHER trial was a multicentre, randomised clinical trial conducted in the UK National Health Service (NHS) which investigated the use of surgery versus nonsurgical intervention (sling) to treat patients with a displaced proximal humeral fracture. Over a period of approximately two and a half years, 250 patients across 32 NHS hospitals were randomised, on an equal basis, to the two arms of the trial. Follow-up of both the primary endpoint (the Oxford Shoulder Score) and the cost-effectiveness endpoints (QALYs for health outcomes, using the EQ-5D-3L questionnaire, together with treatment and rehabilitation costs) took place after six, twelve, and 24 months. Results suggested there was no difference between surgery and sling in terms of effectiveness, but that surgery cost, on average, approximately £1,800 more than sling. Surgery was therefore deemed to be neither more effective, nor more cost-effective, than sling.

We assess the performance of the model using expected net gain, the optimal recruitment rate, r^*, the duration of the optimal recruitment period, T^*, and the optimal number of pairwise observations, $Q^* = r^*T^*/2$. We consider how these metrics are affected by regulatory concerns such as the post-adoption population and the discount factor. We also consider two probabilistic measures. The first is based on Bayesian principles and we refer to it as the 'conditional probability of correct selection' (CPCS). CPCS is defined as the probability of adopting the correct technology, given a specific value of W. Let $\mathcal{D}^{orac}$ be the 'oracle' adoption decision that selects the technology with the highest benefits for the patients post-trial, knowing the true value of W *a priori*. Then $\text{CPCS}(w) = \Pr(\mathcal{D}^* = \mathcal{D}^{orac} \,|\, W = w)$. The probability of correct selection, $\text{PCS} = \mathbb{E}\,[\,\text{CPCS}(W)\,]$, is a more commonly used measure, but we use CPCS for ease of comparison with the next measure.

The second probabilistic measure, which we call 'power', is akin to the frequentist concept of the power of a hypothesis test, the probability of rejecting the null hypothesis that $W = 0$, given $W = w$, in a two-tailed test at the 5% significance level. We plot power curves which show the probability that a 95% CI for the unknown mean does not contain zero, as a function of w.

CPCS and power quantify the probability of correctly adopting a technology, but they differ in two respects. Firstly, CPCS uses prior information, while power considers only the samples collected during the trial. Secondly, CPCS makes use of the optimal adoption decision according to our models, while the power calculations assume a rejection region that guarantees a type I error probability. Further details of the computation of CPCS and power are given in Appendix E.

4.1. Parameter values

The analysis presented here is based primarily on the parameter estimates reported in Forster et al. (2019, Table 1), who estimated parameters for the ProFHER trial by referencing its main publications and consulting with coauthors from the trial. Define surgery as technology N and sling as technology S. We consider the case of offline learning ($\delta_{\text{on}} = 0$) and set $p_{\text{N}} = 0.39$, $\mu_0 = £0$, $n_0 = 2$, and $\sigma_X = £4,400$. We consider the cases of both fixed patient pool and fixed horizon, with parameters $H = 15$ years, $\Delta = 1$ year, $\zeta = 7,000/$ year, so that $P = \zeta H = 105,000$. Our choice of $H = 15$ differs from the choice $H = 6$ in Forster et al. (2019) and reflects the fact that gains from an adoption decision are likely acquired beyond a period of 6 years.

Following Forster et al. (2019), we set $I_{\text{N}} = £0$, $I_{\text{S}} = £0$, $c = £2,040$ and fix the recruitment rate to be $\check{r} = 94$ patients /year. Forster et al. (2019) did not model site-specific setup costs, which we handle as follows. We assume a setup cost function of the form: $c_{\text{cap}}(r) = c_{\text{fixed}} + c_r r$. We assume that half of the setup costs of the trial were fixed and the other half represent marginal costs, obtaining the following estimates: $c_{\text{fixed}} = c_{\text{cap}}(\check{r})/2 = £480,000$; $c_r = c_{\text{cap}}(\check{r})/2/\check{r} = £5,080$ per patient per year. Finally, we let $\rho = 3.44\%$/year, equivalent to the 3.5% annual discount rate recommended by NICE (2013) and used in Corbacho et al. (2016). We refer to analysis based on these parameter values as the 'base case' analysis.

4.2. Numerical experiments

In section 4.2.1, we fix the recruitment rate at $\check{r} = 94$ patients / year and we optimise the recruitment duration, considering both undiscounted and discounted rewards. In section 4.2.2, we optimise both the recruitment rate and duration. For clarity of exposition, we use subscripts to denote the adoption decision type. For instance, V^*_{H} and V^*_{P} refer to the maximum expected net gain for fixed horizon and fixed patient pool, respectively.

Inspection of (9) shows that, when $\mu_0 = 0$, online and offline learning have the same optimal one-shot design. Hence, given our assumption that $\mu_0 = 0$, results presented here are equally applicable to the cases of online and offline learning. The optimal design can differ for online and offline when $\mu_0 \neq 0$ or when a trial is sequential (see section 5).

4.2.1. Fixed recruitment rate ($\check{r} = 94$/year), optimal choice of recruitment period duration, T. Figure 3 shows the expected net gain as a function of the duration of the recruitment period, T, measured in months, when the recruitment rate is fixed (left subfigure: fixed horizon; right subfigure: fixed patient pool). The base case is represented by the dashed line and the undiscounted version of the base case by the continuous line. T^* is marked with a circle or square. The duration of the ProFHER trial itself is marked with a vertical dotted line.

Figure 3 **The expected net gain versus the recruitment period duration for the ProFHER application with fixed horizon (left), and fixed patient pool (right).**

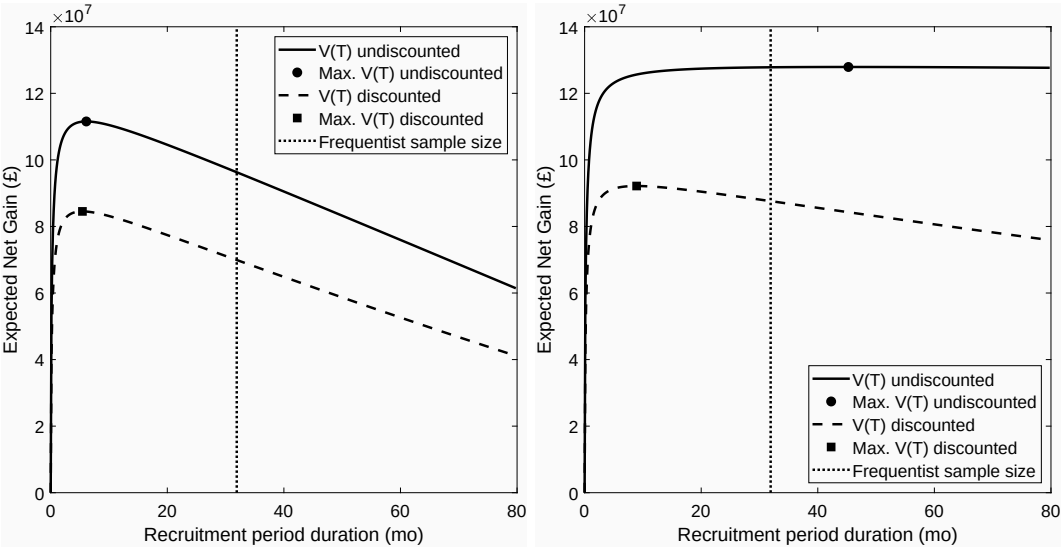

We point out two important observations. Firstly, both the maximum expected net gain of the trial and the optimal recruitment period duration are highly dependent on the adoption decision type. In this example, the expected net gain is approximately 10% higher for fixed patient pool and the optimal duration of the recruitment period is over 1.6 times larger than for fixed horizon. Secondly, the expected net gain function is more sensitive to changes in T in the fixed horizon case than the fixed patient pool case: extending the length of the trial by recruiting more patients reduces the size of the population to benefit in the fixed horizon case, but not in the fixed patient pool case. This result holds more strongly when rewards are not discounted.

Sensitivity to the size of the post-trial population, $P(T)$. Figure 4 shows how the optimal number of pairwise allocations increases as a function of H (left subfigure, for the fixed horizon case) and P (right subfigure, for the fixed patient pool case). Plotted with dotted lines are the asymptotic approximations that were derived in section 3.6. Horizontal dotted lines correspond to the approximations with discounted rewards, increasing dotted lines to those without discounted rewards. The approximations with undiscounted rewards are close to the actual values in the range plotted, with a better fit for fixed patient pool. This is due to $I_{\mathrm{N}} = I_{\mathrm{S}} = \mu_0 = 0$, and $P(T^*)\sigma_0 \gg cr$. The approximations for discounted rewards are accurate for fixed patient pool, but the range of values of H for fixed horizon is not large enough.

Figure 5 presents the CPCS and power curves for the optimal recruitment duration in the base case and a version in which the post-trial population (H and P) is doubled. The vertical lines represent the smallest relevant difference for the frequentist sample size calculations.[4] Because

[4] The sample size of the ProFHER trial was based on setting a type I error probability of $\alpha = 0.05$, power of 0.8, and a specified smallest clinically relevant difference for the primary outcome. The number of pairwise allocations is

Figure 4 **The optimal number of pairwise allocations increases with H for fixed horizon (left) and with P for fixed patient pool (right) in the ProFHER application.**

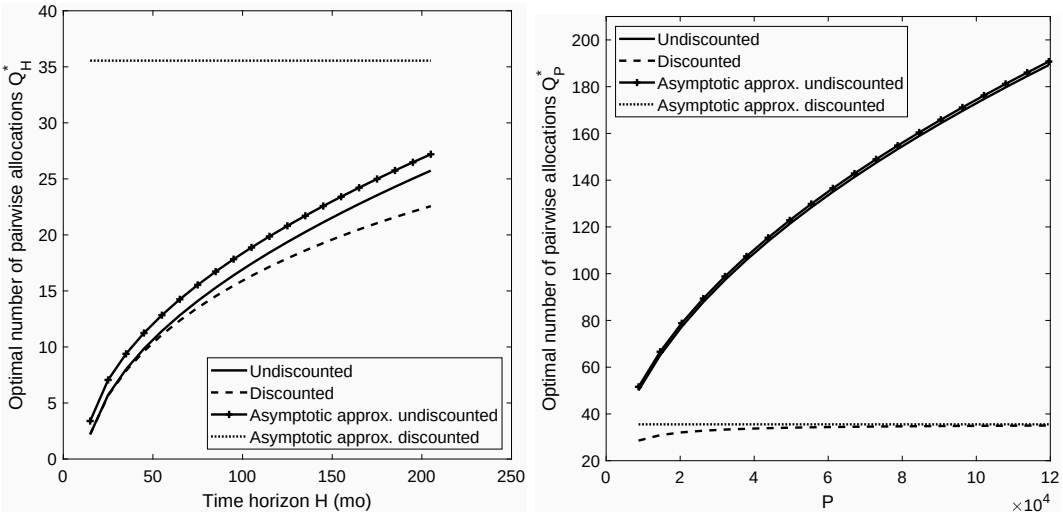

$I_N = I_S = £0$ (so that $\alpha_N(T) = \alpha_S(T) = 0$), it is almost never optimal to choose $\mathcal{D}^* = \text{M}$. This means that the CPCS has a kink at $w = 0$, whereas it will have a kink at $w = \alpha_N(T)$ and $w = -\alpha_S(T)$ when $\alpha_N(T) \neq 0$ and $\alpha_S(T) \neq 0$. Both CPCS and power are higher for a fixed patient pool compared to fixed horizon in this example.

Figure 5 shows that, for the application presented here, at the optimal T, the CPCS beyond the smallest relevant difference is high (i.e. it is above 0.87 for fixed horizon and above 0.93 for fixed patient pool). According to the CPCS, making the correct selection with high probability for the given smallest relevant difference is worth the expenditure in a sufficiently long trial. The power plot shows a completely different story: at the optimal T, the power for the smallest relevant difference is below 0.22 for fixed horizon, and below 0.32 for fixed patient pool. This illustrates why the frequentist decision rule based on the type I error probability of 0.05 can lead health technology adoption decisions which do not maximise value, as argued by Claxton (1999): maximising the trial's value, as defined in (11), does not necessarily gather enough information to satisfy widely acknowledged standards in trial design.[5] It might also gather more than enough.

Sensitivity to the fraction of patients on the new technology, p_N. When $I_N = I_S = \mu_0 = 0$, p_N has no effect on the expected net gain or optimal trial length (see section 3.5). Here we consider the

given by $n = \sigma_X^2 (q_{\alpha/2} + q_\beta)^2 / \delta^2$, where β is the type II error probability, δ is the smallest relevant difference, and q_x is the $1-x$ quantile of a standard normal random variable (Lachin 1981). The smallest relevant difference in terms of INMB that equates the sample size calculations of the actual trial is £1105.

[5] The difference in the CPCS and power when the time horizon is doubled illustrates another interesting feature of our model. For fixed horizon, the increase in the number of pairwise allocations is 13 (30%). This represents a large increase in pairwise allocations, but the maximum increase in the CPCS is only about 3%. This emphasises the point that, although the increase in the CPCS is not very large, it represents a large improvement in the expected benefits to patients.

Figure 5 **CPCS (left) and power (right) curves for the ProFHER application with fixed horizon and fixed patient pool.**

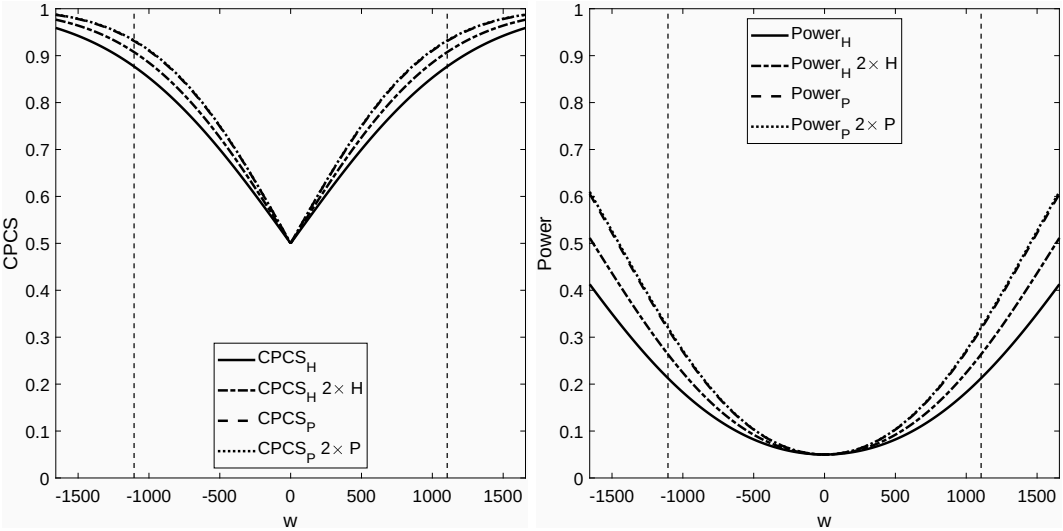

Figure 6 **The effect of p_{N} on the optimal expected net gain and number of pairwise allocations is small for low switching costs in the ProFHER application.**

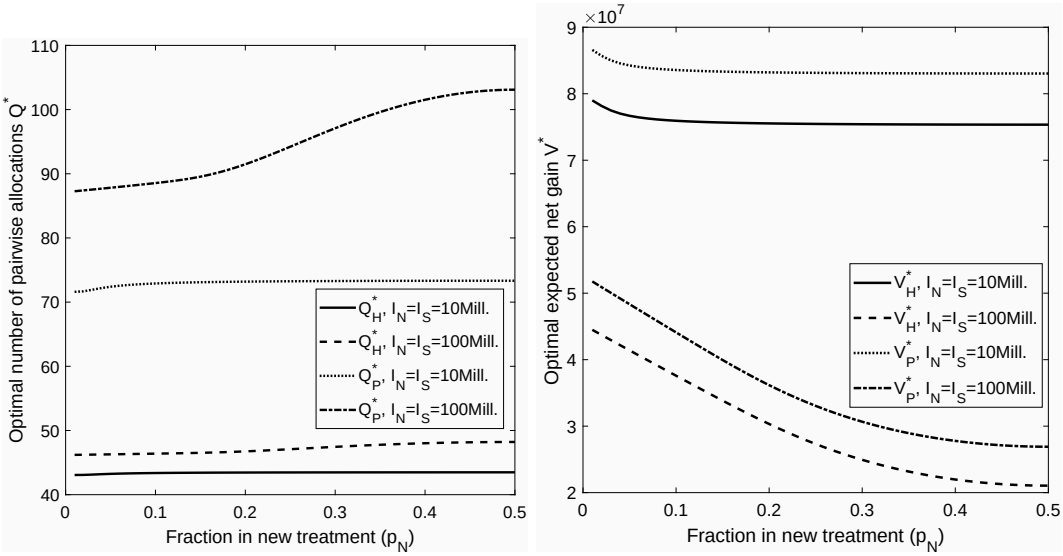

cases of $I_{\mathrm{N}} = I_{\mathrm{S}} = £10\text{Mill}$ and $I_{\mathrm{N}} = I_{\mathrm{S}} = £100\text{Mill}$. Figure 6 shows that, for the fixed horizon, changes in p_{N} have a negligible effect on Q^* for $I_{\mathrm{N}} = I_{\mathrm{S}} = £10\text{Mill}$ and they have a small effect for $I_{\mathrm{N}} = I_{\mathrm{S}} = £100\text{Mill}$. V^* is more sensitive to changes in p_{N}, especially at small values, and more so when the switching costs are high. For fixed patient pool, Q^* is sensitive to p_{N} for large switching costs. V^* is sensitive to p_{N}, especially at small values.

4.2.2. Optimal choice of recruitment period duration, T and recruitment rate, r. We relax the assumption of a fixed recruitment rate and optimise over both T and r. We do not present the case of undiscounted fixed patient pool because it requires optimization of a single

Figure 7 **Marginal value and cost of the recruitment rate showing the optimal levels for three versions of the model.**

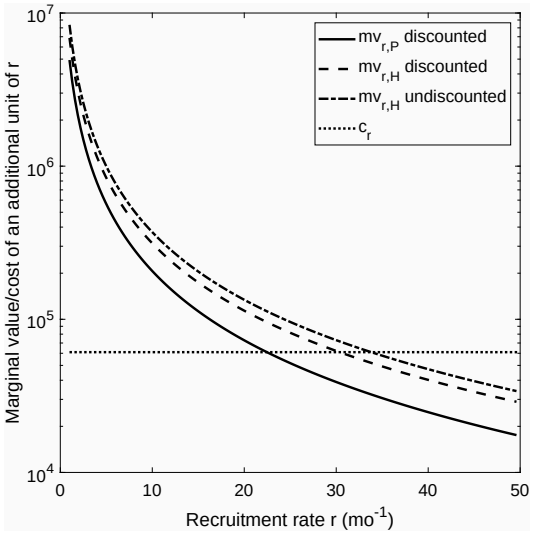

decision variable (it falls under case III in section 3.4). Figure 7 plots the marginal value of the trial's recruitment rate, mv_r, computed as the derivative of $V(T^*, r)$ excluding setup costs with respect to r, at the optimal choice of trial duration, T^*, together with the constant marginal cost, c_r, as a function of r. The intersection of marginal value and marginal cost functions corresponds to the optimal recruitment rate: 22.6 patients /month for fixed patient pool, 30.5 for fixed horizon, and 34.0 for undiscounted fixed horizon. This corresponds to between three and five times the recruitment rate used in the ProFHER trial itself ($\check{r} = 7.8$ patients /month). The corresponding optimal recruitment period durations, T^*, are 5.1 months (57 patient pairs), 2.7 months (42 patient pairs), and 2.9 months (50 patient pairs).

Figure 7 can assist a trial manager with the decision about whether or not to open an additional site. At the actual recruitment rate of the ProFHER trial ($\check{r} = 7.8$ patients /month), Figure 7 shows that the marginal value of increasing r exceeds the marginal cost by £440,000 per additional recruit per month for fixed horizon and £292,000 for fixed patient pool. A trial manager charged with maximising the value of the trial can think of these marginal values as representing the maximum willingness to pay (WTP) for an additional unit of recruitment per month and compare them with the cost of opening an additional site. It is optimal to open an additional site if the cost of doing so is lower than the WTP.

Sensitivity to changing the trial horizon. For ease of exposition, we fix $P = \zeta H$ and analyse the effect of a change in H for both fixed horizon and fixed patient pool. Figure 8 presents the results for the optimal number of pairwise allocations, recruitment rate, and recruitment duration, along with the appropriate asymptotic approximations derived in section 3.6. Similar to Figure 4, the

Figure 8 **Optimal number of pairwise allocations (left), recruitment rate (middle), and recruitment duration (right) as a function of the time horizon.**

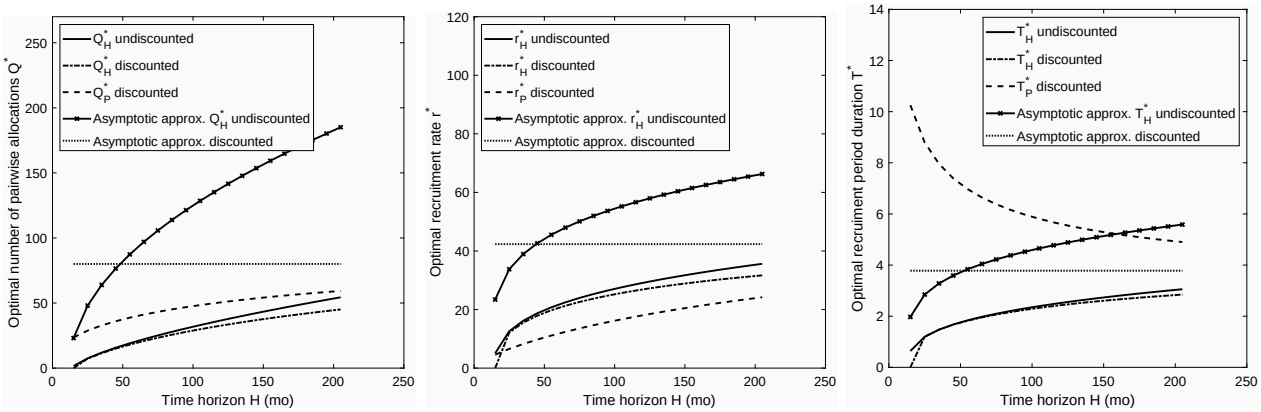

optimal number of pairwise allocations is increasing in H. The more interesting result is how the additional patients are obtained in terms of T and r. The optimal recruitment rate is increasing for both fixed horizon and fixed patient pool. However, the optimal recruitment duration is increasing for fixed horizon but decreasing for fixed patient pool: T^*_{H} approaches the asymptote from below, while T^*_{P} approaches it from above.[6]

4.3. Optimizing the recruitment rate for convex increasing setup costs

The analysis so far has assumed that the trial's setup cost function is linear in the recruitment rate. In practice, it is likely that decreasing returns to scale operate: opening new sites is likely to become increasingly difficult and the recruitment rate of new sites will be smaller if more productive ones are opened earlier. While the cost function may not be known a priori, a reasonable estimate of it may be made based on site-specific factors such past experience with setup costs, clinical investigators, incidence rates and so on.

Using the recruitment data from the ProFHER study, together with some additional assumptions, we estimate the setup cost function, $c_{\mathrm{cap}}(r)$, and compute the optimal trial design. Handoll et al. (2015, Fig. 5, page 37) report the number of patients recruited at each site. For illustrative purposes, we assume that the cost of opening each site has the same cost, but that the recruitment rates of the sites vary.

Given the assumptions made for the application, the cost of opening each of the 35 sites was about £13,700.[7] Assuming the sites are opened from the most productive (highest recruitment

[6] The asymptotic approximation for the case of fixed horizon, undiscounted rewards, fits the actual optimal values poorly because the marginal cost of sampling, c, is much smaller than the marginal cost recruitment rate, c_r. In fact, the approximation that assumes $c - \delta_{\mathrm{on}}(1-2p_{\mathrm{N}})\mu_0/2 = 0$ in Lemma EC.4 of the appendix fits the optimal values much better in the range of values considered here.

[7] The cost of opening sites for oncology trials is estimated at $20,000-$30,000 (Fassbender 2016).

Figure 9 **Cost of recruitment rate assuming a constant variable cost per site (left) and the corresponding marginal cost of recruitment rate with the WTP for an additional unit of recruitment (right)**

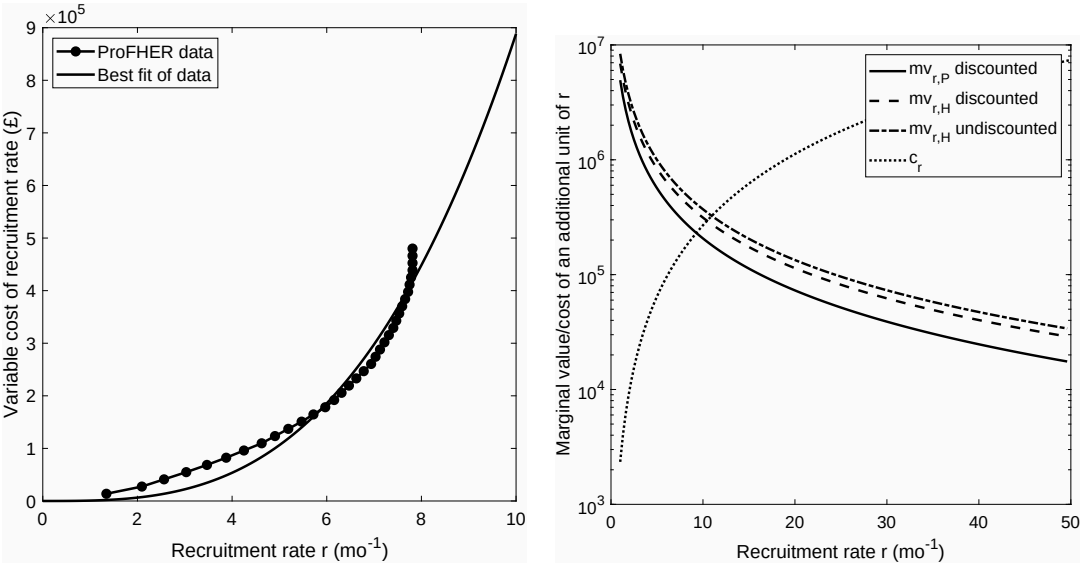

rate) to least productive, the left panel of Figure 9 presents the variable cost of recruitment rate with the best fit (minimum mean squared error) function of the form $f(x) = \alpha x^{\beta}$. The estimated setup cost function is $c_{\text{cap}}(r) = 480,000 + 766r^{3.06}$, so that the marginal cost of a unit increase in recruitment is proportional to $r^{2.06}$.

The right panel of Figure 9 plots this marginal cost function against the marginal value functions. The optimal recruitment rates are 9.3/month for fixed patient pool, 10.5/month for discounted fixed horizon, and 11.0/month for undiscounted fixed horizon. The optimal recruitment rates are higher than the actual recruitment rate. However, compared with the results in 4.2.2, the difference between the optimal recruitment rate and the actual trial is much smaller. The corresponding optimal durations are 8.1 months (38 patient pairs), 4.7 months (25 patient pairs), and 5.2 months (28 patient pairs), which are much smaller than the 125 patient pairs enrolled in the actual trial.

5. Analysis of response adaptive extensions

The one-shot design does not permit interim analysis of the data as it accumulates. Here we consider two extensions which enable the trial to be extended, or shortened, according to interim evidence. We consider how to adapt fully the duration of the recruitment period, holding the recruitment rate fixed. We also consider how to adapt both the recruitment rate and recruitment period duration, taking into account additional costs that such a design may incur.

5.1. Fully response adaptive trial with a fixed recruitment rate

For some special cases, parameter transformation may be used to obtain a version of the one-shot design which fits the model of sequential experimentation proposed by Chick et al. (2017), a design

which assumed $p_{\mathrm{N}} = 0$. We fix the recruitment rate and define decision epochs $t = 0, \ldots, Q_{\max} - 1$. At each decision epoch, measurements of health outcome and treatment cost are observed (with appropriate delay) for a pair of patients randomised to each arm of the trial. An action $a_t \in \{0, 1\}$ is made to continue sampling ($a_t = 1$) or stop sampling ($a_t = 0$). Once recruitment has stopped, we observe the outcomes of patients in the 'pipeline' – those who have been randomised, but whose outcomes have yet to be observed – prior to making the adoption decision. A policy, π, maps available knowledge (prior information, plus acquired data) to an action at each decision epoch. Let $\mathbb{E}_\pi[\cdot]$ denote the expectation induced by policy π. If $a_0 = 0$, the trial does not run and the adoption decision is made immediately, based on the prior information alone.

The model in section 2 assumes that T and r are continuous, whereas the decision epochs are discrete. We transform the model into a discrete time version by defining the following parameters: $\tau = \lceil r\Delta/2 \rceil$ is the number of patient pairs enrolled in the trial during the follow-up period of length Δ; $\tilde{\rho} = e^{2\rho/r} - 1$ is the discrete discount rate per patient pair. Given this setup, $Q \in \{0, 1, \ldots, Q_{\max}\}$ and the expected net gain for the sequential trial is:

$$\begin{aligned} V(\pi) = -\mathbf{1}_{Q>0} c_{\mathrm{cap}}(r) + \mathbb{E}_\pi \left[\sum_{t=0}^{Q-1} \frac{\delta_{\mathrm{on}}(1 - 2p_{\mathrm{N}}) X_{t+1} - 2c}{(1+\tilde{\rho})^t} \right] \\ + \mathbb{E}_\pi \left[\frac{\mathbf{1}_{\mathcal{D}=\mathrm{N}}((1-p_{\mathrm{N}}) P_\rho(2Q/r) W - I_{\mathrm{N}}) + \mathbf{1}_{\mathcal{D}=\mathrm{S}}(-p_{\mathrm{N}} P_\rho(2Q/r) W - I_{\mathrm{S}})}{(1+\tilde{\rho})^{\mathbf{1}_{Q>0}(Q+\tau)}} \right]. \end{aligned} \tag{17}$$

An optimal sequential trial design is a policy π^* that maximises $V(\pi)$.

If we assume a fixed patient pool ($P(T) = P$), we may identify four special cases which permit this problem to be stated and solved using the methods of Chick et al. (2017): (1) current practice is one treatment only ($p_{\mathrm{N}} = 0$); (2) rewards are undiscounted and one of the two technologies must be adopted ($\tilde{\rho} = 0$ and $\mathcal{D} \in \{\mathrm{N}, \mathrm{S}\}$); (3) rewards are undiscounted and there are no switching costs ($\tilde{\rho} = 0,\ I_N = I_S = 0$); (4) there are no switching costs and the prior mean is equal to zero, but rewards need not be discounted ($\tilde{\rho} \geq 0,\ I_N = I_S = \mu_0 = 0$). Appendix D provides full details. It also corrects a misstatement in Alban et al. (2018), who presented preliminary analysis of these cases. The scenarios which enable fully sequential trials here are aesthetically linked to the two special cases for comparative statics noted in the last paragraph of section 3.5.

5.2. Response adaptive recruitment rate and duration

There may be added value in making the recruitment rate a 'response adaptive' decision variable. For example, for fixed horizon trials which report promising initial data, it might be beneficial to learn more quickly by increasing the rate at which patients enter the trial. Changing the recruitment rate mid-trial might be costly, however. For example, it might entail additional work to open new sites or train up personnel. We present a block sequential approach, in which a fixed and finite number of decision points (interim analyses) are available for adapting the recruitment rate.

Figure 10 **Timeline of decisions and information.**

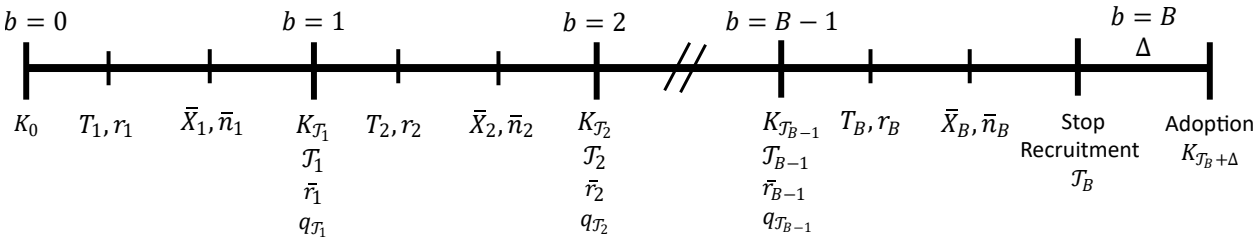

5.2.1. Decision variables and dynamic programming formulation. Let the trial proceed in a sequence of up to B stages, each of which commences with an interim look at the data. B is an exogenous parameter whose value is fixed prior to the start of the trial. At the beginning of each stage, a recruitment rate and duration are selected for the next batch of patients, as a function of the information accumulated. This process is repeated until the Bth stage, where the data from the pipeline patients is observed after the usual delay, and a technology is adopted.

Figure 10 shows a timeline for the trial. Let $b = 0, 1, 2, \ldots, B-1$ index decision points, where the batch duration, T_{b+1}, and batch recruitment rate, r_{b+1}, for the $b+1$st batch are selected on the basis of information obtained thus far. The number of patient pairs recruited in batch $b+1$ is $T_{b+1}r_{b+1}/2$. If $T_{b+1} = 0$, the trial stops, so that T_i, r_i for $i > b+1$ are set to 0.

At the end of the bth batch, for $b = 1, 2, \ldots, B$, the elapsed time is $\mathcal{T}_b = \sum_{i=1}^{b} T_i$, and the maximum rate of recruitment to that time is $\bar{r}_b = \max\{r_i : i = 1, 2, \ldots, b\}$. Define $\mathcal{T}_0 = 0$, $\bar{r}_0 = 0$, and let $K_0 = (\mu_0, n_0)$ denote the prior information set. The information set $K_{\mathcal{T}_b} = (\mu_{\mathcal{T}_b}, n_{\mathcal{T}_b})$ at the end of the bth batch is computed as a function of the information set at the prior batch $K_{\mathcal{T}_{b-1}}$ and the sample average of the $\bar{n}_b$ observations, $\bar{X}_b$, observed during the bth batch. When recruitment ends, the adoption decision is taken at time $\mathcal{T}_B + \Delta$, with information set $K_{\mathcal{T}_B+\Delta}$.

To keep track of the number of patients in the pipeline, define $q_{\mathcal{T}_b}(T)$ as the number of patient pairs in the pipeline at time $\mathcal{T}_b$, whose outcomes will be observed within duration T. The total number of patients in the pipeline is thus $q_{\mathcal{T}_b}(\Delta)$. The function $q_{\mathcal{T}_b}$ is a piecewise linear function of the recruitment rates and durations of batches to that time.

We assume the trial incurs the cost $(c_{\text{cap}}(r_{b+1}) - c_{\text{cap}}(\bar{r}_{b-1}))^+$ at decision point b for selecting recruitment rate r_{b+1}. This assumes that a non-zero cost is incurred if and only if the recruitment rate reaches a new maximum.

Define $V_b(K, \mathcal{T}, \bar{r}, q)$, $b = 0, 1, \ldots, B$, as the value-to-go functions for the information set K, time since the start of the trial $\mathcal{T}$, maximum recruitment rate to date $\bar{r}$, and number of patient pairs in the pipeline q. The terminal reward is the expected health outcomes for the post-adoption population. The dynamic programming formulation of our block sequential trial problem is

$$V_B(K, \mathcal{T}, \bar{r}, q) = e^{-\rho\Delta \mathbf{1}_{\mathcal{T}>0}} \mathbb{E}\left[\mathbf{1}_{\mathcal{D}=\text{N}}((1-p_{\text{N}})P_\rho(\mathcal{T})W - I_{\text{N}}) + \mathbf{1}_{\mathcal{D}=\text{S}}(-p_{\text{N}}P_\rho(\mathcal{T})W - I_{\text{S}}) \middle| K \right],$$

$$V_b(K,\mathcal{T},\bar{r},q) = \max_{\substack{0\le T\le T_{\max}-\mathcal{T}\\ 0\le r\le r_{\max}}} -(c_{\text{cap}}(r)-c_{\text{cap}}(\bar{r}))^+ - c\tilde{T}_\rho(T)r + \delta_{\text{on}}(1-2p_{\text{N}})\tilde{T}_\rho(T)r\mathbb{E}\left[W \mid K\right]/2 \\ + e^{-\rho T}\mathbb{E}\left[V_{b+1}(U(K,\bar{X},\bar{n}_{b+1}),\mathcal{T}+T,\max\{\bar{r},r\},\tilde{U}(q,T,r))\middle| K\right], \tag{18}$$

where the expectation is taken over the variables W and $\bar{X}$. Recall that $\bar{X}_{b+1} \mid W \sim \mathcal{N}(W, \sigma_X^2/\bar{n}_{b+1})$ and $W \mid K_{\mathcal{T}_b} \sim \mathcal{N}(\mu_{\mathcal{T}_b}, \sigma_X^2/n_{\mathcal{T}_b})$, where $n_{\mathcal{T}_b}$ is the effective number of samples that have been observed until $\mathcal{T}_b$, and $\bar{n}_{b+1}$ is the number of samples observed in batch $b+1$ and is computable from information at time $\mathcal{T}_b$ as well as T and r. Thus, $n_{\mathcal{T}_{b+1}} = n_{\mathcal{T}_b} + \bar{n}_{b+1}$.

The terminal reward function, V_B, is discounted only if the trial runs, and accounts for the optimal adoption decision after samples have arrived with delay. The function U maps a prior distribution $K_{\mathcal{T}_b}$, sample mean $\bar{X}_{b+1}$ and number of samples $\bar{n}_{b+1}$ to a posterior distribution using Bayes' rule. The function $\tilde{U}$ outputs $q_{\mathcal{T}_{b+1}}$ given $q_{\mathcal{T}_b}$, batch $b+1$ sampling rate r_{b+1} and duration T_{b+1}. Through backward induction in (18) we can obtain the expected net gain of the trial $V_0(K_0,0,0,0)$. The state space of this model is large and obtaining an exact solution using backward induction is difficult owing to the curse of dimensionality. In Appendix D.2 we propose a forward-looking heuristic to find a policy with large expected net gain.

6. Discussion and conclusion

Innovating the design of clinical research has high priority across many jurisdictions. In the United States, the NIH's recent initiative to increase efficiency, accountability and transparency in clinical trials covers a wide range of areas, running from improved training for investigators through to reducing delays (Hudson et al. 2016). A similar exercise is being undertaken in Europe, where the European Union will introduce the 'Clinical Trial Regulation' (European Union 2014), which seeks to harmonise the process of assessment and supervision of trials across the nation states, with the aim of increasing their transparency and efficiency. And in the United Kingdom, the NIHR is funding a range of studies which have novel methodological designs that aim to deliver research results with greater efficiency. In part, these initiatives appear designed to increase technical efficiencies in the way trials operate under the existing frequentist paradigm. However, echoing Claxton and Posnett (1996) and those who have followed, the current drive for more 'value-based' health care also raises questions about whether there exists a role for more 'value-based' clinical research in the health technology assessments of the future.

Our literature survey has shown how, complementing these initiatives to innovate clinical trial design, Bayesian decision-theoretic models have grown to challenge traditional frequentist approaches. Our model has contributed to this literature in a number of ways. To begin with, it can handle scenarios in which clinical practice is mixed prior to the commencement of the trial, a common occurrence in many areas of clinical research. Trial teams considering designing more

value-based trials should be aware that mixed practice can affect optimal decisions and also that, under particular parameter values, it can safely be ignored (sections 3.5 and 4.2.1). Moreover, optimal decisions concerning the recruitment rate and duration are affected by the manner in which the population of patients to benefit from the trial is defined. The application of section 4 showed how the value of the trial is highly sensitive to the optimal decision of trial length under a fixed horizon, but is less sensitive under a fixed patient pool. Further, we showed that optimising the trial recruitment duration of our model can lead to a sample size which is less than (fixed horizon) *or* greater than (fixed patient pool) the sample size of a frequentist trial. So, it is not necessarily the case that a value-based design such as the one we solve will lead to a trial which is of shorter duration. We also showed that trials which recruit more quickly, so as to make an earlier decision, provide larger expected benefits to patients. This analysis can provide useful insights at the design stage of the trial: our results suggest that providing a grant that ensures a higher recruitment rate and a shorter trial length can accrue higher rewards for patients. We have provided both qualitative and quantitative comparative statics results, the former of which, to the best of our knowledge, have not been obtained previously. We have also obtained asymptotic results that complement the comparative statics analysis. Notably, we show that, generally but not always, the optimal number of enrolled patients is increasing with the post-trial population and decreasing with effective sample size of the prior distribution.

We briefly consider future research. Firstly, as seen in the application of section 4, the fixed horizon case can lead to a sample size which is low when compared with that which would be obtained from a frequentist design. This could lead to problems of incomplete or non-existent adoption. Our work also raises interesting questions about outcomes-based reimbursement contracts for a public sector procurer and a private technology provider. There most likely exists a large gap between how regulators would view 'innovation' and 'efficiency' in clinical trial design, and how Bayesian decision-theory views it. Finally, while this work focused on jointly considering clinical trial design and technology adoption decisions, some of the results associated with the regulatory context might find application in very different sectors, particularly when A/B tests have observations which occur long after exposure to the test stimulus, or when the value of a choice of A or B is highly time sensitive (so that delayed decisions result in lower exploitation value).

Acknowledgments

We thank Stephen Brealey of the York Trials Unit, University of York, and members of the ProFHER team, for helpful discussions and assistance in preparing the parameter values used in section 4. Alban and Chick acknowledge the support of the European Union through the MSCA-ESA-ITN project (676129). We thank seminar attendees at the University of Sheffield, MIT, INSEAD, and University of Lausanne for feeback.

References

Ahuja V, Birge JR (2016) Response-adaptive designs for clinical trials: Simultaneous learning from multiple patients. *European Journal of Operational Research* 248(2):619–633.

Alban A, Chick SE, Forster M (2018) Extending a Bayesian decision-theoretic approach to value-based sequential clinical trial design. Rabe M, Juan A, Mustafee N, Skoogh A, Jain S, Johansson B, eds., *Proceedings of the 2018 Winter Simulation Conference*, 2459–2470 (Piscataway, NJ: IEEE, Inc.).

Angus D, Musthafa A, Clermont G, Griffin M, Linde-Zwirle W, Dremsizov T, Pinsky M (2001) Quality-adjusted survival in the first year after the acute respiratory distress syndrome. *Am J Respir Crit Care Med.* 163(6):1389–94.

Berry D, Eick S (1995) Adaptive assignment versus balanced randomization in clinical trials: a decision analysis. *Statistics in Medicine* 14(3):231–246.

Berry DA, Ho C (1988) One-sided sequential stopping boundaries for clinical trials: a decision-theoretic approach. *Biometrics* 44:219–227.

Betrán AP, Ye J, Moller A, Zhang J, Gülmezoglu AM, Torloni MR (2016) The increasing trend in ceasarean section rates: global, regional and national estimates: 1990–2014. *PLoS One* 11(2):e0148343.

Branke J, Chick S, Schmidt C (2007) Selecting a selection procedure. *Management Science* 53(12):1916–1932.

Brennan A, Kharroubi S, O'Hagan A, Chilcott J (2007) Calculating partial expected value of information in cost-effectiveness models. *Medical Decision Making* 27:448–470.

Bubeck S, Cesa-Bianchi N (2012) Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5(1):1–122.

Chakma J, Sun GH, Steinberg JD, Sammut SM, Jagsi R (2014) Asia's ascent - global trends in biomedical R & D expenditures. *The New England Journal of Medicine* 370:3–6.

Chick SE, Forster M, Pertile P (2017) A Bayesian decision-theoretic model of sequential experimentation with delayed response. *Journal of the Royal Statistical Society, Series B* 79(5):1439–1462.

Chick SE, Gans N, Yapar O (2018) Bayesian sequential learning for clinical trials of multiple correlated medical interventions. https://ssrn.com/abstract=3184758, Accessed 25 Feb 2020.

Chick SE, Wu Y (2005) Selection procedures with frequentist expected opportunity cost bounds. *Operations Research* 53(5):867–878.

Claxton K (1999) The irrelevance of inference: a decision-making approach to the stochastic evaluation of health care technologies. *Journal of Health Economics* 18(3):341–364.

Claxton K, Lacey LF, Walker SG (2000) Selecting treatments: a decision theoretic approach. *Journal of the Royal Statistical Society, Series A* 163:211–215.

Claxton K, Posnett J (1996) An economic approach to clinical trial design and research priority-setting. *Health Economics* 5:513–524.

Corbacho B, Duarte A, Keding A, et al. (2016) Cost effectiveness of surgical *versus* non-surgical treatment of adults with displaced fractures of the proximal humerus. *Bone and Joint Journal* 92-B(2):152–159.

Costa ML, Achten J, Parsons NR, et al. (2014) Percutaneous fixation with Kirschner wires versus volar locking plate fixation in adults with dorsally displaced fracture of distal radius: randomised controlled trial. *BMJ* 349:4807.

DeGroot M (1970) *Optimal Statistical Decisions* (New York: McGraw-Hill), First edition.

DiMasi J, Grabowski H, Hansen R (2016) Innovation in the pharmaceutical industry: new estimates of R&D costs. *Journal of Health Economics* 47:20–33.

Draper D (2013) Discussion of the paper by Hampson and Jennison. *Journal of the Royal Statistical Society, Series B* 75(1):48.

Eckermann S, Willan AR (2007) Expected value of information and decision making in HTA. *Health Economics* 16:195–209.

European Union (2014) Clinical trial regulation EU no. 536/2014. https://www.ema.europa.eu/en/human-regulatory/research-development/clinical-trials/clinical-trial-regulation, Accessed 25 Feb 2020.

Fassbender M (2016) Quintiles: Clinical trial complexity demands innovative approaches. https://www.outsourcing-pharma.com/Article/2016/06/01/Quintiles-clinical-trial-complexity-demands-innovative-approaches, Accessed 29 Jan 2020.

FDA (2015) Patents and exclusivity. US Food and Drug Administration, https://www.fda.gov/downloads/drugs/developmentapprovalprocess/smallbusinessassistance/ucm447307.pdf, Accessed 25 Feb 2020.

FDA (2019) Interacting with the FDA on complex innovative trial designs for drugs and biological products. US Food and Drug Administration, https://www.fda.gov/regulatory-information/search-fda-guidance-documents/interacting-fda-complex-innovative-trial-designs-drugs-and-biological-products, Accessed 17 Feb 2020.

Ferguson ND, et al. (2013) Integrating mortality and morbidity outcomes using quality-adjusted life years in critical care trials. *Am J Respir Crit Care Med.* 187(3):256–261.

Forster M, Brealey S, Chick S, et al. (2019) Cost-Effective Clinical Trial Design: Application of a Bayesian Sequential Stopping Rule to the ProFHER Pragmatic Trial. Discussion Papers 1/19, Department of Economics, University of York.

Frazier PI, Powell WB, Dayanik S (2008) A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization* 47(5):2410–2439.

Girotra K, Terwiesch C, Ulrich KT (2007) Valuing R&D projects in a portfolio: Evidence from the pharmaceutical industry. *Management Science* 53(9):1452–1466.

Gittins J, Pezeshk H (2000) A behavioural Bayes method for determining the size of a clinical trial. *Drug Information Journal* 34:355–363.

Gold MR, Siegel JE, Russell LB, Weinstein MC (1996) *Cost-effectiveness in Health and Medicine* (Oxford University Press).

Griffin S, Welton NJ, Claxton K (2010) Exploring the research decision space: the expected value of information for sequential research designs. *Medical Decision Making* 30:155–162.

Handoll H, Brealey S, Rangan A, et al. (2015) The ProFHER (PROximal Fracture of the Humerus: Evaluation by Randomisation) trial - a pragmatic multicentre randomised controlled trial evaluating the clinical effectiveness and cost-effectivesness of surgical compared with non-surgical treatment for proximal fracture of the humerus in adults. *Health Technology Assessment* 19:1–280.

Hudson KL, Lauer MS, Collins FS (2016) Toward a new era of trust and transparency in clinical trials. *Journal of the American Medical Association* 316(13):1353–1354.

Jacob WF, Kwak YH (2003) In search of innovative techniques to evaluate pharmaceutical R&D projects. *Technovation* 23(4):291–296.

Jennison C, Turnbull BW (1989) Interim analyses: the repeated confidence interval approach. *Journal of the Royal Statistical Society Series B* 51(3):305–361.

Jobjörnsson S, Forster M, Pertile P, Burman C (2016) Late-stage pharmaceutical R&D and pricing policies under two-stage regulation. *Journal of Health Economics* 50:298–311.

Kouvelis P, Milner J, Tian Z (2017) Clinical trials for new drug development: Optimal investment and application. *Manuf. Serv. Oper. Manag.* 19(3):437–452.

Lachin JM (1981) Introduction to sample size determination and power analysis for clinical trials. *Control. Clin. Trials* 2:93–113.

Long EF, Vaidya N, Brandeau ML (2008) Controlling co-epidemics: Analysis of HIV and tuberculosis infection dynamics. *Operations Research* 56(6):1366–1381.

Lucas Jr RE (1971) Optimal management of a research and development project. *Management Science* 17(11):679–697.

NICE (2012) The guidelines manual: Process and methods [pmg6]. UK National Institute for Health and Care Excellence, https://www.nice.org.uk/process/pmg6/chapter/assessing-cost-effectiveness, Accessed 20 February 2020.

NICE (2013) Guide to the methods of technology appraisal. https://www.nice.org.uk/process/pmg9/chapter/foreword, Accessed 9 Feb 2020, London.

NICE (2014) Developing NICE guidelines: The manual. UK National Institute for Health and Care Excellence, Accessed March 20, 2018, https://www.nice.org.uk/process/pmg20/chapter/incorporating-economic-evaluation.

NIHR (2018) NIHR Efficacy and mechanism evaluation. https://www.nihr.ac.uk/explore-nihr/funding-programmes/efficacy-and-mechanism-evaluation.htm, Accessed 8 Feb 2020.

NIHR (2020) NIHR Annual Efficient Studies. Accessed Feb 27, 2020, https://www.nihr.ac.uk/documents/ad-hoc-funding-calls-for-ctu-projects/20141.

Paul SM, Mytelka DS, Dunwiddie CT, Persinger CC, Munos BH, Linborg SR, Schact AL (2014) How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nature Reviews Drug Discovery* 9:203–214.

Pertile P, Forster M, La Torre D (2014) Optimal Bayesian sequential sampling rules for the economic evaluation of health technologies. *Journal of the Royal Statistical Society, Series A* 177(2):419–438.

Raiffa H, Schlaifer R (1961) *Applied Statistical Decision Theory* (Boston: Harvard University Graduate School of Business Administration), first edition.

Rangan A, Handoll H, Brealey S, et al. (2015) Surgical vs nonsurgical treatment of adults with displaced fratures of the proximal humerus: the PROFHER randomized clinical trial. *Journal of the American Medical Association* 313(10):1037–1047.

Russo D (2020) Simple Bayesian algorithms for best arm identification. *Operations Research* 68:in press.

Ryzhov IO, Frazier PI, Powell WB (2012) The knowledge gradient algorithm for a general class of online learning problems. *Operations Research* 60(1):180–195.

Schwartz ES (2004) Patents and R&D as real ooptions. *Economic Notes* 33(1):23–54.

Strong M, Oakley JE, Brennan A, Breeze P (2015) Estimating the expected value of sample information using the probabilistic sensitivity analysis sample: A fast, nonparametric regression-based method. *Medical Decision Making* 35:570–583.

Villar SS, Bowden J, Wason J (2015a) Multi-armed Bandit models for the optimal design of clinical trials: Benefits and challenges. *Stat. Sci.* 30(2):199–215.

Villar SS, Wason J, Bowden J (2015b) Response-adaptive randomization for multi-arm clinical trials using the forward looking Gittins index rule. *Biometrics* 71:969–978.

Willan AR (2008) Optimal sample size determination from an industry perspective based on the expected value of information. *Clin. Trials* 5:587–594.

Willan AR, Pinto EM (2005) The value of information and optimal clinical trial design. *Stat. Med.* 24:1791–1806.

Williamson SF, Jacko P, Villar SS, Jaki T (2017) A Bayesian adaptive design for clinical trials in rare diseases. *Computational Statistics & Data Analysis* 113:136–153.

Appendices

Appendix A summarises the principal notation in the main paper. Appendix B gives technical details for the derivatives of the focal objective function. Appendix C justifies mathematical claims in section 3. Appendix D shows how the one-shot trial design of the main paper can be amenable to running a fully sequential trial in certain special cases. Appendix E provides formalism for the probabalistic performance measures in the numerical experiments of section 4.

Appendix A: Table of principal notation

Table EC.1 Table of principal notation.

Parameter	*Description*
$p_{\mathrm{N}} \in [0, 1/2]$	Fraction of patients treated with technology N under the current practice
$E_{\mathrm{N}}, E_{\mathrm{S}} \in \mathbb{R}$	Effectiveness of technologies N and S, respectively
$C_{\mathrm{N}}, C_{\mathrm{S}} \in \mathbb{R}_{\geq 0}$	Patient-level costs of using technologies N and S, respectively
$\lambda \in \mathbb{R}_{\geq 0}$	Monetary value of one unit of effectiveness (e.g., £30,000 / QALY)
$X \in \mathbb{R}$ (random variable)	Incremental net monetary benefit of technology N over S
$W \in \mathbb{R}$	Unknown expected value of X
$\sigma_X^2 \in \mathbb{R}_{>0}$	Known variance of X
$\mu_0 \in \mathbb{R}$, $\sigma_0^2 \in \mathbb{R}_{>0}$	Mean and variance of prior distribution for W
$n_0 = \sigma_X^2 / \sigma_0^2$	Effective sample size of prior distribution
$T_{\max} \in \mathbb{R}_{>0}$	Maximum time duration of recruitment in trial
$\Delta \in \mathbb{R}_{\geq 0}$, $\Delta < T_{\max}$	Delay in observing realisation of pairwise allocation (in time units)
$\zeta \in \mathbb{R}_{>0}$	Incidence rate of the condition in the population
$r_{\max} \in [0, \zeta]$	Capacity of rate of recruitment
$Q_{\max} \in \mathbb{N}$	Maximum number of pairwise allocations recruited to the trial
$T \in [0, T_{\max}]$	Recruitment period duration (decision variable)
$r \in [0, r_{\max}]$	Rate of recruitment to the trial (decision variable)
$\mathcal{D} \in \{\mathrm{M}, \mathrm{N}, \mathrm{S}\}$	Adoption decision to implement the current practice (mix of technologies) M, the new technology N, or the standard technology S (decision variable)
Z_{Tr}	Posterior mean to be obtained, given μ_0 and Tr time units of patients to be observed
$P(T) \in \mathbb{R}_{\geq 0}$	Number of patients to receive implemented technology once adoption decision is made at $T + \Delta$
H	Maximum time horizon for decision, in the case of fixed horizon
P	Number of patients affected by adoption decision, in the case of fixed patient pool
ρ_{yr}	Annual discount rate, e.g., 3.5% for UK NICE
$\rho \in [0, 1)$	Continuous time discount rate, $\rho = \ln(1 + \rho_{\mathrm{yr}})$
$\tilde{T}_\rho(T)$	Effective discounted recruitment period duration
$P_\rho(T)$	Effective discounted number of patients to receive implemented technology once adoption decision is made at $T + \Delta$
δ_{on}	1 = 'online learning'; 0 = 'offline learning'
$c \in \mathbb{R}_{\geq 0}$	Recruitment cost for an additional participant
$c_{\mathrm{cap}}(r)$	Setup cost of the trial with recruitment rate r
c_r	Marginal cost of an additional unit of recruitment rate
$I_{\mathcal{D}} \in \mathbb{R}_{\geq 0}$	Fixed cost of switching to technology $\mathcal{D}$ from standard technology
$\alpha_{\mathrm{N}}(T), \alpha_{\mathrm{S}}(T) \in \mathbb{R}$	Expected costs per patient if technology N or S is adopted
$\Psi(z)$	Standard normal loss function, $\Psi(z) \equiv \mathbb{E}\left[(Z - z)^+\right] = \phi(z) - z(1 - \Phi(z))$

Appendix B: Derivation of partial derivatives of V

The first order conditions for interior solutions of (11), assuming that $P(T)$ and $c_{\text{cap}}(r)$ are differentiable, are given by

$$\frac{\partial V(T,r)}{\partial T}=0 \text{ and } \frac{\partial V(T,r)}{\partial r}=0.$$

In this appendix, we show the main steps to find the partial derivatives of V with respect to T and r. We first introduce some derivatives that will be used repeatedly:

$$\frac{\partial \alpha_{\text{N}}(T)}{\partial T}=-\frac{\alpha_{\text{N}}(T)}{P_\rho(T)}\frac{dP_\rho(T)}{dT}, \quad \frac{\partial \alpha_{\text{S}}(T)}{\partial T}=-\frac{\alpha_{\text{S}}(T)}{P_\rho(T)}\frac{dP_\rho(T)}{dT},$$
$$\frac{\partial \sigma_Z}{\partial T}=\sqrt{\frac{n_0\sigma_X^2 r}{(2n_0+rT)^3T}}, \qquad \frac{\partial \sigma_Z}{\partial r}=\sqrt{\frac{n_0\sigma_X^2 T}{(2n_0+rT)^3 r}}.$$

The following equations are useful relationships in deriving the partial derivatives:

$$\frac{\partial((\alpha_{\text{N}}(T))/\sigma_Z)}{\partial T}=-\frac{\alpha_{\text{N}}(T)}{\sigma_Z}\left(\frac{1}{P_\rho(T)}\frac{dP_\rho(T)}{dT}+\frac{1}{\sigma_Z}\frac{\partial \sigma_Z}{\partial T}\right)$$
$$\frac{\partial((\alpha_{\text{S}}(T))/\sigma_Z)}{\partial T}=-\frac{\alpha_{\text{S}}(T)}{\sigma_Z}\left(\frac{1}{P_\rho(T)}\frac{dP_\rho(T)}{dT}+\frac{1}{\sigma_Z}\frac{\partial \sigma_Z}{\partial T}\right)$$
$$\frac{\partial \sigma_Z}{\partial T}=\frac{\sigma_X^2 r}{\sigma_Z(2n_0+rT)^2}=\frac{n_0\sigma_Z}{(2n_0+rT)T}$$
$$\frac{\partial \sigma_Z}{\partial r}=\frac{\sigma_X^2 T}{\sigma_Z(2n_0+rT)^2}=\frac{n_0\sigma_Z}{(2n_0+rT)r}.$$

Consider first

$$\begin{aligned}
&\frac{\partial}{\partial T}\left(e^{-\rho(T+\Delta)}P_\rho(T)\sigma_Z\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\right)\\
&=-\rho e^{-\rho(T+\Delta)}P_\rho(T)\sigma_Z\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)+e^{-\rho(T+\Delta)}\frac{dP_\rho(T)}{dT}\sigma_Z\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\\
&\quad+e^{-\rho(T+\Delta)}P_\rho(T)\frac{\partial \sigma_Z}{\partial T}\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\\
&\quad+e^{-\rho(T+\Delta)}P_\rho(T)\sigma_Z\left(\Phi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)-1\right)\left(-\frac{\alpha_{\text{N}}(T)}{\sigma_Z}\left(\frac{1}{P_\rho(T)}\frac{dP_\rho(T)}{dT}+\frac{1}{\sigma_Z}\frac{\partial \sigma_Z}{\partial T}\right)-\frac{\mu_0}{\sigma_Z^2}\frac{\partial \sigma_Z}{\partial T}\right)\\
&=e^{-\rho(T+\Delta)}\left[\left(\sigma_Z\frac{dP_\rho(T)}{dT}+P_\rho(T)\frac{\partial \sigma_Z}{\partial T}-P_\rho(T)\sigma_Z\rho\right)\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\right.\\
&\quad\left.+\left(\frac{\alpha_{\text{N}}(T)}{\sigma_Z}\left(\sigma_Z\frac{dP_\rho(T)}{dT}+P_\rho(T)\frac{\partial \sigma_Z}{\partial T}\right)-\frac{P_\rho(T)\mu_0}{\sigma_Z}\frac{\partial \sigma_Z}{\partial T}\right)\left(1-\Phi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\right)\right].
\end{aligned}$$

Similarly,

$$\begin{aligned}
&\frac{\partial}{\partial T}\left(e^{-\rho(T+\Delta)}P_\rho(T)\sigma_Z\Psi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right)\\
&=e^{-\rho(T+\Delta)}\left[\left(\sigma_Z\frac{dP_\rho(T)}{dT}+P_\rho(T)\frac{\partial \sigma_Z}{\partial T}-P_\rho(T)\sigma_Z\rho\right)\Psi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right.\\
&\quad\left.+\left(\frac{\alpha_{\text{S}}(T)}{\sigma_Z}\left(\sigma_Z\frac{dP_\rho(T)}{dT}+P_\rho(T)\frac{\partial \sigma_Z}{\partial T}\right)+\frac{P_\rho(T)\mu_0}{\sigma_Z}\frac{\partial \sigma_Z}{\partial T}\right)\left(1-\Phi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right)\right].
\end{aligned}$$

Thus, we get

$$\frac{\partial V(T,r)}{\partial T} = e^{-\rho T}\left(\frac{1}{2}\delta_{\text{on}} r(1-2p_{\text{N}})\mu_0 - cr\right) \tag{EC.1}$$
$$+ (1-p_{\text{N}})e^{-\rho(T+\Delta)}\left[\left(\sigma_Z\frac{dP_\rho(T)}{dT} + P_\rho(T)\frac{\partial\sigma_Z}{\partial T} - P_\rho(T)\sigma_Z\rho\right)\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\right.$$
$$\left. + \left(\frac{\alpha_{\text{N}}(T)}{\sigma_Z}\left(\sigma_Z\frac{dP_\rho(T)}{dT} + P_\rho(T)\frac{\partial\sigma_Z}{\partial T}\right) - \frac{P_\rho(T)\mu_0}{\sigma_Z}\frac{\partial\sigma_Z}{\partial T}\right)\left(1-\Phi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\right)\right]$$
$$+ p_{\text{N}}e^{-\rho(T+\Delta)}\left[\left(\sigma_Z\frac{dP_\rho(T)}{dT} + P_\rho(T)\frac{\partial\sigma_Z}{\partial T} - P_\rho(T)\sigma_Z\rho\right)\Psi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right.$$
$$\left. + \left(\frac{\alpha_{\text{S}}(T)}{\sigma_Z}\left(\sigma_Z\frac{dP_\rho(T)}{dT} + P_\rho(T)\frac{\partial\sigma_Z}{\partial T}\right) + \frac{P_\rho(T)\mu_0}{\sigma_Z}\frac{\partial\sigma_Z}{\partial T}\right)\left(1-\Phi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right)\right].$$

The partial derivative with respect to r follows similar steps to the ones to find the derivatives with respect to T, so we only present the final result:

$$\frac{\partial V(T,r)}{\partial r} = -\frac{\partial c_{\text{cap}}(r)}{\partial r} - c\tilde{T}_\rho(T) + \delta_{\text{on}}\tilde{T}_\rho(T)(1-2p_{\text{N}})\mu_0/2$$
$$+ e^{\rho(T+\Delta)}P_\rho(T)\sqrt{\frac{n_0\sigma_X^2 T}{(2n_0+rT)^3 r}}\left[(1-p_{\text{N}})\phi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right) + p_{\text{N}}\phi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right]$$

Appendix C: Proofs of mathematical claims

C.1. Proofs in section 3.3

Proof of Prop. 1. Weierstrass' theorem states that the optimal solution of $\max_{x\in S} f(x)$ exists if f is upper semi-continuous and S is closed and bounded (Andreasson et al. 2007, section 4.2). A function $f: S\to\mathbb{R}$ is upper semi-continuous at x_0 if $\limsup_{x\to x_0} f(x) \le f(x_0)$, or equivalently, if for every $\epsilon>0$ there is a neighborhood S' around x_0 such that $f(x)\le f(x_0)+\epsilon$ for all $x\in S'$.

The domain of $V(T,r)$ is $D=\{(T,r)\mid 0\le T\le T_{\max},\ 0\le r\le r_{\max}\}$ which is closed and bounded. Hence, to prove the existence of a solution, it is sufficient to show that $V(T,r)$ is upper semi-continuous.

It is easy to check that $P_\rho(T)$ is upper semi-continuous given that $P(T)$ is upper semi-continuous. Because $c_{\text{cap}}(r)$ is lower semi-continuous, $-c_{\text{cap}}(r)$ is upper semi-continuous. It follows that $V(T,r)$ is upper semi-continuous in $\{(T,r)\mid 0<T\le T_{\max},\ 0<r\le r_{\max}\}$ and we only need to show that $V(T,r)$ is upper semi-continuous when either $T=0$ or $r=0$.

For $T=0$, first consider the neighborhood when T is exactly zero. Then, $V(0,r+\delta_r)=V(0,r)$ for any δ_r in the domain D by the definition in (5a). Now, let $\delta_T>0$ and consider the following inequality using (5b):

$$\begin{aligned}V(\delta_T,r+\delta_r) &= -(c_{\text{cap}}(r+\delta_r)+c(r+\delta_r)\tilde{T}_\rho(\delta_T)) + \delta_{\text{on}}[(r+\delta_r)\tilde{T}_\rho(\delta_T)/2](1-2p_{\text{N}})\mu_0\\ &\quad + e^{-\rho(\delta_T+\Delta)}\mathbb{E}\left[(P_\rho(\delta_T)(1-p_{\text{N}})W-I_{\text{N}})^+ + (-P_\rho(\delta_T)p_{\text{N}}W-I_{\text{S}})^+\right]\\ &\le \delta_{\text{on}}[(r+\delta_r)\tilde{T}_\rho(\delta_T)/2](1-2p_{\text{N}})\mu_0\\ &\quad + e^{-\rho\delta_T}e^{-\rho\Delta}\mathbb{E}\left[(P_\rho(0)(1-p_{\text{N}})W-I_{\text{N}})^+ + (-P_\rho(0)p_{\text{N}}W-I_{\text{S}})^+\right].\end{aligned}$$

The inequality holds because $c_{\text{cap}}(r)\ge 0$ and $P_\rho(T)$ is non-increasing, so $P_\rho(\delta_T)\le P_\rho(0)$. Now, we use $e^{-\rho\delta_T}\le(1+\rho e^{\rho|\delta_T|}|\delta_T|)$ and $\mathbb{E}[(P_\rho(0)(1-p_{\text{N}})W-I_{\text{N}})^+ + (-P_\rho(0)p_{\text{N}}W-I_{\text{S}})^+]\le V(0,r)$ to get:

$$\begin{aligned}V(\delta_T,r+\delta_r) &\le \delta_{\text{on}}[r_{\max}\tilde{T}_\rho(\delta_T)/2](1-2p_{\text{N}})\mu_0 + (1+\rho e^{\rho|\delta_T|}|\delta_T|)V(0,r)\\ &\le V(0,r)+\epsilon,\end{aligned}$$

and the last inequality holds for any $\epsilon > 0$ if we choose $|\delta_T|$ small enough. Hence, V is upper semi-continuous at $T = 0$.

For $r = 0$, consider first the neighborhood when r is exactly zero. Then, $V(T + \delta_T, 0) = V(T, 0)$ for any δ_T in the domain of $V(T, r)$, by the definition of V in (5a). Now, let $\delta_r > 0$ and consider the following inequality that follows similar steps as the previous case with $T = 0$:

$$\begin{aligned}
V(T+\delta_T,\delta_r) &= -(c_{\text{cap}}(\delta_r) + c\delta_r\tilde{T}_\rho(T+\delta_T)) + \delta_{\text{on}}[(\delta_r)\tilde{T}_\rho(T+\delta_T)/2](1-2p_{\text{N}})\mu_0 \\
&\qquad + e^{-\rho(T+\delta_T+\Delta)}\mathbb{E}\left[(P_\rho(T+\delta_T)(1-p_{\text{N}})W - I_{\text{N}})^+ + (-P_\rho(T+\delta_T)p_{\text{N}}W - I_{\text{S}})^+\right] \\
&\le \delta_{\text{on}}[\delta_r\tilde{T}_\rho(\delta_T)/2](1-2p_{\text{N}})\mu_0 \\
&\qquad + e^{-\rho\delta_T}e^{-\rho(T+\Delta)}\mathbb{E}\left[(P_\rho(0)(1-p_{\text{N}})W - I_{\text{N}})^+ + (-P_\rho(0)p_{\text{N}}W - I_{\text{S}})^+\right] \\
&\le \delta_{\text{on}}[\delta_r\tilde{T}_\rho(T_{\max})/2](1-2p_{\text{N}})\mu_0 + (1+\rho e^{\rho|\delta_T|}|\delta_T|)V(T,0) \\
&\le V(T,0) + \epsilon,
\end{aligned}$$

and the last inequality holds for any $\epsilon > 0$ in the neighborhood with $|\delta_r|$ and $|\delta_T|$ small enough. Notice that the first inequality requires that $P(T)$ is non-increasing not only upper semi-continuous. Thus, we have shown that V is upper semi-continuous for $r = 0$. □

C.2. Proofs in section 3.4

Proof of Prop. 2. It is sufficient to show that $V(T, r)$ is the same for all members of the the set $\mathcal{S} = \{(T, r) \in [0, T_{\max}] \times [0, r_{\max}] : rT/2 = Q\}$. If $Q = 0$, then $V(r, T) = V(0, 0)$ for all members $\mathcal{S}$. This follows directly from the definition of V in (5a). Now suppose that $Q > 0$, and let $\sigma_Z^2 = \sigma_X^2 Q/(n_0(2n_0 + Q))$. Then, we obtain, for all members of the set $\mathcal{S}$, the same expected net gain:

$$V(T,r) = \delta_{\text{on}}Q(1-2p_{\text{N}})\mu_0/2 - cQ - c_{\text{cap}} + P\sigma_Z\left[(1-p_{\text{N}})\Psi\left(\frac{\alpha_{\text{N}}-\mu_0}{\sigma_Z}\right) + p_{\text{N}}\Psi\left(\frac{\alpha_{\text{S}}+\mu_0}{\sigma_Z}\right)\right]. \quad \square$$

Proof of Prop. 3. Let (T^*, r^*) be an optimal solution to (11). First, observe that if $T^* = 0$ or $r^* = 0$, then $(0, r_{\max})$ is also an optimal solution by the definition of V in (5a).

Now, consider the case $T^* > 0$ and $r^* > 0$. We show that the alternative trial design $(r^*T^*/r_{\max}, r_{\max})$ achieves an expected net gain at least as good as the optimal solution (T^*, r^*). Note that the online learning term, $\delta_{\text{on}}[r\tilde{T}_\rho(T)/2](1-2p_{\text{N}})\mu_0$, and the cost of the trial, $-(c_{\text{cap}} + cr\tilde{T}_\rho(T))$, in (5b) are smaller or equal to the corresponding costs under the alternative solution $(r^*T^*/r_{\max}, r_{\max})$:

$$\begin{aligned}
&\delta_{\text{on}}[r^*\tilde{T}_\rho(T^*)/2](1-2p_{\text{N}})\mu_0 - (c_{\text{cap}} + cr^*\tilde{T}_\rho(T^*)) \\
&\qquad \ge \delta_{\text{on}}[r^*\tilde{T}_\rho(r^*T^*/r_{\max})/2](1-2p_{\text{N}})\mu_0 - (c_{\text{cap}} + cr^*\tilde{T}_\rho(r^*T^*/r_{\max})).
\end{aligned}$$

By construction, the number of patients recruited under the alternative design is the same as under the optimal solution and the posterior mean variance, σ_Z^2, is the same under both solutions. The non-increasing assumption of $P(T)$ implies that $P_\rho(T)$ is non-increasing, and, hence, $P_\rho(T^*) \le P_\rho(r^*T^*/r_{\max})$, $\alpha_{\text{N}}(T^*) \ge \alpha_{\text{N}}(r^*T^*/r_{\max})$, and $\alpha_{\text{S}}(T^*) \ge \alpha_{\text{S}}(r^*T^*/r_{\max})$. Because $\Psi(\cdot)$ is decreasing, we have established that

$$\begin{aligned}
&P_\rho(T^*)\sigma_Z\left[(1-p_{\text{N}})\Psi\left(\frac{\alpha_{\text{N}}(T^*)-\mu_0}{\sigma_Z}\right) + p_{\text{N}}\Psi\left(\frac{\alpha_{\text{S}}(T^*)+\mu_0}{\sigma_Z}\right)\right] \\
&\qquad \le P_\rho(r^*T^*/r_{\max})\sigma_Z\left[(1-p_{\text{N}})\Psi\left(\frac{\alpha_{\text{N}}(r^*T^*/r_{\max})-\mu_0}{\sigma_Z}\right) + p_{\text{N}}\Psi\left(\frac{\alpha_{\text{S}}(r^*T^*/r_{\max})+\mu_0}{\sigma_Z}\right)\right],
\end{aligned}$$

which are the remaining terms in the expression of V. Thus, we conclude that $V(r^*T^*/r_{\max}, r_{\max}) \geq V(T^*, r^*)$. Because (T^*, r^*) is optimal, we have established that $V(r^*T^*/r_{\max}, r_{\max}) = V(T^*, r^*)$ and that an optimal solution with $r = r_{\max}$ exists. □

Proof of Prop. 4. Let (T^*, r^*) be an optimal solution and consider the alternative solution $(T_{\max}, r^*T^*/T_{\max})$. It is straightforward to show that the alternative solution is also optimal by following the same procedure as in the proof of Prop. 3. □

C.3. Comparative statics results in section 3.5

In this appendix, we present the algebra that leads to the results of comparative statics of (11) presented in section 3.5. For simplicity, we present the results assuming a fixed recruitment rate, $\max_{T\in[0,T_{\max}]} V(T,r)$, and undiscounted rewards. The results are, however, as general as presented in section 3.5.

- *Post-trial population.* For fixed patient pool, the optimal expected net gain is increasing in P:

$$\begin{aligned}\frac{\partial V(T,r)}{\partial P} =& (1-p_{\mathrm{N}})\sigma_Z \Psi\left(\frac{\alpha_{\mathrm{N}}-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\sigma_Z \Psi\left(\frac{\alpha_{\mathrm{S}}+\mu_0}{\sigma_Z}\right) \\ &+ (1-p_{\mathrm{N}})\alpha_{\mathrm{N}}\left(1-\Phi\left(\frac{\alpha_{\mathrm{N}}-\mu_0}{\sigma_Z}\right)\right) + p_{\mathrm{N}}\alpha_{\mathrm{S}}\left(1-\Phi\left(\frac{\alpha_{\mathrm{S}}+\mu_0}{\sigma_Z}\right)\right) > 0.\end{aligned}$$

  Using the implicit function theorem, the sign of $\partial T^*/\partial P$ is the same as the sign of:

$$\begin{aligned}\frac{\partial^2 V(T,r)}{\partial T \partial P} =& \frac{\partial \sigma_Z}{\partial T}\left[(1-p_{\mathrm{N}})\phi\left(\frac{\alpha_{\mathrm{N}}-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\phi\left(\frac{\alpha_{\mathrm{S}}+\mu_0}{\sigma_Z}\right)\right. \\ &\left. + (1-p_{\mathrm{N}})\frac{\alpha_{\mathrm{N}}(\alpha_{\mathrm{N}}-\mu_0)}{\sigma_Z^2}\phi\left(\frac{\alpha_{\mathrm{N}}-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\frac{\alpha_{\mathrm{S}}(\alpha_{\mathrm{S}}+\mu_0)}{\sigma_Z^2}\phi\left(\frac{\alpha_{\mathrm{S}}+\mu_0}{\sigma_Z}\right)\right].\end{aligned}$$

  The first two terms in brackets have a positive effect on the optimal trial length, because the value of information is higher when P is higher. The third and fourth terms come from the change in the adoption decision rule because α_{N} and α_{S} decrease as P increases. When $\alpha_{\mathrm{N}} < \mu_0$ there is a negative contribution from the third term but positive from all other terms. Similarly, when $\alpha_{\mathrm{S}} < -\mu_0$ there is a negative contribution from the fourth term but negative from all other terms. In general, we cannot conclude that the optimal T is increasing with P, but this is often the case. For instance, if $-\alpha_{\mathrm{S}} \leq \mu_0 \leq \alpha_{\mathrm{N}}$, or equivalently, it is a priori optimal to adopt M, then T^* is increasing with P.

  For fixed horizon, we perform the sensitivity analysis on parameter H. The results are very similar to fixed patient pool. The expected net gain is again increasing in H:

$$\begin{aligned}\frac{\partial V(T,r)}{\partial H} =& \zeta\left[(1-p_{\mathrm{N}})\sigma_Z \Psi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\sigma_Z \Psi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\right. \\ &\left. + (1-p_{\mathrm{N}})\alpha_{\mathrm{N}}(T)\left(1-\Phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right)\right) + p_{\mathrm{N}}\alpha_{\mathrm{S}}(T)\left(1-\Phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\right)\right] > 0\end{aligned}$$

  The derivative that determines the direction of change of the optimal trial length is

$$\begin{aligned}\frac{\partial^2 V(T,r)}{\partial T \partial H} =& \zeta\frac{\partial \sigma_Z}{\partial T}\left[(1-p_{\mathrm{N}})\phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\right. \\ &\left. + (1-p_{\mathrm{N}})\frac{\alpha_{\mathrm{N}}(T)(\alpha_{\mathrm{N}}(T)-\mu_0)}{\sigma_Z^2}\phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\frac{\alpha_{\mathrm{S}}(T)(\alpha_{\mathrm{S}}(T)+\mu_0)}{\sigma_Z^2}\phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\right] \\ &- \frac{\zeta(1-p_{\mathrm{N}})\alpha_{\mathrm{N}}(T)^2}{\sigma_Z(H-\Delta-T)}\phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right) - \frac{\zeta p_{\mathrm{N}}\alpha_{\mathrm{S}}(T)^2}{\sigma_Z(H-\Delta-T)}\phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right).\end{aligned}$$

Comparing this expression to that for $\frac{\partial^2 V(T,r)}{\partial T \partial P}$, we have two additional terms with a negative effect which capture the reduction in the post-trial patients for a longer trial. The direction of change of T^* is not definitive through comparative statics.

- *Fraction of patients in new technology p_N.* It is is straightforward to check that

$$\begin{aligned}\frac{\partial V(T,r)}{\partial p_{\mathrm{N}}} = &-2\delta_{\mathrm{on}} r T \mu_0 + P(T)\sigma_Z \left(\Psi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right) - \Psi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right)\right) \\ &+ P(T)\left[\alpha_{\mathrm{S}}(T)\left(1-\Phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\right) - \alpha_{\mathrm{N}}(T)\left(1-\Phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right)\right)\right].\end{aligned}$$

There are three summands in this expression. The first is the online learning effect. The second is the effect due to the post-trial benefits. The third is a correction due to the change of the optimal adoption decision rule. The magnitude of the second summand depends on the magnitude of $P(T)\sigma_Z$, while the magnitude of the third summand depends on the magnitudes of $P(T)\alpha_{\mathrm{N}}(T)$ and $P(T)\alpha_{\mathrm{S}}(T)$. When $P(T)\sigma_Z$ dominates, the effect depends on the expected gains of adopting technologies N or S; if $\Psi\left((\alpha_{\mathrm{S}}(T)+\mu_0)/\sigma_Z\right) > \Psi\left((\alpha_{\mathrm{N}}(T)-\mu_0)/\sigma_Z\right)$, then an increase in p_{N} (increase in the number of people who would benefit from the adoption of S) would increase the expected net gain.

The second derivative $\partial^2 V(T,r)/\partial T \partial p_{\mathrm{N}}$ is a long expression that does not elucidate any interesting insights. We present a special case with some interesting results, $I_{\mathrm{N}} = I_{\mathrm{S}} = 0$. Note that $\alpha_{\mathrm{N}}(T) = \alpha_{\mathrm{S}}(T) = 0$ and we obtain

$$\frac{\partial^2 V(T,r)}{\partial T \partial p_{\mathrm{N}}} = -2\delta_{\mathrm{on}} r \mu_0 - 2\frac{dP(T)}{dT}\mu_0\left(1-\Phi\left(\frac{\mu_0}{\sigma_Z}\right)\right).$$

From this expression, observe that under some additional conditions the optimal trial length becomes independent of p_{N}. One possibility is $\mu_0 = 0$. The second possibility is under fixed patient pool and offline learning. However, the optimal expected net gain is not independent of p_{N} under such conditions.

- *Effective number of samples for the prior distribution.* To analyse the sensitivity of V to n_0, we compute

$$\frac{\partial V(T,r)}{\partial n_0} = -\frac{P(T)\sigma_Z^2(2n_0+rT)}{2n_0(n_0+rT)}\left[(1-p_{\mathrm{N}})\phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\right] < 0.$$

The optimal expected net gain is decreasing in n_0 because n_0 is a measure of the confidence in the beliefs, and therefore, the lower the confidence, the higher the rewards obtained from learning through a trial.

To analyse the sensitivity on the optimal trial length, we use the fixed patient pool adoption decision type for simplicity and obtain the following derivative:

$$\begin{aligned}\frac{\partial^2 V(T,r)}{\partial T \partial n_0} = &P\frac{\partial^2 \sigma_Z}{\partial T \partial n_0}\left[(1-p_{\mathrm{N}})\phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right) + p_{\mathrm{N}}\phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\right] \\ &+ (1-p_{\mathrm{N}})P\phi\left(\frac{\alpha_{\mathrm{N}}(T)-\mu_0}{\sigma_Z}\right)\frac{(\alpha_{\mathrm{N}}(T)-\mu_0)^2}{\sigma_Z^3}\frac{\partial \sigma_Z}{\partial T}\frac{\partial \sigma_Z}{\partial n_0} \\ &+ p_{\mathrm{N}}P\phi\left(\frac{\alpha_{\mathrm{S}}(T)+\mu_0}{\sigma_Z}\right)\frac{(\alpha_{\mathrm{S}}(T)+\mu_0)^2}{\sigma_Z^3}\frac{\partial \sigma_Z}{\partial T}\frac{\partial \sigma_Z}{\partial n_0}.\end{aligned}$$

The second and third terms are negative because $\partial \sigma_Z/\partial n_0 < 0$, which correspond to the effect of a larger n_0 requiring a smaller sample size to achieve the required confidence in a decision. These

two terms correspond to the intuitive effect of a less informative prior requiring more evidence to be gathered. However, an additional effect is found in the first term which is positive for $n_0 < rT/4$ because $\partial^2\sigma_Z/\partial T\partial n_0 = \sigma_Z(-4n_0+rT)/4(2n_0+rT)^2T$. Therefore, it is possible that T^* increases with n_0 in some range where $n_0 < rT^*/4$.

C.4. Proofs for asymptotic results of Prop. 5 in section 3.6.1

Claims are proved separately for the four cases of section 3.4 in Lemmas EC.1 to EC.4. The combination of the four lemmas completes the proof of Prop. 5. The proofs assume that $c-\delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2>0$ and that there is no discounting, $\rho=0$. The proofs further assume that $c_{\text{cap}}(r)=c_{\text{fixed}}+c_r r$ where $c_r=0$ for cases I and II with constant setup costs and where $c_r>0$ for cases III and IV with variable costs.

C.4.1. Case I: constant setup costs, no discounting, fixed patient pool. If $c-\delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 \leq 0$ in this case, the effective cost of sampling is less than or equal to zero, and, given the trial is run, it is optimal to sample infinitely for any P. We, therefore, analyse the more interesting setting where $c-\delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2>0$. Because Q^* grows unbounded with P (shown below), we let $r_{\max}=\infty$ and $T_{\max}=\infty$, i.e., there is no upper bound on the decision variables.

LEMMA EC.1. *If $c_{cap}(r)=c_{cap}$, $\rho=0$, $P(T)=P$ and $c-\delta_{\text{on}}(1-2p_N)\mu_0/2>0$, then*

$$\lim_{P\to\infty}\frac{Q^*}{\sqrt{P}}=\left(\frac{\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{4(c-\delta_{\text{on}}(1-2p_N)\mu_0/2)}\right)^{1/2},$$
$$\lim_{P\to\infty}\frac{V^*}{P}=(1-p_N)\sigma_0\Psi(-\mu_0/\sigma_0)+p_N\sigma_0\Psi(\mu_0/\sigma_0).$$

Proof of Lemma EC.1. For functions f and g, we use the notation $f(P)\sim g(P)$ to denote $\lim_{P\to\infty} f(P)/g(P)=1$. Without loss of generality, the proof assumes that r is fixed and T is the decision variable. For clarity, we make the dependence of T^* on P explicit with $T^*(P)$. We need to show that

$$rT^*(P)\sim\left(\frac{\sqrt{n_0\sigma_X^2}\phi(\sqrt{n_0}\mu_0/\sigma_X)P}{c-\delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2}\right)^{1/2}.$$

We know that $T^*(P)$ either satisfies the first order condition $\partial V(T,r)/\partial T=0$ or $T^*(P)=0$. However, $T^*(P)=0$ is not optimal for sufficiently large P and $T^*(P)$ satisfies the first order condition. The expression $\partial V(T,r)/\partial T=0$ was derived in Appendix B. By rearranging terms, we obtain

$$P\left(\frac{1-p_{\text{N}}}{\sqrt{2\pi}}e^{-\frac{n_0(2n_0+rT^*(P))}{2\sigma_X^2 rT^*(P)}(\alpha_{\text{N}}-\mu_0)^2}+\frac{p_{\text{N}}}{\sqrt{2\pi}}e^{-\frac{n_0(2n_0+rT^*(P))}{2\sigma_X^2 rT^*(P)}(\alpha_{\text{S}}+\mu_0)^2}\right) \tag{EC.2}$$
$$=\sqrt{\frac{(2n_0+rT^*(P))^3T^*(P)}{n_0\sigma_X^2 r}}(cr-\delta_{\text{on}}r(1-2p_{\text{N}})\mu_0/2).$$

The left-hand side approaches infinity as $P\to\infty$, so that the right-hand side needs to approach infinity as $P\to\infty$. Hence, $\lim_{P\to\infty}T^*(P)=\infty$, and the right-hand side satisfies the following relationship:

$$\sqrt{\frac{(2n_0+rT^*(P))^3T^*(P)}{n_0\sigma_X^2 r}}(cr-\delta_{\text{on}}r(1-2p_{\text{N}})\mu_0/2)\sim\frac{c-\delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2}{\sqrt{n_0\sigma_X^2}}r^2(T^*(P))^2$$

Because $\alpha_N = I_N/(1-p_N)P$ and $\alpha_S = I_S/p_N P$, which approach zero as $P \to \infty$, the left-hand side of (EC.2) satisfies the following relationship:

$$P\left(\frac{1-p_N}{\sqrt{2\pi}}e^{-\frac{n_0(2n_0+rT^*(P))}{2\sigma_X^2 rT^*(P)}(\alpha_N-\mu_0)^2}+\frac{p_N}{\sqrt{2\pi}}e^{-\frac{n_0(2n_0+rT^*(P))}{2\sigma_X^2 rT^*(P)}(\alpha_S+\mu_0)^2}\right)\sim\frac{P}{\sqrt{2\pi}}e^{-\frac{n_0}{2\sigma_X^2}\mu_0^2}.$$

Combining results, we obtain

$$\frac{P}{\sqrt{2\pi}}e^{-\frac{n_0}{2\sigma_X^2}\mu_0^2}\sim\frac{c-\delta_{on}(1-2p_N)\mu_0/2}{\sqrt{n_0\sigma_X^2}}r^2(T^*(P))^2$$

and, by rearranging terms, the desired result for $Q^*/\sqrt{P}$.

Using the previous result, the asymptotic result for $V^* = V(T^*)$ is straightforward to find. Let σ_Z^* be σ_Z evaluated at T^*. Because $\lim_{P\to\infty} T^* = \infty$, we know that $\lim_{P\to\infty}\sigma_Z^* = \sigma_X/\sqrt{n_0} = \sigma_0$. Then,

$$\begin{aligned}
V^* =& -c_{cap}(r) - crT^* + \delta_{on}\ (1-2p_N)\mu_0 rT^*/2 + P\sigma_Z^*\left((1-p_N)\Psi\left(\frac{\alpha_N-\mu_0}{\sigma_Z^*}\right)+p_N\Psi\left(\frac{\alpha_S+\mu_0}{\sigma_Z^*}\right)\right)\\
V^* \sim& -crT^* + \delta_{on}(1-2p_N)\mu_0 rT^*/2 + P\sigma_Z^*\left((1-p_N)\Psi\left(\frac{-\mu_0}{\sigma_Z^*}\right)+p_N\Psi\left(\frac{\mu_0}{\sigma_Z^*}\right)\right)\\
V^* \sim& -crT^* + \delta_{on}(1-2p_N)\mu_0 rT^*/2 + P\sigma_0\left((1-p_N)\Psi\left(\frac{-\mu_0}{\sigma_0}\right)+p_N\Psi\left(\frac{\mu_0}{\sigma_0}\right)\right)\\
V^* \sim& (-c+\delta_{on}(1-2p_N)\mu_0/2)\left(\frac{\sqrt{n_0\sigma_X^2}\phi(\sqrt{n_0}\mu_0/\sigma_X)}{(c-\delta_{on}(1-2p_N)\mu_0/2)r^2}\right)^{1/2}\sqrt{P}+P\sigma_0\left((1-p_N)\Psi\left(\frac{-\mu_0}{\sigma_0}\right)+p_N\Psi\left(\frac{\mu_0}{\sigma_0}\right)\right)\\
V^* \sim& P\sigma_0\left((1-p_N)\Psi\left(\frac{-\mu_0}{\sigma_0}\right)+p_N\Psi\left(\frac{\mu_0}{\sigma_0}\right)\right),
\end{aligned}$$

and the last line is the statement in the lemma. □

C.4.2. Case II: constant setup costs, no discounting, fixed horizon. Because we assume undiscounted rewards, case II is necessarily fixed horizon. If $c-\delta_{on}(1-2p_N)\mu_0/2 \le 0$ in this case, then the first order condition for T^* is never satisfied, and it is optimal to sample infinitely. We therefore analyse the setting where $c-\delta_{on}(1-2p_N)\mu_0/2 > 0$. In addition, we allow $T_{max} = \infty$, i.e., there is no upper bound on T, because $T^* \to \infty$ as $P \to \infty$, as shown below. Prop. 3 shows that $r^* = r_{max}$ is optimal. By multiplying T^* with $r_{max}/2$, we obtain asymptotic results for Q^*.

LEMMA EC.2. *If $c_{cap}(r) = c_{cap}$, $\rho = 0$, $P(T) = \zeta(H - T - \mathbf{1}_{T>0}\Delta)$, $c-\delta_{on}(1-2p_N)\mu_0/2 > 0$, and $T_{max} = \infty$, then*

$$\begin{aligned}
\lim_{H\to\infty}\frac{T^*}{\sqrt{\zeta(H-\Delta)}} &= \left(\frac{\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{(c-\delta_{on}(1-2p_N)\mu_0/2)r_{max}^2+\zeta r_{max}\sigma_0(\Psi(\mu_0/\sigma_0)+(1-p_N)\mu_0/\sigma_0)}\right)^{1/2}, \qquad \text{(EC.3)}\\
\lim_{H\to\infty}\frac{V^*}{\zeta(H-\Delta)} &= (1-p_N)\sigma_0\Psi(-\mu_0/\sigma_0)+p_N\sigma_0\Psi(\mu_0/\sigma_0).
\end{aligned}$$

Proof of Lemma EC.2. For functions f and g, we use the notation $f(H) \sim g(H)$ to denote $\lim_{H\to\infty} f(H)/g(H) = 1$. For clarity, we make the dependence of T^* on H explicit with $T^*(H)$. We need to show that

$$T^*(H) \sim \left(\frac{\sqrt{n_0\sigma_X^2}\phi(\sqrt{n_0}\mu_0/\sigma_X)\zeta(H-\Delta)}{(c-\delta_{on}(1-2p_N)\mu_0/2)r^2+\zeta r\sigma_X/\sqrt{n_0}(\Psi(\mu_0\sqrt{n_0}/\sigma_X)+(1-p_N)\mu_0\sqrt{n_0}/\sigma_X)}\right)^{1/2}.$$

We know that $T^*(H)$ either satisfies the first order condition $dV(T^*(H))/dT = 0$ or $T^*(H) = 0$. Because $T^*(H) = 0$ is not optimal for sufficiently large H, $T^*(H)$ satisfies $dV(T^*(H))/dT = 0$. Using the expression derived in Appendix B, $dV(T^*(H))/dT = 0$ is equivalent to

$$\begin{aligned} cr - \delta_{\text{on}} r(1-2p_{\text{N}})\mu_0/2 =& (1-p_{\text{N}})\Bigg[\zeta(H - T^*(H) - \Delta)\frac{\partial \sigma_Z^*(H)}{\partial T}\phi\left(\frac{\alpha_{\text{N}}(T^*(H)) - \mu_0}{\sigma_Z^*(H)}\right) \\ &- \sigma_Z^*(H)\zeta\Psi\left(\frac{\alpha_{\text{N}}(T^*(H)) - \mu_0}{\sigma_Z^*(H)}\right) - \alpha_{\text{N}}(T^*(H))\zeta\left(1 - \Phi\left(\frac{\alpha_{\text{N}}(T^*(H)) - \mu_0}{\sigma_Z^*(H)}\right)\right)\Bigg] \\ &+ p_{\text{N}}\Bigg[\zeta(H - T^*(H) - \Delta)\frac{\partial \sigma_Z^*(H)}{\partial T}\phi\left(\frac{\alpha_{\text{S}}(T^*(H)) + \mu_0}{\sigma_Z^*(H)}\right) \\ &- \sigma_Z^*(H)\zeta\Psi\left(\frac{\alpha_{\text{S}}(T^*(H)) + \mu_0}{\sigma_Z^*(H)}\right) - \alpha_{\text{S}}(T^*(H))\zeta\left(1 - \Phi\left(\frac{\alpha_{\text{S}}(T^*(H)) + \mu_0}{\sigma_Z^*(H)}\right)\right)\Bigg], \end{aligned}$$

where $\sigma_Z^*(H)$ is σ_Z evaluated at $T^*(H)$. Because the left side of the equation is bounded for any H, the right side must be as well. Thus, $(H - T^*(H) - \Delta)\partial\sigma_Z^*(H)/\partial T$ is bounded and implies that $\lim_{H\to\infty} T^*(H) = \infty$, $\lim_{H\to\infty}(H - T^*(H) - \Delta) = \infty$, and $(H - \Delta) \sim (H - T^*(H) - \Delta)$, with the use of expression for $\partial\sigma_Z/\partial T$ derived in Appendix B. Using this observation on the above equation, we obtain the following relationship:

$$\begin{aligned} cr - \delta_{\text{on}} r(1-2p_{\text{N}})\mu_0/2 + \zeta\frac{\sigma_X}{\sqrt{n_0}}\left(\Psi\left(\frac{\mu_0\sqrt{n_0}}{\sigma_X}\right) + (1-p_{\text{N}})\frac{\mu_0\sqrt{n_0}}{\sigma_X}\right) \\ \sim \zeta(H - T^*(H) - \Delta)\sqrt{\frac{n_0\sigma_X^2 r}{(2n_0 + rT^*(H))^3 T^*(H)}}\phi\left(\frac{\mu_0\sqrt{n_0}}{\sigma_X}\right). \end{aligned}$$

By rearranging and some additional asymptotic approximations, we obtain

$$\left(cr - \delta_{\text{on}} r(1-2p_{\text{N}})\mu_0/2 + \zeta\frac{\sigma_X}{\sqrt{n_0}}\left(\Psi\left(\frac{\mu_0\sqrt{n_0}}{\sigma_X}\right) + (1-p_{\text{N}})\frac{\mu_0\sqrt{n_0}}{\sigma_X}\right)\right)\frac{r(T^*(H))^2}{\sqrt{n_0}\sigma_X\phi\left(\frac{\mu_0\sqrt{n_0}}{\sigma_X}\right)} \sim \zeta(H - \Delta).$$

The desired result is obtained by rearranging the terms and setting $r = r_{\max}$, which we know to be an optimal decision from Prop. 3. Using this result and following a similar procedure as in the proof of lemma EC.1, we can obtain the asymptotic result for V^*:

$$\begin{aligned} V^* =& -c_{\text{cap}}(r) - crT^* + \delta_{\text{on}}\ (1-2p_{\text{N}})\mu_0 rT^*/2 + P(T^*)\sigma_Z^*\left((1-p_{\text{N}})\Psi\left(\frac{\alpha_{\text{N}}(T^*) - \mu_0}{\sigma_Z^*}\right) + p_{\text{N}}\Psi\left(\frac{\alpha_{\text{S}}(T^*) + \mu_0}{\sigma_Z^*}\right)\right) \\ V^* \sim& -crT^* + \delta_{\text{on}}(1-2p_{\text{N}})\mu_0 rT^*/2 + \zeta(H - T^* - \Delta)\sigma_0\left((1-p_{\text{N}})\Psi\left(\frac{-\mu_0}{\sigma_0}\right) + p_{\text{N}}\Psi\left(\frac{\mu_0}{\sigma_0}\right)\right) \\ V^* \sim& \zeta(H - \Delta)\sigma_0\left((1-p_{\text{N}})\Psi\left(\frac{-\mu_0}{\sigma_0}\right) + p_{\text{N}}\Psi\left(\frac{\mu_0}{\sigma_0}\right)\right). \quad \square \end{aligned}$$

C.4.3. Case III: affine setup costs, no discounting, fixed patient pool. If $c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 \le 0$, then the first order condition is never satisfied and it is optimal to sample infinitely. In that scenario, a solution to (11) does not exist unless r is constrained and $r^* = r_{\max}$. We therefore analyse the setting where $c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 > 0$. Because the optimal recruitment rate grows unboundedly with P (as shown below), we assume that $r_{\max} = \infty$, i.e., no upper bound on r. To obtain asymptotic results on Q^*, we only need to multiply the results of r^* with $T_{\max}/2$, as Prop. 4 shows that $T = T_{\max}$ is optimal.

LEMMA EC.3. *If $c_{cap}(r) = c_{fixed} + c_r r$ where $c_r > 0$, $\rho = 0$, $P(T) = P$ and $c - \delta_{\text{on}}(1-2p_N)\mu_0/2 > 0$ then*

$$\begin{aligned} \lim_{P\to\infty}\frac{r^*}{\sqrt{P}} &= \left(\frac{\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{[cT_{\max} + c_r - \delta_{\text{on}}(1-2p_N)\mu_0 T_{\max}/2]}\right)^{1/2}, \qquad \text{(EC.4)} \\ \lim_{P\to\infty}\frac{V^*}{P} &= \sigma_0[\Psi(\mu_0/\sigma_0) + (1-p_N)\mu_0/\sigma_0]. \end{aligned}$$

Proof of Lemma EC.3. Prop. 4 shows that it is optimal to have $T^* = T_{\max}$. The optimal r has to satisfy the first order optimality condition $\partial V(T_{\max}, r)/\partial r = 0$ for large enough P. The optimality condition is equivalent to

$$P\left(\frac{1-p_{\mathrm{N}}}{\sqrt{2\pi}}e^{-\frac{n_0(2n_0+r^*(P)T_{\max})}{2\sigma_X^2 r^*(P)T_{\max}}(\alpha_{\mathrm{N}}-\mu_0)^2}+\frac{p_{\mathrm{N}}}{\sqrt{2\pi}}e^{-\frac{n_0(2n_0+r^*(P)T_{\max})}{2\sigma_X^2 r^*(P)T_{\max}}(\alpha_{\mathrm{S}}+\mu_0)^2}\right)$$
$$=\sqrt{\frac{(2n_0+r^*(P)T_{\max})^3 r^*(P)}{n_0\sigma_X^2 T_{\max}}}(c_r+cT-\delta_{\mathrm{on}}T_{\max}(1-2p_{\mathrm{N}})\mu_0/2).$$

The optimality condition is the same as (EC.2) with the exception of exchanging $T^*(P)$ with $r^*(P)$, exchanging r with $T_{\max}$, and an additional c_r on the right-hand side. The rest of the proof follows as the proof of Lemma EC.1 and is not reproduced here. □

C.4.4. Case IV: affine setup costs, no discounting, fixed horizon. If $c-\delta_{\mathrm{on}}(1-2p_{\mathrm{N}})\mu_0/2<0$ in this case, the effective cost of sampling is less than or equal to zero, and, given the trial is run, it is optimal to sample infinitely for any H. We, therefore, analyse the more interesting setting where $c-\delta_{\mathrm{on}}(1-2p_{\mathrm{N}})\mu_0/2\geq 0$. Notice that, unlike the previous cases, $c-\delta_{\mathrm{on}}(1-2p_{\mathrm{N}})\mu_0/2=0$ is considered here and gives interesting results that differ substantially from the scenario when $c-\delta_{\mathrm{on}}(1-2p_{\mathrm{N}})\mu_0/2>0$. Because both T^* and r^* grow unbounded with P (shown below), we let $T_{\max}=\infty$ and $r_{\max}=\infty$, i.e., there is no upper bound on the decision variables.

LEMMA EC.4. *If* $\rho=0$, $c_{cap}(r)=c_{fixed}+c_r r$, *where* $c_r>0$, $P(T)=\zeta(H-T-\mathbf{1}_{T>0}\Delta)$, $c-\delta_{\mathrm{on}}(1-2p_N)\mu_0/2\geq 0$, *and* $T_{\max}=\infty$ *and* $r_{\max}=\infty$, *then*

$$\lim_{H\to\infty}\frac{V^*}{\zeta(H-\Delta)}=(1-p_N)\sigma_0\Psi(-\mu_0/\sigma_0)+p_N\sigma_0\Psi(\mu_0/\sigma_0).$$

In addition, if $c-\delta_{\mathrm{on}}(1-2p_N)\mu_0/2\geq 0$, *then*

$$\lim_{H\to\infty}\frac{r^*}{(\zeta(H-\Delta))^{1/4}}=\lim_{H\to\infty}\frac{KT^*}{(\zeta(H-\Delta))^{1/4}}=\left(\frac{K^2\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{c+\delta_{\mathrm{on}}(1-2p_N)\mu_0/2}\right)^{1/4},$$

if $c-\delta_{\mathrm{on}}(1-2p_N)\mu_0/2>0$, *and*

$$\lim_{H\to\infty}\frac{r^*}{(\zeta(H-\Delta))^{1/3}}=\lim_{H\to\infty}\frac{KT^*}{(\zeta(H-\Delta))^{1/3}}=\left(\frac{K^2\sqrt{n_0\sigma_X^2}\phi(\mu_0/\sigma_0)}{c_r}\right)^{1/3},$$

if $c-\delta_{\mathrm{on}}(1-2p_N)\mu_0/2=0$, *where* $K=\sigma_X\zeta(\Psi(\mu_0/\sigma_0)+(1-p_N)\mu_0/\sigma_0)/(\sqrt{n_0}c_r)$.

Proof of Lemma EC.4. For large enough H, the optimal trial design satisfies $r^*T^*>0$ and the first order optimality conditions $\partial V(r^*,T^*)/\partial T=0$ and $\partial V(r^*,T^*)/\partial r=0$. Using the expressions derived in Appendix B and the property $\partial\sigma_Z/\partial r=(r/T)\partial\sigma_Z/\partial T$, the first order optimality conditions are equivalent to

$$(c-\delta_{\mathrm{on}}(1-2p_{\mathrm{N}})\mu_0/2)r^*+(1-p_{\mathrm{N}})\left[\zeta\alpha_{\mathrm{N}}(T^*)\left(1-\Phi\left(\frac{\alpha_{\mathrm{N}}(T^*)-\mu_0}{\sigma_Z^*}\right)\right)+\sigma_Z^*\zeta\Psi(\left(\frac{\alpha_{\mathrm{N}}(T^*)-\mu_0}{\sigma_Z^*}\right)\right] \quad \text{(EC.5)}$$
$$+p_{\mathrm{N}}\left[\zeta\alpha_{\mathrm{S}}(T^*)\left(1-\Phi\left(\frac{\alpha_{\mathrm{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right)+\sigma_Z^*\zeta\Psi(\left(\frac{\alpha_{\mathrm{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right]$$
$$=\zeta(H-T^*-\Delta)\frac{\sigma_Z^*}{dT}\left[(1-p_{\mathrm{N}})\phi\left(\frac{\alpha_{\mathrm{N}}(T^*)-\mu_0}{\sigma_Z^*}\right)+p_{\mathrm{N}}\phi\left(\frac{\alpha_{\mathrm{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right]$$

and

$$(c_r + (c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2)T^*)\frac{r^*}{T^*} = \zeta(H - T - \Delta)\frac{\partial \sigma_Z^*}{\partial T}\left[(1-p_{\text{N}})\phi\left(\frac{\alpha_{\text{N}}(T^*)-\mu_0}{\sigma_Z^*}\right) + p_{\text{N}}\phi\left(\frac{\alpha_{\text{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right], \tag{EC.6}$$

where σ_Z^* is σ_Z evaluated at (T^*, r^*). Notice that the right-hand side of both conditions are the same so we can formulate a third condition:

$$\begin{aligned}(1-p_{\text{N}})\left[\zeta\alpha_{\text{N}}(T^*)\left(1-\Phi\left(\frac{\alpha_{\text{N}}(T^*)-\mu_0}{\sigma_Z^*}\right)\right) + \sigma_Z^*\zeta\Psi(\left(\frac{\alpha_{\text{N}}(T^*)-\mu_0}{\sigma_Z^*}\right)\right] \\ + p_{\text{N}}\left[\zeta\alpha_{\text{S}}(T^*)\left(1-\Phi\left(\frac{\alpha_{\text{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right) + \sigma_Z^*\zeta\Psi(\left(\frac{\alpha_{\text{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right] = c_r\frac{r^*}{T^*}\end{aligned} \tag{EC.7}$$

We first show that $\lim_{H\to\infty} H - T^* - \Delta = \infty$ by means of contradiction. Assuming that $\lim_{H\to\infty} H - T^* - \Delta < \infty$ implies that $\lim_{H\to\infty} T^* = \infty$. Then, the right-hand side of (EC.5) approaches zero as $H \to \infty$. But the left-hand side is larger than zero so that (EC.5) cannot be satisfied for large enough H and we have reached a contradiction.

Next, we show that $\lim_{H\to\infty} T^* = \infty$ by contradiction. Assuming that $\lim_{H\to\infty} T^* < \infty$, (EC.5) implies that $\lim_{H\to\infty} r^* = \infty$. We reach a contradiction in (EC.7) because the left-hand side is bounded and the right-hand side is not.

Because we have established that $\lim_{H\to\infty} T^* = \infty$, (EC.7) implies $\lim_{H\to\infty} r^* = \infty$. In fact, (EC.7) implies that

$$\begin{aligned}r^* &= \left[\zeta(1-p_{\text{N}})\left[\alpha_{\text{N}}(T^*)\left(1-\Phi\left(\frac{\alpha_{\text{N}}(T^*)-\mu_0}{\sigma_Z^*}\right)\right) + \sigma_Z^*\Psi\left(\frac{\alpha_{\text{N}}(T^*)-\mu_0}{\sigma_Z^*}\right)\right]\right. \\ &\quad \left. + \zeta p_{\text{N}}\left[\alpha_{\text{S}}(T^*)\left(1-\Phi\left(\frac{\alpha_{\text{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right) + \sigma_Z^*\Psi\left(\frac{\alpha_{\text{S}}(T^*)+\mu_0}{\sigma_Z^*}\right)\right]\right]\frac{T^*}{c_r} \\ &\sim KT^*,\end{aligned} \tag{EC.8}$$

where $K = (\zeta\sigma_X\Psi(-\mu_0\sqrt{n_0}/\sigma_X)/(c_r\sqrt{n_0})$. The asymptotic equivalence follows due to $\lim_{H\to\infty} T^* = \infty$ and $\lim_{H\to\infty} r^* = \infty$.

The rest of the proof has to consider the two cases $c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 = 0$ and $c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 > 0$ separately.

Assume first that $c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 > 0$. Combining (EC.6) and (EC.8), we obtain

$$K(c_r + (c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2)T^*) \sim \zeta(H - T^* - \Delta)\sqrt{\frac{n_0\sigma_X^2 K}{(2n_0 + K(T^*)^2)^3}},$$

and with further simplifications

$$T^* \sim \left(\frac{\sqrt{n_0}\sigma_X\phi(\sqrt{n_0}\mu_0/\sigma_X)}{K^2(c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2)}\zeta(H-\Delta)\right)^{1/4}.$$

From (EC.8) it follows

$$r^* \sim \left(\frac{K^2\sqrt{n_0}\sigma_X\phi(\sqrt{n_0}\mu_0/\sigma_X)}{c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2}\zeta(H-\Delta)\right)^{1/4},$$

and we are done with the proof for the case $c - \delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2 > 0$.

Now, assume $c-\delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2=0$. Combining (EC.6) and (EC.8), we obtain

$$Kc_r \sim \zeta(H-T^*-\Delta)\sqrt{\frac{n_0\sigma_X^2 K}{(2n_0+K(T^*)^2)^3}},$$

and with further simplifications

$$T^* \sim \left(\frac{\sqrt{n_0}\sigma_X\phi(\sqrt{n_0}\mu_0/\sigma_X)}{K^2c_r}\zeta(H-\Delta)\right)^{1/3}.$$

From (EC.8) it follows

$$r^* \sim \left(\frac{K\sqrt{n_0}\sigma_X\phi(\sqrt{n_0}\mu_0/\sigma_X)}{c_r}\zeta(H-\Delta)\right)^{1/3},$$

and we are done with the proof for the case $c-\delta_{\text{on}}(1-2p_{\text{N}})\mu_0/2=0$.

Using the asymptotic behaviour of T^* and r^*, it is easy to show that

$$V^* \sim \zeta(H-\Delta)\left((1-p_{\text{N}})\sigma_0\Psi(-\mu_0/\sigma_0)+p_{\text{N}}\sigma_0\Psi(\mu_0/\sigma_0)\right). \quad \square$$

C.5. Proofs for asymptotic results in section 3.6.2

Proof of Prop. 6 It is easy to check that $\lim_{P\to\infty}P_\rho=\zeta/\rho$ with fixed patient pool and that $\lim_{H\to\infty}P_\rho(T)=\zeta/\rho$ with fixed horizon, where convergence is uniform on the bounded domain $0\le T\le T_{\max}$. Similarly, $\alpha'_{\text{N}}=I_{\text{N}}\rho/((1-p_{\text{N}})\zeta)=\lim_{P\to\infty}\alpha_{\text{N}}=\lim_{H\to\infty}\alpha_{\text{N}}(T)$, and $\alpha'_{\text{S}}=I_{\text{S}}\rho/(p_{\text{N}}\zeta=\lim_{P\to\infty}\alpha_{\text{S}}=\lim_{H\to\infty}\alpha_{\text{S}}(T)$. Consider first $rT>0$:

$$\begin{aligned}
|V_\infty(T,r)-V(T,r)| &= e^{-\rho(T+\Delta)}\sigma_Z\left|\frac{\zeta}{\rho}-P_\rho(T)\right|\left((1-p_{\text{N}})\left|\Psi\left(\frac{\alpha'_{\text{N}}-\mu_0}{\sigma_Z}\right)-\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\right|\right.\\
&\qquad\left.+p_{\text{N}}\left|\Psi\left(\frac{\alpha'_{\text{S}}+\mu_0}{\sigma_Z}\right)-\Psi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right|\right)\\
&\le \sigma_0\left|\frac{\zeta}{\rho}-P_\rho(T)\right|\left((1-p_{\text{N}})\left|\Psi\left(\frac{\alpha'_{\text{N}}-\mu_0}{\sigma_Z}\right)-\Psi\left(\frac{\alpha_{\text{N}}(T)-\mu_0}{\sigma_Z}\right)\right|\right.\\
&\qquad\left.+p_{\text{N}}\left|\Psi\left(\frac{\alpha'_{\text{S}}+\mu_0}{\sigma_Z}\right)-\Psi\left(\frac{\alpha_{\text{S}}(T)+\mu_0}{\sigma_Z}\right)\right|\right),
\end{aligned}$$

where the inequality uses $\sigma_Z\le\sigma_0$ and $e^{-\rho(T+\Delta)}\le 1$. Because $\Psi(\cdot)$ is a bounded function in the domain, there is a constant C, independent of T and r, such that

$$|V_\infty(T,r)-V(T,r)|\le C\left|\frac{\zeta}{\rho}-P_\rho(T)\right|.$$

Because $P_\rho(T)$ converges uniformly to ζ/ρ, it follows that $V(T,r)$ converges uniformly to $V_\infty(T,r)$ when $rT>0$. Consider now $rT=0$:

$$\begin{aligned}
|V_\infty(T,r)-V(T,r)| &= \left|\frac{\zeta}{\rho}\max\{0,\,(1-p_{\text{N}})\mu_0-I_{\text{N}}\rho/\zeta,\,-p_{\text{N}}\mu_0-I_{\text{S}}\rho/\zeta\}\right.\\
&\qquad\left.-P_\rho(T)\max\{0,\,(1-p_{\text{N}})\mu_0-I_{\text{N}}/P_\rho(T),\,-p_{\text{N}}\mu_0-I_{\text{S}}/P_\rho(T)\}\right|\\
&\le \left|\frac{\zeta}{\rho}-P_\rho(T)\right|\max\{0,\,(1-p_{\text{N}})\mu_0-I_{\text{N}}\rho/\zeta,\,-p_{\text{N}}\mu_0-I_{\text{S}}\rho/\zeta\}.
\end{aligned}$$

Again, because $P_\rho(T)$ converges uniformly to ζ/ρ, it follows that $V(T,r)$ converges uniformly to $V_\infty(T,r)$ when $rT=0$ and we are done proving that $V(T,r)$ converges uniformly to $V_\infty(T,r)$ as $P\to\infty$ or $H\to\infty$.

Due to uniform convergence, it follows that $\lim_{P\to\infty} V_{\mathrm{P}}^* = \lim_{H\to\infty} V_{\mathrm{H}}^* = \max_{T,r} V_\infty(T,r)$. We cannot guarantee a unique maximiser of $V_\infty(T,r)$. However, uniform convergence also guarantees that, if a unique maximiser $(T_\infty, r_\infty) = \arg\max_{T,r} V_\infty(T,r)$ exists, then $\lim_{P\to\infty} T_{\mathrm{P}}^* = \lim_{H\to\infty} T_{\mathrm{H}}^* = T_\infty$ and $\lim_{P\to\infty} r_{\mathrm{P}}^* = \lim_{H\to\infty} r_{\mathrm{H}}^* = r_\infty$. Thus, $\lim_{P\to\infty} Q_{\mathrm{P}}^* = \lim_{H\to\infty} Q_{\mathrm{H}}^* = r_\infty T_\infty/2$. □

Appendix D: Additional analysis of response-adaptive extensions

D.1. Some special cases which allow fully sequential trials

The sequential version of the one-stage optimal trial design problem in section 2.5, for the case of a fixed rate of recruitment r, was written in section 5.1 as having the objective function in (17). If we can show that this problem is equivalent (after suitable reparameterization) to an existing fully sequential trial, then special cases of the one-shot trials proposed in section 2.5 can be run as fully adaptive trials, at least for the case of a fixed rate of recruitment. In a fully sequential trial, the decision to continue the trial or to stop sampling is made after the outcome of each patient pair is observed, as for an optimal stopping time problem.

This section affirms that this can be done by reparameterizing (17) for several special cases to match the model of Chick et al. (2017). For these special cases, then, we can allow the duration of the trial T to be response adaptive to observed outcomes, for any given fixed rate of recruitment $r>0$, with a simple transformation of a few parameters. The scenarios which enable fully sequential trials here are aesthetically linked to the two special cases for comparative statics noted in the last paragraph of section 3.5.

To this end, we recall the objective function of the fully sequential trial in Chick et al. (2017).

$$V(\pi) = \mathbb{E}_\pi\left[\left\{\sum_{t=0}^{Q-1} \frac{-c+\delta_{\mathrm{on}} X_{t+1}}{(1+\tilde{\rho})^t}\right\} + \frac{\mathbf{1}_{\mathcal{D}=\mathrm{N}}(P_\rho W - I)}{(1+\tilde{\rho})^{\mathbf{1}_{Q>0}(Q+\tau)}}\right]. \tag{EC.9}$$

Here, the policy π is a nonanticipative function which selects an action a_t to either continue sampling, or to stop, on the basis of prior knowledge and the t samples observed so far, for $t=0,1,2,\ldots,Q$. After stopping and all data is observed, the new alternative is selected as best if $P_\rho W - I \geq 0$.

We first explore the differences between (17) and (EC.9) related to the fixed costs of the trial. The fixed cost $c_{\mathrm{cap}}(r)$ in (17) only affects the decision of running the trial or not, captured by action a_0. Such decision can also be made by computing the value of the optimal design with zero fixed costs, and then evaluating whether it overcomes the fixed costs. Thus, fixed costs do not disturb the equivalence of the optimal trial design once the decision to observe the first patient pair is taken, even though they do affect the optimal choice of a_0. We now explore equivalence of the sequential sampling models, assuming a_0 indicates the first sample is to be observed, for several special cases.

All the special cases in this subsection assume the fixed pool of post-adoption population, $P(T)=P$, and that the rate of recruitment is fixed. Preliminary results in this section were introduced in Alban et al. (2018). This subsection corrects a misstatement from that paper for a special case below. The case of adapting the rate of recruitment, rather than the number of patient pairs, is discussed in section 5.2.

D.1.1. Current practice is standard treatment ($p_{\mathrm{N}}=0$) When $p_{\mathrm{N}}=0$, (17) becomes

$$V(\pi)=-\mathbf{1}_{Q>0}c_{\mathrm{cap}}(r)+\mathbb{E}_{\pi}\left[\sum_{t=0}^{Q-1}\frac{\delta_{\mathrm{on}}X_{t+1}-2c}{(1+\tilde{\rho})^{t}}\right]+\mathbb{E}_{\pi}\left[\frac{\mathbf{1}_{\mathcal{D}=\mathrm{N}}(P_{\rho}W-I_{\mathrm{N}})}{(1+\tilde{\rho})^{\mathbf{1}_{Q>0}(Q+\tau)}}\right],$$

which is equivalent to (EC.9) by letting $c'=2c$ and $I'=I_{\mathrm{N}}$.

D.1.2. Forcing the decision to adopt one of the two technologies, undiscounted rewards When the adoption decision is forced to be N or S with undiscounted rewards[8] ($\tilde{\rho}=0$), we can use $\mathbf{1}_{\mathcal{D}=\mathrm{S}}=1-\mathbf{1}_{\mathcal{D}=\mathrm{N}}$ such that (17) becomes

$$V(\pi)=-\mathbf{1}_{Q>0}c_{\mathrm{cap}}(r)-p_{\mathrm{N}}P\mu_0-I_{\mathrm{S}}+\mathbb{E}_{\pi}\left[\sum_{t=0}^{Q-1}\delta_{\mathrm{on}}(1-2p_{\mathrm{N}})X_{t+1}-2c+\mathbf{1}_{\mathcal{D}=\mathrm{N}}(PW-(I_{\mathrm{N}}-I_{\mathrm{S}}))\right]$$

If $p_{\mathrm{N}}<1/2$, then we can divide through by $1-2p_{\mathrm{N}}$ and add the constant $P_{\rho}(Q/r)p_{\mathrm{N}}\mu_0-I_{\mathrm{S}}$ without changing the optimal π. Thus, the optimal π maximises:

$$\mathbb{E}_{\pi}\left[\sum_{t=0}^{Q-1}\delta_{\mathrm{on}}X_{t+1}-\frac{2c}{1-2p_{\mathrm{N}}}+\mathbf{1}_{\mathcal{D}=\mathrm{N}}\left(\frac{P}{1-2p_{\mathrm{N}}}W-\frac{I_{\mathrm{N}}-I_{\mathrm{S}}}{1-2p_{\mathrm{N}}}\right)\right],$$

which is in the form of (EC.9) when we let $c'=2c/(1-2p_{\mathrm{N}})$, $P'=P/(1-2p_{\mathrm{N}})$, and $I'=(I_{\mathrm{N}}-I_{\mathrm{S}})/(1-2p_{\mathrm{N}})$.

If $p_{\mathrm{N}}=1/2$, the optimal π maximises

$$\mathbb{E}_{\pi}\left[\sum_{t=0}^{Q-1}2c+\mathbf{1}_{\mathcal{D}=\mathrm{N}}(PW-(I_{\mathrm{N}}-I_{\mathrm{S}}))\right],$$

which is in the form of (EC.9) when $\delta'_{\mathrm{on}}=0$, $c'=2c$, $P'=P$, and $I'=I_{\mathrm{N}}-I_{\mathrm{S}}$.

D.1.3. No switching costs and undiscounted rewards When $I_{\mathrm{N}}=I_{\mathrm{S}}=0$ and rewards are not discounted ($\tilde{\rho}=0$), then it is optimal to adopt N or S, and the transformation in Appendix D.1.2 is also valid in this scenario.

D.1.4. No switching costs and $\mu_0=0$ When $I_{\mathrm{N}}=I_{\mathrm{S}}=\mu_0=0$, whether rewards are discounted or not, we obtain a special case of the above subsection that can additionally accommodate scenarios with discounted rewards. We again use $\mathbf{1}_{\mathcal{D}=\mathrm{S}}=1-\mathbf{1}_{\mathcal{D}=\mathrm{N}}$ in (17) to obtain

$$V(\pi)=-\mathbf{1}_{Q>0}c_{\mathrm{cap}}(r)+\mathbb{E}_{\pi}\left[\sum_{t=0}^{Q-1}\frac{\delta_{\mathrm{on}}(1-2p_{\mathrm{N}})X_{t+1}-2c}{(1+\tilde{\rho})^{t}}+\frac{\mathbf{1}_{\mathcal{D}=\mathrm{N}}(P_{\rho}W)}{(1+\tilde{\rho})^{\mathbf{1}_{Q>0}(Q+\tau)}}\right].$$

If $p_{\mathrm{N}}<1/2$, then we can divide through by $1-2p_{\mathrm{N}}$ and obtain the optimal π by solving (EC.9) with $c'=2c/(1-2p_{\mathrm{N}})$, $P'=P/(1-2p_{\mathrm{N}})$, and $I'=(I_{\mathrm{N}}-I_{\mathrm{S}})/(1-2p_{\mathrm{N}})$. If $p_{\mathrm{N}}=1/2$ we obtain the optimal π by solving (EC.9) with $\delta'_{\mathrm{on}}=0$, $c'=2c$, $P'=P$, and $I'=I_{\mathrm{N}}-I_{\mathrm{S}}$.

D.2. A forward-looking heuristic to optimise the recruitment rate in a response-adaptive design.

The idea of our heuristic is to use the one-shot design to inform the dynamic decisions by reevaluating the one-shot solutions after a relevant amount of information is gathered in the spirit of (Branke et al. 2007, Frazier et al. 2008). After recruiting patients for a fraction $1/B$ of the duration prescribed by the one-shot

[8] Alban et al. (2018) incorrectly claim that this results is also valid with discounted rewards.

design with the prescribed rate, we revisit the available data to adjust the recruitment rate and duration by solving again the one-shot design, this time running for a fraction $1/(B-1)$ of the newly prescribed duration at the newly prescribed rate, given that there are $B-1$ batches remaining. We continue in this manner, with batch $b+1$ running for fraction $1/(B-b)$ of the remaining samples prescribed when solving at time $\mathcal{T}_b$. We now give a detailed explanation of the heuristic leading to the proposed algorithm in Table EC.2.

The one-shot decisions made by our heuristic are very similar to the analysis in section 3 with minor differences in the objective and constraints because the dynamic decisions need to account for patient pairs in the pipeline and different setup cost structure. Define the present expected net gain at time $\mathcal{T}$ given information K, maximum recruitment rate $\bar{r}$, and patient pairs in the pipeline given by q, if the trial uses recruitment rate r for an additional duration T:

$$\begin{aligned}\tilde{V}(T,r,\mathcal{D};K,\mathcal{T},\bar{r},q) = &-(c_{\text{cap}}(r)-c_{\text{cap}}(\bar{r}))^+ - c\tilde{T}_\rho(T)r + \delta_{\text{on}}(\tilde{T}_\rho(T)r/2)(1-2p_{\text{N}})\mathbb{E}\left[W \mid K\right] \\ &+ e^{-\rho(T+\Delta)}\mathbb{E}\left[\mathbf{1}_{\mathcal{D}=\text{N}}((1-p_{\text{N}})P_\rho(\mathcal{T}+T)W - I_{\text{N}}) + \mathbf{1}_{\mathcal{D}=\text{S}}(-p_{\text{N}}P_\rho(\mathcal{T}+T)W - I_{\text{S}}) \mid K,T,r,q\right].\end{aligned} \tag{EC.10}$$

Notice that (EC.10), unlike (5b), only accounts for additional costs of increasing recruitment rate and the expectation is conditioned on having q patient pairs in the pipeline. The analog analysis of section 3 yields the following closed-form equation:

$$\begin{aligned}\tilde{V}(T,r;K,\mathcal{T},\bar{r},q) = &-(c_{\text{cap}}(r)-c_{\text{cap}}(\bar{r}))^+ - c\tilde{T}_\rho(T)r) + \delta_{\text{on}}(\tilde{T}_\rho(T)r/2)(1-2p_{\text{N}})\mu \\ &+ e^{-\rho(T+\Delta)}P_\rho(\mathcal{T}+T)\tilde{\sigma}_Z\left[(1-p_{\text{N}})\Psi\left(\frac{\alpha_{\text{N}}(\mathcal{T}+T)-\mu}{\tilde{\sigma}_Z}\right) + p_{\text{N}}\Psi\left(\frac{\alpha_{\text{S}}(\mathcal{T}+T)+\mu}{\tilde{\sigma}_Z}\right)\right],\end{aligned} \tag{EC.11}$$

where $\tilde{\sigma}_Z = \sigma_X^2(Tr/2+q(\Delta))/(n(n+Tr/2+q(\Delta)))$.

The algorithm in Table EC.2 shows how to use (EC.11) to make dynamic decisions using our forward-looking heuristic. We present it for completeness to show how the one-shot value-based trial can be extended in a natural way to batch sequential trials which adapt both the recruitment rate and duration of sampling. We do not present numerical examples for this case, because the focus is on the one-shot trials, analysis and extended numerical application above, and because adaptive durations for fixed rates have been discussed by (Chick et al. 2017, Alban et al. 2018) and extended in section 5.1 above, and this groundwork lays the work for future application-oriented batch sequential trials work.

Appendix E: Computation of CPCS and power

In section 3.5, we defined the CPCS as the conditional probability of correct selection of a technology for adoption, given a specific value of $W = w$:

$$\text{CPCS}(w) = \Pr(\mathcal{D}^* = \mathcal{D}^{orac} \mid W = w),$$

where $\mathcal{D}^{orac}$ is the oracle's adoption decision. In this section, we show how to compute it.

First of all, note that the optimal adoption decision depends on Z_{Tr}, which given W has the following distribution:

$$Z_{Tr} \mid W \sim \mathcal{N}\left(\frac{n_0\mu_0 + rTW/2}{n_0 + rT/2}\ ,\ \frac{2rT\sigma_X^2}{(2n_0+rT)^2}\right).$$

Table EC.2 Heuristic for adapting recruitment rate and trial duration in B batches motivated by expected value of information / knowledge gradient sampling approach.

1. Initialise: Set K to be the prior distribution for unknown W, $\mathcal{T}=0$, $\bar{r}=0$, $q=0$.
2. For $b=0,1,\ldots,B-1$ batches of patient pairs.
   (a) Find T^*, r^* to optimise (EC.11).
   (b) Allocate samples for batch $b+1$ at rate r^* for duration $T^*/(B-b)$.
   (c) Update next decision time $\mathcal{T} \leftarrow \mathcal{T}+T^*/(B-b)$, the maximum rate $\bar{r} \leftarrow \max\{\bar{r}, r^*\}$, the information set K given prior and data observed while batch $b+1$ samples were being allocated, the number of patient pairs q in pipeline at the end of sampling for batch $b+1$.
3. Wait for remaining outcomes to be observed, select best alternative $\mathcal{D}$ on the basis of all information.

If $W > \alpha_{\mathrm{N}}(T)$, the oracle adoption decision is N. Thus, the CPCS simplifies to

$$\mathrm{CPCS}(w) = \Pr(Z_{Tr} > \alpha_{\mathrm{N}}(T) \mid W = w).$$

We can then compute $\mathrm{CPCS}(w) = 1-\Phi(U_{\mathrm{N}})$, where

$$U_{\mathrm{N}} = \frac{2n_0(\alpha_{\mathrm{N}}(T)-\mu_0) + rT(\alpha_{\mathrm{N}}(T)-w)}{\sigma_X\sqrt{2rT}}.$$

If $w < -\alpha_{\mathrm{S}}(T)$, the oracle's adoption decision is S, and we obtain $\mathrm{CPCS}(w) = \Pr(Z_{Tr} < -\alpha_{\mathrm{S}}(T) \mid W = w) = 1-\Phi(U_{\mathrm{S}})$, where

$$U_{\mathrm{S}} = \frac{2n_0(\alpha_{\mathrm{S}}(T)+\mu_0) + rT(\alpha_{\mathrm{S}}(T)+w)}{\sigma_X\sqrt{2rT}}.$$

Similarly, if $-\alpha_{\mathrm{S}}(T) \le w \le \alpha_{\mathrm{N}}(T)$, the oracle adoption decision is M, and we obtain $\mathrm{CPCS}(w) = \Pr(-\alpha_{\mathrm{S}}(T) \le Z_{Tr} \le \alpha_{\mathrm{N}}(T) \mid W = w) = \Phi(U_{\mathrm{N}}) + \Phi(U_{\mathrm{S}}) - 1$.

The final closed-form expression of CPCS is then

$$\mathrm{CPCS}(w) = \begin{cases} 1-\Phi(U_{\mathrm{N}}), & w > \alpha_{\mathrm{N}}(T) \\ 1-\Phi(U_{\mathrm{S}}), & w < -\alpha_{\mathrm{S}}(T) \\ \Phi(U_{\mathrm{N}}) + \Phi(U_{\mathrm{S}}) - 1, & -\alpha_{\mathrm{S}}(T) \le w \le \alpha_{\mathrm{N}}(T) \end{cases}$$

The computation of power requires the following definitions. Let α be the type I error and q_x be the $1-x$ quantile of a standard normal distribution. A two-sided hypothesis test, rejects the (null) hypothesis that the population mean W is zero if the sample average $\bar{X} = 2/(rT)\sum_{i=1}^{rT/2} X_i$ is larger than $q_{\alpha/2}\sigma_X/\sqrt{rT/2}$, or smaller than $-q_{\alpha/2}\sigma_X/\sqrt{rT/2}$. The power is the probability of rejection given $W = w$:

$$\begin{aligned} \mathrm{power}(w) &= \Pr\left(\bar{X} > \frac{\sigma_X}{\sqrt{rT/2}}q_{\alpha/2} \cup \bar{X} < -\frac{\sigma_X}{\sqrt{rT/2}}q_{\alpha/2} \,\middle|\, W = w\right) \\ &= \Pr\left(\bar{X} > \frac{\sigma_X q_{\alpha/2}}{\sqrt{rT/2}} \,\middle|\, W = w\right) + \Pr\left(\bar{X} < -\frac{\sigma_X q_{\alpha/2}}{\sqrt{rT/2}} \,\middle|\, W = w\right) \\ &= 1-\Phi\left(q_{\alpha/2} - \sqrt{\frac{rt}{2\sigma_X^2}}w\right) + 1-\Phi\left(q_{\alpha/2} + \sqrt{\frac{rt}{2\sigma_X^2}}w\right). \end{aligned}$$

References

Alban A, Chick SE, Forster M (2018) Extending a Bayesian decision-theoretic approach to value-based sequential clinical trial design. Rabe M, Juan A, Mustafee N, Skoogh A, Jain S, Johansson B, eds., *Proceedings of the 2018 Winter Simulation Conference*, 2459–2470 (Piscataway, NJ: IEEE, Inc.).

Andreasson N, Evgrafov A, Patriksson M (2007) *Introduction to Continuous Optimization* (Studentlitteratur AB).

Branke J, Chick S, Schmidt C (2007) Selecting a selection procedure. *Management Science* 53(12):1916–1932.

Chick SE, Forster M, Pertile P (2017) A Bayesian decision-theoretic model of sequential experimentation with delayed response. *Journal of the Royal Statistical Society, Series B* 79(5):1439–1462.

Frazier PI, Powell WB, Dayanik S (2008) A knowledge-gradient policy for sequential information collection. *SIAM Journal on Control and Optimization* 47(5):2410–2439.