



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

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

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Availability:

This version is available at: <https://hdl.handle.net/11585/879132> since: 2022-11-16

Published:

DOI: <http://doi.org/10.1287/mnsc.2022.4540>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Andres Alban, Stephen E. Chick, Martin Forster. (2022). “Value-based clinical trials: selecting recruitment rates and trial lengths in different regulatory contexts”. *Management Science*, pp. 1-20.

The final published version is available online at:

<https://doi.org/10.1287/mnsc.2022.4540>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Value-based clinical trials: selecting recruitment rates and trial lengths 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 Statistical Sciences “Paolo Fortunati”, University of Bologna, Bologna, Italy and Department of Economics
and Related Studies, University of York, York, United Kingdom martin.forster@unibo.it

Health systems are placing increasing emphasis on improving the design and operation of clinical trials, with the aim of making the health technology adoption process more *value-based*. We present a model of a value-based two-armed clinical trial in which both the recruitment rate and trial length are optimized. The model is value-based because it balances the cost of the trial with the expected benefit it generates for patients, valued by the relative health benefits and costs of the technologies. We consider a wide range of regulatory and practical contexts which address how patient health is valued (discount rate, time horizon, pragmatic trials). We present comparative statics and asymptotic analysis, together with a retrospective application to a recent health technology assessment, and an extension for adaptive trials. Results challenge traditional perceptions concerning the efficiency, length, and knowledge that may be gained from clinical research for trial managers or funders charged with delivering *value* efficiently: we highlight trade-offs between trial costs and population health benefits influenced by trial outcomes and the importance of optimizing both recruitment rate and trial duration rather than sample size alone.

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

History: Submitted: 17 Apr 2020; revised: 26 Aug 2021 and 14 Feb 2022; Accepted 1 March 2022

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 health technologies (HTs) and growing pressure on their budgets. A particular concern surrounds the evaluation of the clinical efficacy, effectiveness, and cost-effectiveness of HTs. With a large body of evidence suggesting that there is a productivity crisis in biopharmaceutical research and development ([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 trials.

In this paper, we present a value-based, Bayesian decision-theoretic model of a two-armed clinical trial for HT adoption, with the goal of improving the efficiency of the HT innovation process by optimizing both the recruitment rate and trial duration within different regulatory contexts.

This topic is timely and important. In the United Kingdom (UK), the National Institute for Health Research (NIHR) is funding programs that innovate methodological designs for operating trials more efficiently (NIHR 2018, 2020). In the United States, the National Institutes of Health have launched a drive to improve the quality and ‘efficiency, accountability and transparency’ in clinical research (Hudson et al. 2016). Recently, the European Union (2021) introduced the ‘Clinical Trial Regulation’, which seeks to harmonize the process of assessment and supervision of trials, to increase their transparency and efficiency.

Efficiency in clinical research is not consistently defined. On the one hand, it may be understood in terms of requiring fewer enrolled patients to identify the smallest clinically relevant difference, while maintaining the predefined type I and type II error rates, or in terms of reducing the research cost per patient. On the other hand, efficiency in HT adoption is commonly understood in terms of adopting the most cost-effective technology (NICE 2014). The adoption process compares the expected health benefits net of treatment costs of one technology with another. In this paper, we view efficiency from the perspective of clinical trial operation and HT adoption. We model a so-called value-based trial, whose objective is to maximize health benefits obtained for the patient population that may benefit from the technology adoption, accounting for the costs of running the trial and switching to the chosen HT, as well as the costs of the technologies themselves. Here, health benefits are measured by converting health outcomes to monetary value – one example being a standard conversion of quality-adjusted life years (QALYs) (Sharma et al. 2021). Thus, efficiency here addresses misalignment between objectives for trial design (minimize cost) and health economics (maximize cost effectiveness) by aligning a trial’s design with HT adoption decisions.

Our framework maximizes the value of the trial for the full patient population that may benefit from the technologies under consideration, including those enrolled in the trial, those not enrolled during the trial period, and those in the post-adoption population. We consider pragmatic trials, which compare two technologies that are both in use by some health providers before the trial commences, though explanatory trials, which compare a new technology with a known standard, and which can fit the model if the trial population resembles the population to benefit from the adoption decision. We also consider different regulatory contexts according to the discounting of costs and health outcomes, the concern for enrolled patient outcomes (online or offline learning), and the definition of the post-adoption horizon (one in which the number of patients affected by the adoption decision is fixed, the other in which it is a function of the trial’s length, which can occur when the right to market a new drug is protected by a patent). Because the value of a trial can be time-sensitive in certain regulatory contexts, it may be insufficient to optimize over the sample size alone (the usual approach taken in clinical trial design). We therefore optimize both

the recruitment rate and trial length, and show how doing so can create more value, in expectation, for the health care system.

In comparison with exploration-exploitation methods in other applications, several specific features of the problem make it interesting. Firstly, the regulatory context influences the valuation of benefits accruing to the patient population and, in turn, the expected value of information obtained from patients enrolled in the trial. Secondly, the recruitment capacity may be costly and nonlinear. Hence, the problem is richer than merely selecting an optimal sample size. Thirdly, there may be significant delays between the time a trial participant receives treatment and the time that health outcomes and treatment costs are observed. By addressing these matters, the model assists diverse groups of decision-makers in addressing a range of questions: (i) Clinical trial managers: What is the optimal recruitment rate for the trial? How long should the trial run? (ii) Funders of trials: In assessing several proposed trials, which have higher expected health benefit? Is a trial worth running? What level of resourcing is necessary? (iii) Public sector policymakers: How do the aforementioned regulatory issues affect the value of optimal trial designs? More broadly, our analysis may be of general interest in exploration-exploitation contexts where the value of an exploitation decision is time sensitive, or learning costs are nonlinear in the sampling rate. Our results challenge traditional perceptions concerning the efficiency, length, and knowledge that may be gained from clinical research for trial managers or funders charged with delivering value. We note the importance of optimizing both recruitment rate and trial duration rather than sample size alone from a social planner's perspective.

We review background literature in section 1 and set up the model in section 2. Section 3 provides structural results to characterize the efficient trial, presents comparative statics analysis and asymptotic results which shed insights about the trial's sample size. Section 4 applies our framework to a pragmatic trial from the UK. Section 5 shows how the trial length can be adapted sequentially, as patient outcomes are observed. Section 6 discusses our main results and section 7 concludes. The Online Companion provides supporting material.

1. Background literature

The classical approach to clinical trial design selects the sample size and null hypothesis rejection region to guarantee a type I error and a type II error for the smallest relevant difference worth detecting (Lachin 1981). Claxton and Posnett (1996) and Claxton (1999) criticize this approach because it ignores economic principles, such as the value of sample information and data acquisition costs. They propose, as an alternative, a decision-theoretic approach using rules from cost-effectiveness analysis. Many contributions which advocate a decision-theoretic approach have followed (e.g., Claxton et al. 2000, Gittins and Pezeshek 2000, Eckermann and Willan 2007, Griffin

et al. 2010, Draper 2013, Montazerhodjat et al. 2017). These compare the cost of collecting information with the value that it generates, using so-called value of information calculations (Raiffa and Schlaifer 1961, Kunst et al. 2019).

We contribute to this stream of literature in several ways. Firstly, we show that it is important to model the recruitment rate and trial duration, not just the sample size. Secondly, we consider these decisions in a wide range of regulatory contexts. Thirdly, as well as addressing trials which compare a new technology which is not in use in a health system with a known standard, our model also addresses trials which address mixed clinical practice (when two HTs are used to treat similar patients in ongoing medical practice), a scenario commonly encountered in clinical research, e.g., use of caesarean section (Betrán et al. 2016) and fracture treatments (Costa et al. 2014). Fourthly, we provide comparative statics and asymptotic analysis.

The majority of decision-theoretic models above concern one-shot trials. More recently, interest has grown in adaptive approaches (Pertile et al. 2014, Ahuja and Birge 2016, Williamson et al. 2017, Chick et al. 2017, 2021, Villar et al. 2018, Villar and Rosenberger 2018, Pallmann et al. 2018, Anderer et al. 2019). The related Bayesian ranking and selection literature has proposed various combinations of discounted and undiscounted rewards and online and offline learning (Branke et al. 2007, Frazier et al. 2008, Ryzhov et al. 2012, Xie et al. 2016). Although our principal focus is on one-shot designs, we also discuss extensions for adaptive trials in section 5.

Although it is natural for a Bayesian approach to maximize expected value, or to minimize 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 0-1 trials (Williamson et al. 2017). We use a Bayesian, value-based framework to be consistent with the UK's NICE (2014) guidance for uncertainty quantification for probabilistic sensitivity analysis for HT assessments and to give an average-case rather than a frequentist worst-case analysis (Inoue and Chick 1998, Robert 2007). Although Bayesian in nature, simulations of our model can give power curves, in line with US FDA (2019) guidance regarding the communication of complex and innovative trial designs.

There exists a range of other approaches to related issues in clinical trial design. Some consider changing the balance of allocation to treatment arms as a function of the past history of allocations (Berry and Eick 1995, Ahuja and Birge 2016). Others maintain a balanced allocation, but allow for the trial to stop at any stage as evidence accumulates (Berry and Ho 1988, Jennison and Turnbull 1989, Hampson and Jennison 2013, Pertile et al. 2014, Chick et al. 2017). We focus here on balanced allocation in a two-armed trial as it is the most common type of trial in practice. We assess the two

arms for a given population, rather than focusing on determining which subpopulation benefits the most from a given treatment as is the case in precision medicine.

It is becoming more common to carry out an economic evaluation of a health technology in addition to assessing its effectiveness in a clinical trial. [Drummond et al. \(2015\)](#) note that economic evaluations are most common in jurisdictions where a single authority is charged with making reimbursement decisions for new health technologies. They note that approximately half of the countries in the European Union, together with Canada, Australia and New Zealand, requested the economic evaluation of pharmaceutical products and, sometimes, other health technologies (the ProFHER case study described below is an example). [Panteli et al. \(2015\)](#) review evidence-based approaches to the reimbursement and pricing of pharmaceutical products in 36 European countries and find that measurement of both effectiveness and cost is important in the vast majority: 19 countries considered costs in their evidence-based decision-making process and 15 more of them considered costs on a case-by-case basis ([Panteli et al. 2015](#), Table S1). [Emanuel et al. \(2020\)](#) review the purchasing of pharmaceutical products in the United States, Australia, France, Germany, Norway, Switzerland and the UK. They find that all countries, except the United States, have centralized state-level mechanisms in place that are designed to lower the prices of drugs that have received marketing authorization.

To the best of our knowledge, the value-based framework that takes into account research costs is not used to design trials today. However, as already noted, proposals to adopt such a framework date to [Claxton and Posnett \(1996\)](#) and [Claxton \(1999\)](#) and growing interest in application of the framework has led to the publication of two ‘good practice’ guides ([Fenwick et al. 2020](#), [Rothery et al. 2020](#)). Implementing a value-based trial in practice requires collecting cost and QALY data (or other health outcome data which can be converted to monetary values). While many clinical trials do not have QALYs as a primary endpoint or monitor treatment costs, accounting for QALYs in clinical trials has seen increasing attention ([Angus et al. 2001](#), [Ferguson et al. 2013](#), [NICE 2014](#)), as has the use of health economic criteria in informing trial design ([Draper 2013](#), [Flight et al. 2019](#), [NIHR 2020](#)). [Mitchell and Patterson \(2020\)](#) review clinical trials registered on [clinicaltrials.gov](#) that studied a drug and or a biological during 2004–2018 to establish the proportion that collected economic endpoints (broadly defined as those recording resource utilization/costs). Of the 104,885 trials identified, 1,437 (1.37%) included economic endpoints and for phase 2/3 trials, 939 (2.54%) did so. The proportion increased from 1.2% in 2004—2008 to 1.6% in 2014—2018. Costs and QALYs are also part of the broader operations research literature on resource planning for health interventions ([Kaplan and Brandeau 1994](#), [Long et al. 2008](#)).

The societal perspective of value that is adopted in this work contrasts with Bayesian decision-theoretic contributions which consider a trial’s optimal choice of the sample size from an industry

perspective (Gittins and Pezeshk 2000, Willan 2008). There, the terminal reward is a function of the probability of approval by the regulator and the expected market share. 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. Kouvelis et al. (2017) further link data from trials to a theoretical model of recruitment rate optimization to improve a private firm’s profit. Rojas-Cordova and Bish (2018) also take a for-profit firm perspective in optimizing group-adaptive trials with 0-1 outcomes. Our work focuses on uncertainties related to the efficacy (incremental health benefits and costs) of a HT adoption decision relative to the cost of research, rather than on uncertainties in patient accrual, and models the cost of effort to obtain a given recruitment rate on an empirical basis. Bravo et al. (2018) study how the arrival stream of new technologies might influence the choice of statistical threshold for a trial, and endogenize QALY gains as part of optimizing a trial. Their insight about expected QALYs gained following a technology adoption decision can inform our choice of the number of patients to benefit from the decision. A public-private model setting raises issues about contracting and incentive alignment between technology adopter and provider: our model bypasses this issue by focusing on the vast body of clinical research funded by academic and regulatory organizations (NIHR 2019).

2. Mathematical model of a value-based clinical trial

We present a Bayesian decision-theoretic model of a clinical trial comparing two HTs on cost-effectiveness grounds. The objective is to choose both the recruitment rate and recruitment duration so as to maximize the monetary value of health benefits generated for the target population, minus the cost of carrying out the trial and any costs incurred in technology adoption. This section discusses decision variables, outcome measures, the objective function, and regulatory contexts addressed by the model. Appendix A summarizes the principal notation.

2.1. Trial design and decision variables

A clinical trial compares two HTs – N (*new*) and S (*standard*) – on cost-effectiveness grounds. We consider standard two-arm trials, in which no patients are treated with N before commissioning (such as explanatory trials), as well as trials in which there is mixed practice (such as many pragmatic trials). Define $p_N \in [0, 1/2]$ as the fraction of the eligible population that is treated with N prior to the trial’s inception, so that the fraction $1 - p_N$ is treated with S. Without loss of generality, assume that a greater proportion of patients receive S (one can rename the two HTs to satisfy this constraint).

The trial comprises three decisions: the adoption decision, \mathcal{D} , the recruitment duration, T , and the recruitment rate, r . The adoption decision $\mathcal{D} \in \{S, N, M\}$ is made at trial conclusion with the information obtained in the trial and determines the treatment for post-adoption patients. If $\mathcal{D} = S$,

all patients are treated with the standard HT; if $\mathcal{D} = \text{N}$, all patients are treated with the new HT; $\mathcal{D} = \text{M}$ implies that patients are treated according to the mix that was in place prior to the trial.

The trial design decisions made at the start of the trial are T and r . The recruitment duration, T , is constrained to lie in the interval $[0, T_{\max}]$. Regulatory or funding requirements might constrain T_{\max} . The recruitment rate, r , is constrained to lie in $[0, r_{\max}]$. The maximum 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. We assume that patients are randomized evenly to the two arms of the trial in pairs. Let $Q = Tr/2$ denote the sample size (of patient pairs, not of patients). In practice we may require $Q \in [0, Q_{\max}]$ for $Q_{\max} \leq T_{\max}r_{\max}/2$, but this constraint is implicit to simplify exposition.

2.2. Value-based criterion for a HT 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 monetary net benefit of using treatment N over S – the *incremental net monetary benefit* (INMB, [Fenwick et al. 2020](#)):


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

where $E_j \in \mathbb{R}$ and $C_j \in \mathbb{R}$ measure expected effectiveness (health or quality of life improvement) and cost of HT $j \in \{\text{N}, \text{S}\}$, 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 ([Sharma et al. 2021](#)). We use £20,000/QALY, consistent with the National Institute for Health and Care Excellence in England and Wales ([NICE 2013](#)). 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, a pilot study, or expert opinion. Observations of the INMB of each patient pair 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: (i) running the trial and treating post-trial patients with the HT recommended by its results, and (ii) 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, where the trial starts recruitment at time zero. The vertical axis plots the incidence and recruitment rates of patients. Areas labeled are: (A) patients recruited to the trial; (B) patients not recruited to the trial during the recruitment period; (C) patients in the waiting period; and (D) post-adoption patients treated with the recommended HT.

Figure 1 Each area represents the number of patients in the following classes: (A) recruited to the trial; (B) not recruited to the trial during the recruitment period; (C) treated with current practice during waiting period; (D) benefiting from the adoption decision.



2.3.1. Patients recruited to the trial. During the trial, $Tr/2$ pairs of patients are randomized. For patient pair $i \in 1, \dots, Q_{\max}$, we can observe a noisy observation X_i of INMB. The trial incurs a variable cost, c , per patient. We assume that $X_i | W \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(W, \sigma_X^2)$, where σ_X^2 is the known sampling variance. The proportion $(1 - 2p_N)/2$ is the fraction of patients receiving treatment N under randomization who would have received S had the trial not taken place. Thus, the expected net gain for the Tr patients in the trial, net of per-patient costs, is $Tr((1 - 2p_N)/2)\mathbb{E}[W] - cTr$.

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 treatment 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 receive treatment according to the practice in place before the start of the trial. 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 HT adoption. We assume that $I_M = 0$ because neither HT is adopted, and $I_N, I_S \geq 0$. Let P be the number of post-adoption patients, assumed to be known at the start of the trial. If $\mathcal{D} = N$, the expected net gain for these patients is $P(1 - p_N)\mathbb{E}[W] - I_N$, where $P(1 - p_N)$ is the number of patients who, absent the trial, would be treated with HT S. If $\mathcal{D} = S$, the expected net gain is $-Pp_N\mathbb{E}[W] - I_S$. If $\mathcal{D} = M$, the expected net gain is zero because the trial did not change practice. The expected net gain for this portion of the population is $\mathbb{E}[\mathbf{1}_{\mathcal{D}=N}(P(1 - p_N)W - I_N) + \mathbf{1}_{\mathcal{D}=S}(-Pp_NW - I_S) | T, r]$, where $\mathbf{1}_F$ is the indicator function, equal to one if F is true and zero otherwise. \mathcal{D} is a random variable because it depends on the yet unobserved patients. We condition on T and r to clarify that the expectation depends on the trial design because \mathcal{D} is determined by the outcomes of $Tr/2$ patient pairs.

2.3.5. Trial setup costs. The setup and operation of a trial incur costs. We define a setup cost, $c_{\text{cap}}(r)$, to model the cost of initiating a trial (including patient outreach and opening recruitment sites), retention efforts, data analysis, and dissemination of results. For example, $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r^2$ is a setup cost model with a fixed cost and an increasing marginal cost of capacity. This cost is assumed to be independent of the number of patients in the trial. Including variable per-patient costs for patients completing the trial, the total trial cost is $c_{\text{cap}}(r) + cTr$.

Trials may have issues with recruitment and retention (Walters et al. 2017). Mechanisms to increase recruitment and retention, with a view to achieving a defined recruitment target, can relate to protocol (refine, simplify informed consent), sites (replace marginally performing sites, add sites, regularly communicate with sites), physicians (professional incentives), patient engagement (outreach channels, convenience, explanations) and so on (Achilleas et al. 2010, Bower et al. 2014, Walters et al. 2017, Lock 2019). With this in mind, r and c are assumed to be modeled as per patient that delivers data, and $c_{\text{cap}}(r)$ models the cost to achieve a given average rate r .

2.3.6. Expected net gain. Define the expected net gain of a trial design by $V(T, r)$. If a trial is not run, the post-adoption population is the entire population and the expected net gain is:

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

If the trial recruits at least one pair of patients, then $T, r > 0$ and the expected net gain is the sum of the expected net gains of the enrolled and post-adoption patients, minus the cost of the trial:

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

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

2.4. The regulatory context

We extend the model to handle two more issues concerned with the regulatory context: the size of the population affected by the HT adoption decision and whether rewards are discounted or not.

¹ Online learning may be particularly relevant in trials for rare diseases because the enrolled patients represent a significant fraction of the patient population.

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 non-increasing in T . We model two scenarios:

1. There is 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 $P(T) = \zeta(H - T - \mathbf{1}_{T > 0}\Delta)$.

2. There is a *fixed pool* of patients, $P(T) = P$, that is not dependent on the length of the trial.

For new drug approvals, one example of the first scenario is the standard 20-year period of patent protection that applies to new pharmaceutical products for the 159 countries which make up the [World Trade Organization \(2006\)](#). Examples of the second scenario include jurisdictions which grant exclusive marketization rights for a HT for a defined period post-authorization ([US FDA 2015](#)), often because it is argued that the period granted by patent protection is not long enough.² For health technologies more widely, we note that reimbursement decisions can use review horizons that are evidence-specific. For example, the National Institute for Health and Care Excellence in England and Wales sets a time for review which varies according to both the available evidence and knowledge about ‘when ongoing research will be reported’ ([NICE 2018](#)).

2.4.2. Discounted expected net gain. The preceding assumes that costs and benefits during the trial are not discounted, an assumption which may be realistic for some jurisdictions but not others.³ With continuous time discount rate $\rho > 0$, define the discounted recruitment period duration as $\tilde{T}_\rho(T)$ and the discounted number of post-adoption patients as $P_\rho(T)$, under the assumption that $P(T)$ patients arrive at a constant rate, ζ , over a duration of time $P(T)/\zeta$:

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

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

We can handle discounting by substituting $P(T)$ with $P_\rho(T)$ and T with $\tilde{T}_\rho(T)$ in (2) and (3):

$$V(T, 0) = V(0, r) = \mathbb{E}[\mathbf{1}_{\mathcal{D}=\text{N}}(P_\rho(0)(1 - p_{\text{N}})W - I_{\text{N}}) + \mathbf{1}_{\mathcal{D}=\text{S}}(-P_\rho(0)p_{\text{N}}W - I_{\text{S}}) \mid T, r], \quad (4a)$$

$$V(T, r) = -(c_{\text{cap}}(r) + c\tilde{T}_\rho(T)r) + \delta_{\text{on}}(\tilde{T}_\rho(T)r/2)(1 - 2p_{\text{N}})\mathbb{E}[W] \\ + e^{-\rho(T+\Delta)}\mathbb{E}[\mathbf{1}_{\mathcal{D}=\text{N}}((1 - p_{\text{N}})P_\rho(T)W - I_{\text{N}}) + \mathbf{1}_{\mathcal{D}=\text{S}}(-p_{\text{N}}P_\rho(T)W - I_{\text{S}}) \mid T, r], \quad (4b)$$

where (4b) applies when the trial runs ($Tr > 0$).

² Marketing authorization agreements vary by country and the nature of the health technology; to provide an exhaustive list is beyond the scope of this paper. We note here some prominent examples, which illustrate the variety of agreements in place for new medicines: [Rome et al. \(2021\)](#) estimate that, in the United States, market exclusivity lasts between 13 and 17 years; [Lexchin \(2021\)](#) estimate a period of approximately 1.5 years in Canada; the 27 states of the European Union operate a period of ten years for orphan medicines.

³ Although traditional clinical trials do not tend to discount health benefits and costs – focusing as they do on the estimation of average treatment effects for clinical outcomes – economic evaluations do. For example, in England and Wales, the NICE ‘reference case’, which defines the baseline methodology for health technology assessment, requires that a 3.5% discount rate be used to discount future streams of costs and benefits, with sensitivity analysis using a rate of 1.5% ([NICE 2013](#), Sec. 5.6). A recent review of the health care systems in 31 countries found that all of them discounted costs and benefits, with 90% recommending sensitivity analysis ([Sharma et al. 2021](#)).

2.5. One-shot optimal trial design problem

Our focal optimal value-based trial design problem is

$$\begin{aligned} V^* &= \max_{T,r} V(T,r) \\ \text{s.t. } & T \in [0, T_{\max}], r \in [0, r_{\max}], \end{aligned} \quad (5)$$

where $V(T,r)$ is defined by (4a) when the trial does not run and is otherwise defined by (4b). The expectations which determine $V(T,r)$ are with respect to the prior distribution for W , as described in section 2.2, and the likelihood for observations 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 identify the optimal adoption decision. Next, we explore optimal values for the recruitment duration T^* and rate r^* . Some mild assumptions guarantee the existence of a solution. We then provide comparative statics and explore the asymptotics of optimal trial designs.

3.1. Optimal adoption decision

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

$$Z_{Tr} \equiv \mathbb{E} [W \mid X_1, \dots, X_{Tr/2}].$$

Let $n_0 = \sigma_X^2/\sigma_0^2$ be the effective sample size of the prior distribution and define $\sigma_Z^2 = \sigma_X^2(Tr/2)/[n_0(n_0 + Tr/2)]$. Then it can be shown that (DeGroot 1970)

$$\begin{aligned} Z_{Tr} &= \frac{n_0\mu_0 + \sum_{i=1}^{Tr/2} X_i}{n_0 + Tr/2} \sim \mathcal{N}(\mu_0, \sigma_Z^2), \\ \text{and } W \mid Z_{Tr} &\sim \mathcal{N}(Z_{Tr}, \sigma_X^2/(n_0 + Tr/2)). \end{aligned} \quad (6)$$


Thus, Z_{Tr} is a sufficient statistic for all the information obtained in the trial, and \mathcal{D} is solely determined by Z_{Tr} instead of the sequence X_i of all observations.

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

$$\mathcal{D} = \arg \max \left\{ \underbrace{0}_{\text{M}}, \underbrace{P_\rho(T)(1-p_N)Z_{Tr} - I_N}_{\text{N}}, \underbrace{-P_\rho(T)p_N Z_{Tr} - I_S}_{\text{S}} \right\},$$

where the terms inside the arg max operator are the post-adoption expected net gains for each of M, N, and S, respectively. Let $\alpha_N(T) = I_N/((1-p_N)P_\rho(T))$ and $\alpha_S(T) = I_S/(p_N P_\rho(T))$ be the

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



expected per patient switching costs given adoption of technology N or S, respectively. To avoid undefined expressions, we set

$$\alpha_S(T) = \begin{cases} 0, & I_S = 0 \\ \infty, & p_N = 0 \text{ or } P_\rho(T) = 0 \end{cases} \quad \text{and} \quad \alpha_N(T) = \begin{cases} 0, & I_N = 0 \\ \infty, & P_\rho(T) = 0. \end{cases}$$

The optimal adoption decision divides the open interval for posterior beliefs into three regions, delineated by $\alpha_N(T)$ and $\alpha_S(T)$: if $Z_{Tr} > \alpha_N(T)$, it is optimal to adopt N; if $Z_{Tr} < -\alpha_S(T)$, it is optimal to adopt S, otherwise it is optimal to continue with the current mix. We refer to $\alpha_N(T)$ and $\alpha_S(T)$ as *indifference points*, because one is indifferent between N and M when $Z_{Tr} = \alpha_N(T)$, and one is indifferent between S and M when $Z_{Tr} = -\alpha_S(T)$. 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 $-p_N P_\rho(T)$ when S is adopted and $(1 - p_N) P_\rho(T)$ when N is adopted.

3.2. Simplified objective function

Given the optimal adoption decision, we now simplify the expected net gain to exclude the explicit dependence on \mathcal{D} . By conditioning on Z_{Tr} , setting $x^+ = \max\{0, x\}$, and using the tower property of expectations, (4b) simplifies to the following expression:

$$V(T, r) = -(c_{\text{cap}}(r) + c\tilde{T}_\rho(T)r) + \delta_{\text{on}}(\tilde{T}_\rho(T)r/2)(1 - 2p_N)\mathbb{E}[Z_{Tr}] + e^{-\rho(T+\Delta)}\mathbb{E}[(P_\rho(T)(1 - p_N)Z_{Tr} - I_N)^+ + (-P_\rho(T)p_N Z_{Tr} - I_S)^+]. \quad (7)$$

Define the normal linear loss function $\Psi(z) \equiv \mathbb{E}[(Z - z)^+] = \phi(z) - z(1 - \Phi(z))$, where Z is a standard normal random variable with cumulative distribution function Φ and probability density function ϕ . When $Tr > 0$, (7), and therefore (4b), can be computed as follows:

$$V(T, r) = -(c_{\text{cap}}(r) + c\tilde{T}_\rho(T)r) + \delta_{\text{on}}(\tilde{T}_\rho(T)r/2)(1 - 2p_N)\mu_0 + e^{-\rho(T+\Delta)}P_\rho(T)\sigma_Z \left[(1 - p_N)\Psi\left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z}\right) + p_N\Psi\left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z}\right) \right]. \quad (8)$$

The probabilities of adopting N, S, and M are, respectively, $1 - \Phi((\alpha_N(T) - \mu_0)/\sigma_Z)$, $1 - \Phi((\alpha_S(T) + \mu_0)/\sigma_Z)$, and $\Phi((\alpha_N(T) - \mu_0)/\sigma_Z) + \Phi((\alpha_S(T) + \mu_0)/\sigma_Z) - 1$. When $Tr = 0$,

$$V(T, r) = \max\{0, (1 - p_N)P_\rho(0)\mu_0 - I_N, -p_N P_\rho(0)\mu_0 - I_S\}. \quad (9)$$

These simplifications streamline the analysis below to solve (5).

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, i.e., 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, i.e., the post-adoption 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 models (fixed patient pool and fixed horizon) that were introduced in section 2.4.1 both satisfy these conditions, as do the example cost of capacity models in section 2.3.5. The proof of the following proposition is presented in Appendix C.1.

PROPOSITION 1. *If $c_{\text{cap}}(r)$ 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 (5) exists.*

A closed-form solution to (5) 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 optimization algorithm at several random points.⁴ The next section shows that the optimization can be made to be univariate in some contexts.

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

Table 1 presents a taxonomy of the fixed sample size clinical trials which may be solved using our model. We label the categories as Cases I to IV, according to whether the setup cost, c_{cap} , is constant and whether rewards are undiscounted with a fixed patient pool. The cases are described in sections 3.4.1 to 3.4.4 where we show that, for Cases I to III, the problem reduces to the optimal choice of a single variable, namely T , r or the product of the two. For Case IV – which applies when setup costs are not constant and either discounting is used or the patient pool is not fixed (or both) – both T and r must be chosen optimally.

Most traditional clinical trials do not model time discounting (suggesting the taxonomy’s left column), but health technology assessments do (NICE 2013, Sharma et al. 2021). Because value-based trials model technology adoption, the right column is typical. Traditional patent protection is also in the right column. Although many trials have fixed initial budgets, the inclusion of additional sites to improve recruitment entails greater setup costs (US HHS 2014). This justifies that most trials are in the second row. Thus, Case IV is our main focus. We discuss the importance of appropriately modeling the case of the trial in section 4.3.3.

⁴In numerical experiments, we have not found more than two local optima for the product Tr .

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. Optimize sample size (section 3.4.1, Proposition 2)	II. Fix $r^* = r_{\max}$, optimize T (section 3.4.2, Proposition 3)
Otherwise	III. Fix $T^* = T_{\max}$, optimize r (section 3.4.3, Proposition 4)	IV. Optimize both r and T (section 3.4.4)

3.4.1. Case I: constant setup costs, undiscounted rewards and fixed patient pool.

Choose the decision variables T and r so their product optimizes the sample size, Q^* . This gives some flexibility in selecting T and r , as long as Q is optimized.

PROPOSITION 2. *If $c_{cap}(r) = c_{cap}$, $\rho = 0$, $P(T) = P$, then all members of the set $\mathcal{S} = \{(T, r) \in [0, T_{\max}] \times [0, r_{\max}] : Tr/2 = Q^*\}$, for some optimal Q^* number of patient pairs, solve (5).*

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

3.4.2. Case II: constant setup 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 allowed, by choosing $r^* = r_{\max}$ and to optimize over T .

PROPOSITION 3. *If $c_{cap}(r) = c_{cap}$, there is an optimal solution (T^*, r^*) to (5) with $r^* = r_{\max}$:*

$$V^* = \max_{T \in [0, T_{\max}]} V(T, r_{\max}). \quad (10)$$

3.4.3. Case III: non-constant setup 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 optimize the recruitment rate.

PROPOSITION 4. *If $P(T) = P$, $\rho = 0$, there is an optimal solution (T^*, r^*) to (5) with $T^* = T_{\max}$.*

$$V^* = \max_{r \in [0, r_{\max}]} V(T_{\max}, r). \quad (11)$$

3.4.4. Case IV: nonconstant setup costs and discounted rewards and/or fixed horizon. It is necessary to optimize both T and r to solve (5).

Table 2 Comparative statics results for the value function and decision variables in selected trials.

Parameter (b)	1	2	3	4
	Case I-IV trials (Proven) dV^*/db	Case IV trials (Numerical, % of instances) $dQ^*/db > 0$	$dT^*/db > 0$	$dr^*/db > 0$
P (fixed pool)	≥ 0	98.7 ¹	1.3	99.3
H (fixed horizon)	≥ 0	100.0 ¹	100.0	100.0
ρ	≤ 0	0.0	2.0	54.9
Δ	≤ 0	0.0	49.8	0.0
n_0	≤ 0	96.2 ²	99.9	11.7
σ_X	≥ 0	97.6 ³	21.8	100.0
c	≤ 0	0.0	0.0	0.0
c_r	≤ 0	0.2	99.8	0.0

Notes: Column 1 presents proven results for the derivative of the value function with respect to the parameters listed. Columns 2 to 4 present the percentage of instances for which a positive derivative is obtained in our numerical analysis of Case IV trials. ¹Positive when $\alpha_N(T) = \alpha_S(T) = 0$. ²In Case I, when $n_0 > Q^*/2$, the sign can only be negative. ³Provably positive for Case I and III trials.

3.5. Comparative statics

We use comparative statics to investigate the impact of marginal changes in selected parameters, b , on the maximized expected net gain, V^* , and the optimal values of the decision variables Q^* , T^* , and r^* . Comparative statics for V^* are obtained by assuming that the optimal values of the decision variables lie in the interior of the domain and satisfy the relevant first order necessary conditions. By the envelope theorem:

$$\frac{dV^*}{db} = \frac{\partial V(T^*, r^*)}{\partial b}.$$

Comparative statics for Q^* , T^* and r^* are obtained using the implicit function theorem and the relevant first order conditions evaluated at optimal values. For example, for a Case IV trial:

$$\frac{dT^*}{db} = \frac{\begin{vmatrix} -\partial^2 V / \partial T \partial b & \partial^2 V / \partial T \partial r \\ -\partial^2 V / \partial r \partial b & \partial^2 V / \partial r^2 \end{vmatrix}}{|\mathbf{A}|}, \quad (12)$$

where $|\mathbf{A}|$ is the determinant of the Hessian of (5). Equation (12) simplifies to $da^*/db = -(\partial^2 V(a^*) / \partial a \partial b) / (\partial^2 V(a^*) / \partial a^2)$ for Cases I–III, where $a^* = Q^*$ for Case I trials, T^* for Case II trials, and r^* for Case III trials.

Sections 3.5.1 and 3.5.2 present results for the parameters which we believe are the most relevant and interesting. Results for additional parameters are discussed in section 3.5.3.

3.5.1. Results for the maximized expected net gain. Results for dV^*/db are shown for selected parameters of interest in column 1 of Table 2. They are derived analytically in Appendix C.3 and are the same for Cases I–IV. The maximized expected net gain is increasing in P and H because a larger population translates into a larger value of the trial. It is decreasing in ρ and Δ because a larger discount rate or delay represents a smaller effective population. It is

decreasing in n_0 because more confidence in prior beliefs (a larger n_0 resulting from a smaller value of σ_0^2 relative to σ_X^2) reduces the value of information. For a fixed n_0 , an increase in σ_X increases σ_0^2 , resulting in an increase in the value of information. Increasing the enrollment (c) or setup (c_r) costs of the trial naturally decreases the maximized expected net gain.

3.5.2. Results for the decision variables. It is not possible to unambiguously sign analytical expressions for the comparative statics of the decision variables (for reasons outlined in Appendix C.3). To draw insights, we numerically solved a large number of problem instances, using as a baseline order of magnitude the parameter values from the application of section 4. We then varied these by one order of magnitude above and below, computing numerical derivatives of interest for each parameter combination using expressions such as (12). We used $c_{\text{cap}}(r) = c_r r^\theta$ as the setup cost function, with $\theta = 1, 2$, or 3 for the sensitivity analysis, so as to model nonlinear, convex, increasing costs. Columns 2–4 of Table 2 show the results for Case IV trials, presented as the percentage of problem instances where the derivatives were positive and excluding instances for which any optimal decision variable was at the boundary of its domain (these comprised 59% of the 137,781 instances that we evaluated). Appendix D provides further details. We discuss the results for Case IV trials in this section and present results for Cases I–III in Appendix D.

The analysis raises some interesting insights. For example, one might have expected dV^*/db and dQ^*/db to have the same sign: intuitively, a larger V^* justifies a larger expenditure on enrolled patients to make a better adoption decision. Although this intuition is confirmed by a large percentage of problem instances, it does not hold for all (see column 2 of Table 2). Further, results for T^* and r^* are sometimes intuitive – such as those for c and c_r – but not always – such as the derivative of T^* with respect to P and H .

For example, results for the comparative statics of the sample size with respect to the post-adoption population (dQ^*/dP (fixed patient pool) and dQ^*/dH (fixed horizon)) may be positive or negative owing to two opposing forces. Increasing the size of the population means that observations in the trial have more value (Q^* increases). However, it also decreases the expected per patient switching costs $\alpha_N(T)$ and $\alpha_S(T)$, making the adoption decision less costly (Q^* decreases). If $\alpha_N(T) = \alpha_S(T) = 0$, the latter force disappears and Q^* is increasing in P and H . Our results show that the second force is weaker, suggesting that large switching costs are required to reverse the direction of change.⁵ For fixed horizon trials, both T^* and r^* are increasing with H . In contrast, for fixed patient pool trials, T^* is generally decreasing in H , but r^* is generally increasing.⁶

⁵ If switching costs are sufficiently large, the trial would have a negative expected net gain and not be performed.

⁶ Appendix G.1 illustrates that, as the size of the population increases, the optimal recruitment duration approaches an asymptote from below in fixed horizon trials and from above in fixed patient pool trials.

An increase in the discount rate, ρ , has a negative effect on Q^* in all problem instances, but it can have a positive or negative effect on T^* and r^* . Increasing the discount rate decreases the present value of post-adoption patients, inducing similar opposing forces to those just discussed for the population size. But costs are also discounted, so there is an additional force acting in the positive direction. We found an increase in r^* in 99.9% of fixed patient pool instances but only in 9.7% of fixed horizon instances, leading to an overall figure of 54.9%. Q^* is decreasing in the delay, Δ , for all instances because Δ reduces the post-adoption population. T^* is increasing for 99.2% of fixed patient pool instances and none of the fixed horizon instances.

The optimal values of all three decision variables are decreasing in the per-patient cost, c . An increase in the marginal cost of the recruitment rate, by increasing c_r , decreases Q^* owing to a decrease in the recruitment rate. 99.8% of instances show an increase in recruitment duration.

In most instances, dQ^*/dn_0 is positive. This may be counter-intuitive, because it implies that a less informative prior (a smaller n_0) leads to the collection of less information (a lower sample size). A larger n_0 implies a narrower concentration of the prior distribution. Thus, in order to change the adoption decision from that recommended by the prior mean, a larger sample size may be needed. However, a sufficiently large n_0 can lead to a smaller sample size, as we observe in many Case I and III instances in Appendix D, because the narrower concentration of the prior distribution requires fewer observations to provide enough evidence to make an adoption decision. Appendix C.3.3 shows that the optimal sample size Q^* decreases in n_0 if $n_0 > Q^*/2$ in Case I trials.

The sampling standard deviation σ_X has a positive effect on sample size in the majority of Case IV instances. This follows the intuition that a larger sample size is required when observations are noisier. In almost all instance, σ_X has a positive effect on the recruitment rate.

There are two results where the simulated statics are not consistent (near 100%, near 0%, or near 50% owing to a split for fixed horizon and fixed patient pool). In a majority of instances, σ_X has a negative effect on the recruitment duration, except for 21.8% of instances, of which the vast majority have a large μ_0 in absolute value. And dr^*/dn_0 is usually negative, except for 11.7% of instances with the smallest n_0 and largest μ_0 tested. Because a trial is often run because there is uncertainty as to which alternative is best, it is plausible that μ_0 will be close to zero in practice, in which case we expect $dT/d\sigma_X$ and dr^*/dn_0 to be negative.

3.5.3. Results for other parameters. The signs of derivatives with respect to μ_0 and p_N are indefinite due to symmetry. For example, an increase in p_N increases the size of the population that benefits from adopting S, but it decreases the size of the population that benefits from adopting N. There are two special parameter settings in which the optimal trial design is the same for any p_N . The first is when $I_N = I_S = \mu_0 = 0$; the second is when $I_N = I_S = 0$ with undiscounted rewards,

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 an application in Appendix G.1.3.

The sensitivity of decision variables to changes in other parameters can be assessed using the comparative statics results in this section. For example, inspection of (8) shows that a shift from offline to online learning is equivalent to a shift of $(1 - 2p_N)\mu_0/2$ in c . Increasing the value of one unit of effectiveness, λ , can be assessed using the results for σ_X .⁷

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 shown with the ProFHER application in Appendix G.

3.6.1. Undiscounted rewards. We study the limiting behavior of V^* and Q^* as $P \rightarrow \infty$ for a fixed patient pool and as $H \rightarrow \infty$ for a fixed horizon. We study these limits while ignoring the constraints $T \leq T_{\max}$ and $r \leq r_{\max}$ unless they are necessary for the existence of a solution to (5). We break the assumptions of the model on the bounds of the decision variables to allow for the derivation of the theoretical properties that provide insights on the rate of growth of the expected net gain and decision variables: for Cases I and II, we fix $r = r_{\max} < \infty$ and show that $T^* \rightarrow \infty$; for Case III, we fix $T = T_{\max} < \infty$ and show that $r^* \rightarrow \infty$; for Case IV, we show that both T^* and r^* grow unboundedly. With finite bounds, we can show that $T^* \rightarrow T_{\max}$ and $r^* \rightarrow r_{\max}$. The use of these results as approximations in practice requires that the post-trial population, as well as T_{\max} and r_{\max} , are large enough. Appendix G.2 demonstrates with an application that the approximations can be accurate for reasonable values of these parameters.

Prop. 5 presents this asymptotic behavior for the case of undiscounted rewards (i.e., Cases II and IV are necessarily fixed horizon). It assumes that $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 > 0$; otherwise, the limits often do not exist. It further assumes that $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r$ 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 $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r$. Then, for Cases I ($c_r = 0$) and III ($c_r > 0$)*

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

⁷ Assume that $E_N - E_S \sim \mathcal{N}(\mu_E, \sigma_E^2)$ and $C_N - C_S \sim \mathcal{N}(\mu_C, \sigma_C^2)$ with correlation ρ_{EC} for $E_N - E_S$ and $C_N - C_S$. Then, $\mu_0 = \lambda\mu_E - \mu_C$ and $\sigma_X^2/n_0 = \lambda^2\sigma_E^2 + \sigma_C^2 - 2\lambda\sigma_E\sigma_C\rho_{EC}$. Assuming n_0 is fixed, a δ increase in λ shifts μ_0 by $\delta\mu_E$ units and changes σ_X^2 by $n_0\delta\sigma_E(2\lambda\sigma_E + \delta\sigma_E - 2\sigma_C\rho_{EC})$. If the prior belief is that neither technology is more clinically effective (i.e. $\mu_E = 0$, a typical default value), then there is no shift in μ_0 .

For Cases II ($c_r = 0$) and IV ($c_r > 0$)

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

For Case II, where $c_r = 0$,

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

For Case IV, where $c_r > 0$,

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

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.

For a finite value of $P(T)$, asymptotic approximations to the optimal value of the decision variable and maximized 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 $c\hat{T}\hat{r}$, and $\hat{T}\hat{r}$ 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 of H (fixed horizon). Second, both the fixed patient pool and fixed horizon scenarios 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 scenario with positive discount rate, $\rho > 0$. This is only applicable to Cases II and IV (Cases I and III require undiscounted rewards). We denote with subscripts the post-adoption scenario. For instance, $V_H(T, r)$ refers to the expected net gain with fixed horizon, while $V_P(T, r)$ refers to the expected net gain with fixed patient pool. Unlike the asymptotic results for undiscounted rewards, V^* , T^* , r^* , and Q^* do not diverge in the limit of large $P(T)$. Therefore, the results here take a different form. See Appendix C.5 for proofs.

It is easy to check that $\lim_{P \rightarrow \infty} P_\rho = \zeta/\rho$ with fixed patient pool, and $\lim_{H \rightarrow \infty} P_\rho(T) = \zeta/\rho$ with fixed horizon. Define similarly $\alpha'_N = \lim_{P \rightarrow \infty} \alpha_N = \lim_{H \rightarrow \infty} \alpha_N(T) = I_N\rho/((1 - p_N)\zeta)$, and $\alpha'_S = \lim_{P \rightarrow \infty} \alpha_S = \lim_{H \rightarrow \infty} \alpha_S(T) = I_S\rho/(p_N\zeta)$. 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 , $\alpha_N(T)$ with α'_N , and $\alpha_S(T)$ with α'_S in (8) and (9):

$$V_\infty(T, r) = \begin{cases} -(c_{\text{cap}}(r) + cr\tilde{T}_\rho(T)) + \delta_{\text{on}}(r\tilde{T}_\rho(T)/2)(1 - 2p_N)\mu_0 \\ \quad + e^{-\rho(T+\Delta)}\frac{\zeta}{\rho}\sigma_Z \left[(1 - p_N)\Psi\left(\frac{\alpha'_N - \mu_0}{\sigma_Z}\right) + p_N\Psi\left(\frac{\alpha'_S + \mu_0}{\sigma_Z}\right) \right], & \text{if } Tr > 0 \\ \max\{0, (1 - p_N)\zeta\mu_0/\rho - I_N, -p_N\zeta\mu_0/\rho - I_S\}, & \text{if } Tr = 0. \end{cases} \quad (13)$$

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

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

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

In summary, for finite but large $P(T)$, we can approximate (5) with

$$\begin{aligned} & \max_{T, r} V_\infty(T, r) \\ & \text{s.t. } T \in [0, T_{\max}], \quad r \in [0, r_{\max}]. \end{aligned} \tag{14}$$

The function $V_\infty(T, r)$ matches the discounted fixed patient pool expected net gain with $P_\rho = \zeta/\rho$. Thus, all previous results related to discounted fixed patient pool are also valid when solving (14). Both the fixed patient pool and fixed horizon models converge to the same function and maximizers.

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 HT decision. The original analysis used discounted rewards (Corbacho et al. 2016) and we assume nonconstant trial setup costs, so we model this as a Case IV trial, for both fixed horizon and fixed patient pool. Hence, both T and r are optimized.

The ProFHER trial was a multi-center, randomized 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 the course of approximately two and a half years, 250 patients across 35 NHS hospitals were randomized, 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, together with treatment and rehabilitation costs) took place after six, twelve, and 24 months. Results showed no statistically significant difference between surgery and sling in terms of effectiveness, but that surgery cost, on average, approximately £1,800 more than sling. Corbacho et al. (2016) concluded that surgical treatment was ‘not cost effective for the majority of displaced fractures of the proximal humerus involving the surgical neck in the UK’s NHS.’

We assess the performance of the trial design using the optimized expected net gain V^* , recruitment rate r^* , recruitment duration T^* and sample size $Q^* = T^* r^* / 2$. We also use two probabilistic metrics. The first, which is based on Bayesian principles, is the *conditional probability of correct selection* (CPCS). The CPCS is the probability of adopting the correct HT given a specific value

of W . Let \mathcal{D}^{oracle} be the *oracle adoption decision*, which selects the HT with the highest benefits for post-trial patients, knowing the true value of W *a priori*. Then $\text{CPCS}(w) = \Pr(\mathcal{D} = \mathcal{D}^{oracle} \mid W = w)$. 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 showing the probability that a 95% confidence interval for the unknown mean does not contain zero, as a function of w .⁸ Appendix E provides further details.

4.1. Parameter values

Our base case analysis uses the parameter values reported in Forster et al. (2021, Table SM.2), obtained by referencing the ProFHER trial’s main publications and consulting its research team, with some small variations which are described below. Define surgery as the new technology, N, and sling as the standard, S. We consider offline learning ($\delta_{on} = 0$) and set $p_N = 0.39$, $\mu_0 = \mathcal{L}0$, $n_0 = 2$, and $\sigma_X = \mathcal{L}4,400$. We consider both fixed horizon and fixed patient pool trials, with parameters $H = 15$ years, $\Delta = 1$ year, $\zeta = 7,000/\text{year}$, and $P = \zeta H = 105,000$. We set $I_N = \mathcal{L}0$, $I_S = \mathcal{L}0$, $c = \mathcal{L}2,040$ and let the continuous discount rate be $\rho = 3.44\%/\text{year}$, equivalent to the 3.5% annual discount rate recommended by NICE (2013) and used in Corbacho et al. (2016). We estimate a nonlinear setup cost function $c_{cap}(r)$, described in section 4.2.⁹ For clarity of exposition, we use subscripts to denote the post-adoption scenario. For instance, V_H^* and V_P^* refer to the maximum expected net gain for fixed horizon and fixed patient pool trials, respectively.


4.2. Estimating the setup cost function

Cost data by line-item exist for the fixed costs, site-specific costs, and patient variable costs of large drug trials (US HHS 2014). ‘Cost-per-site’ data, obtained by averaging trial costs over all sites, does not model potentially nonlinear setup costs. Additional factors suggest that $c_{cap}(r)$ has terms that increase at least linearly in r . For example, if trial sites with greater potential for the recruitment of patients are prioritized by trial managers, then marginal setup costs, $dc_{cap}(r)/dr$, increase in the recruitment rate if more sites with lower potential are included in the trial.

We used data for the number of patients recruited at each site in the ProFHER trial (Handoll et al. 2015, Fig. 5, page 37) to estimate the setup cost function, $c_{cap}(r)$. We ordered the sites

⁸ CPCS and power quantify the probability of correctly adopting a HT, but they differ in two respects. Firstly, CPCS uses prior information, whereas power considers only data 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.

⁹ Our choices differ from the parameter values in Forster et al. (2021) as follows: $H = 15$ differs from the choice of $H = 6$ (reflecting the fact that gains from an adoption decision are likely to be acquired beyond a period of next potential review in 6 years); we discount costs and benefits, whereas Forster et al. (2021) did not; Forster et al. (2021) did not model site-specific setup costs.

Figure 3 Cost of recruitment rate assuming a constant cost of opening a site.

according to the total number recruited and found the setup cost function which minimized the mean squared error under the following simplifying assumptions: 1. half of the setup costs are fixed and half are marginal; 2. the cost of opening sites is constant; 3. techniques mentioned in section 2.3.5 can be used to achieve the desired recruitment rate (here, the rate at which patient data is eventually observed, net of patients who decline participation or do not attend follow-up visits). The resulting function was estimated to be $c_{\text{cap}}(r) = 480,000 + 766r^{3.06}$. See Figure 3.

4.3. Numerical experiments

Inspection of (8) shows that, when $\mu_0 = 0$, online and offline learning have the same optimal one-shot design and expected net gain. Since we assume that $\mu_0 = 0$ for the ProFHER application, results presented here are equally applicable to the cases of online and offline learning.¹⁰

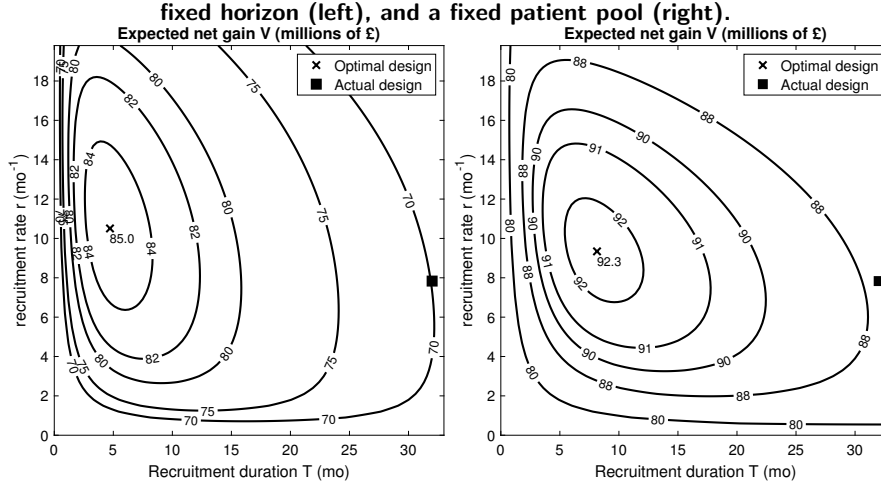
4.3.1. Optimal trial design and expected net gain for the base case parameters.

Figure 4 shows contour plots of the expected net gain as a function of the recruitment duration, T , and the recruitment rate, r (left subfigure: fixed horizon; right subfigure: fixed patient pool). The optimal design (T^*, r^*) is marked with a cross and the design of the actual trial, $\check{T} = 32$ months and $\check{r} = 7.8/\text{month}$, is marked with a square. With a fixed horizon, the optimal design is $T_{\text{H}}^* = 4.7$ months and $r_{\text{H}}^* = 10.5/\text{month}$ ($Q_{\text{H}}^* = 25$ patient pairs) and obtains $V_{\text{H}}^* = \pounds 85.0\text{Mill}$. With a fixed patient pool, the optimal design is $T_{\text{P}}^* = 8.1$ months and $r_{\text{P}}^* = 9.3/\text{month}$ ($Q_{\text{P}}^* = 38$ patient pairs) and obtains $V_{\text{P}}^* = \pounds 92.3\text{Mill}$.

We point out three important implications of these results. Firstly, both the maximized expected net gain of the trial and the optimal choice of the decision variables are sensitive to the post-adoption scenario: the optimal expected net gain and sample size are larger under fixed patient

¹⁰ Appendix G.1 presents additional sensitivity analysis. Appendix G.1.4 shows little apparent variation in results between online and offline learning when $\mu_0 \neq 0$.

Figure 4 Contour plots of the expected net gain as a function of T and r for the ProFHER application with a fixed horizon (left), and a fixed patient pool (right).




pool compared to fixed horizon by £7.3Mill (8.6%) and 13 (52%), respectively.¹¹ Secondly, the expected net gain function is more sensitive to the design under the fixed horizon than the fixed patient pool: extending the length of the trial reduces the size of the post-adoption population in the fixed horizon, but not in the fixed patient pool. This is particularly evident by observing the expected net gain of the actual trial's design, which is £4.7Mill (5.1%) below the optimized expected net gain under fixed patient pool and £15.0Mill (17.7%) under fixed horizon. Thirdly, the expected net gain in both graphs, for both the optimal and actual trials in this example, is more sensitive to a change in trial duration, as compared to a similar relative change in recruitment rate. This is reassuring in case the anticipated recruitment rate differs from the realized recruitment rate, e.g., when the drop-out rate is misestimated.

4.3.2. Selecting the recruitment rate. Figure 5 plots the marginal value of the trial's recruitment rate, $mv_r = d(V(T^*, r) + c_{cap}(r))/dr$, the derivative of the expected net gain excluding setup costs with respect to r , at the optimal choice of trial duration, T^* , together with the marginal cost, $mc_r = dc_{cap}(r)/dr = 2,344r^{2.06}$, as a function of r . The intersection of marginal value and marginal cost functions corresponds to the optimal recruitment rate.

Figure 5 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 5 shows that the marginal value is £440,000 per additional recruit per month for fixed horizon and £292,000 for fixed patient pool. A trial manager charged with maximizing the value of the trial can think of these marginal values as representing the maximum willingness to pay (WTP) per unit increase of recruitment per month. It is optimal to open an additional site if the cost of doing so is lower than the WTP.

¹¹ When $P = \zeta H$, the expected net gain is always larger under fixed patient pool because it treats a larger population post-adoption.

Figure 5 Marginal value and cost of the recruitment rate.

4.3.3. Benefits of optimizing both the recruitment rate and duration according to the regulatory context. In Section 3.4 we stated that Case IV trials – which optimize over both the recruitment rate and duration – are commonly encountered in health technology assessment and in drug trials. This section uses our application to explore the potential benefit of running a trial as a Case IV trial by optimizing both the recruitment rate and the recruitment duration, over and above optimizing only one of these two decision variables. It then explores the potential impact on expected net gain if a modeler assumes a Case I context, as is implicit in traditional clinical trial design settings, when rewards are consistent with a Case IV context.

Figure 6 (left panel) shows the percentage loss in expected net gain when, instead of optimizing both decision variables, we fix the recruitment rate and optimize the recruitment duration (or, equivalently, the sample size). A deviation of five recruitment rate units loses 2% of overall value, and optimizing T only with a low rate of $r = 2$ patients/month leads to 8-15% loss. These are noteworthy amounts if one observes that the Case IV design has an expected net gain of around £100Mill.

Figure 6 (right panel) shows the percentage loss if the modeler assumes a Case I trial instead of a Case IV trial, and optimizes the sample size rather than both the recruitment rate and recruitment duration. Loss is plotted as a function of the recruitment rate, so the recruitment duration is adapted to keep the sample size constant at the optimal value of the Case I trial. The loss is more substantial when we make decisions that assume no discounting and constant setup costs (Case I), when the chosen objective function has discounting and increasing setup costs (Case IV). Here, losses can range from 5% to 70%. Moreover, these assumptions affect the trial design by prescribing sample sizes which are over four times larger.

4.3.4. CPCS and power plots for the optimal design. Figure 7 presents the CPCS and power curves for the optimal designs with the base case parameters and the actual trial. Such power curves may be useful, for example, in communicating with the [US FDA \(2019\)](#). The vertical lines

Figure 6 The percentage loss in expected net gain. The trial is a Case IV trial, but: (left panel) only the recruitment duration is optimized; (right panel) the modeler assumes a Case I trial and optimizes sample size.

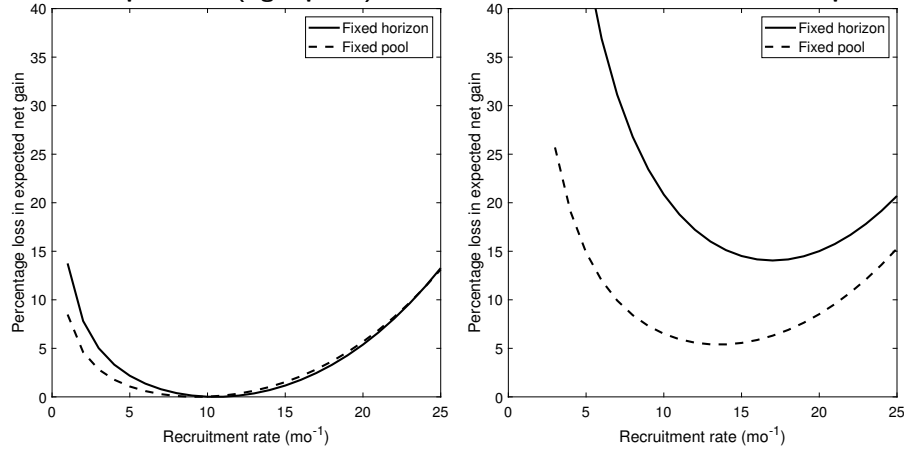
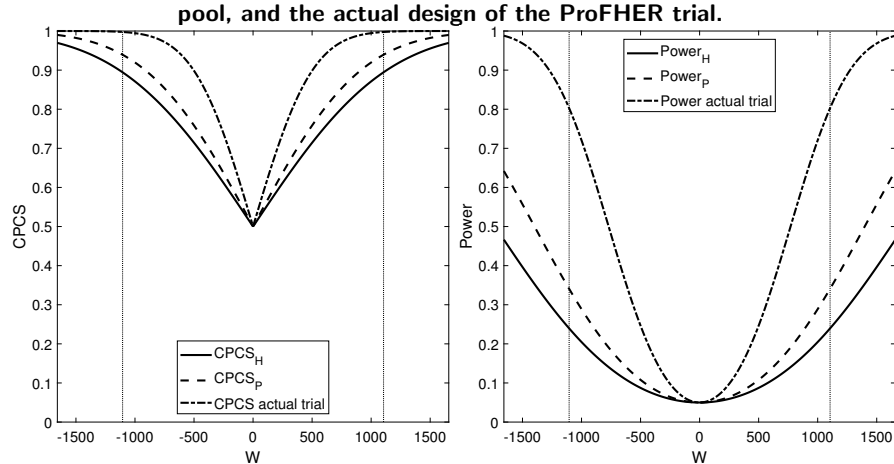


Figure 7 CPCS (left) and power (right) curves for the optimal design under fixed horizon and fixed patient pool, and the actual design of the ProFHER trial.



represent the smallest relevant difference for the frequentist sample size calculations.¹² Because $I_N = I_S = \mathcal{L}0$ (so that $\alpha_N(T) = \alpha_S(T) = 0$), it is almost never optimal to choose $\mathcal{D} = M$. 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 $\alpha_N(T) \neq 0$ or $\alpha_S(T) \neq 0$. Both CPCS and power are higher under the fixed patient pool compared to the fixed horizon due to the larger sample size in our application.

Figure 7 shows that the CPCS of the actual trial at the smallest relevant difference is 99.8%. The optimal designs under a fixed horizon and a fixed patient pool obtain 89.4% and 93.9%, respectively. While the optimal designs do not have a CPCS as high as the actual trial, they still make correct decisions with high probability. The power curve shows a different story. At the smallest relevant

¹² 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 smallest relevant difference in terms of INMB that equates the sample size calculations of the actual trial is $\mathcal{L}1105$.

difference, the actual trial obtains 80% power by the design criteria. The optimal designs obtain 23.9% for fixed horizon, and 34.0% for fixed patient pool.

In this application, the CPCS and power of the optimal value-based design is lower than the frequentist design. This is not always true and is dependent on the parameters of the model. Based on the results in sections 3.5–3.6 and in Appendix G below, the sample size, and hence the CPCS and power, of the optimal value-based design can exceed the frequentist design by, for instance, having a smaller discount rate, a larger post-adoption population (e.g., by extending the duration of market exclusivity), or a smaller per-patient cost (using the comparative statics of Table 2).

More generally, such power curves can also indicate whether a larger sample size might be helpful in achieving a minimum power if that is really required (in which case the trial would not create as much value, given the parameters specified by the modeler), or whether the power is more than satisfactorily addressed, given the value-based focus of the adoption decision. Prop. 5 and Prop. 6 imply that the sample size grows without bound with undiscounted rewards, and therefore, statistical power increases to 1 as $P(T)$ grows without bound and as ρ goes to 0.

4.4. Further observations from the ProFHER trial

Appendix G.1 complements the comparative statics analysis of section 3.5 with sensitivity analysis for the ProFHER trial. Appendix G.2 shows that the asymptotic approximations can be accurate for Case I–III trials, even when $P(T)$ is not too large. We find that the asymptotic approximations for fixed horizon can be slower to converge than for fixed patient pool. The reduction in post-adoption patients for longer trials with a fixed horizon may be driving this. For large post-adoption populations, the asymptotic approximations can be quite good. However, for smaller trials, such as ProFHER, the optimization in (5) without asymptotic approximation is more appropriate.

5. Extension: Response-adaptive trial duration

The one-shot design above does not formally permit interim analysis of the data as it accumulates. Section 1 notes strong interest in adaptive trials and other sequential sampling applications. Here, response adaptive options include (a) adapting the trial duration T as data accumulates for a given recruitment rate r , and optimizing over fixed r ; (b) adapting the recruitment rate r as data accumulates for a given T (or sample size $Q = Tr/2$), optimizing over fixed T (or Q), or (c) adapting r and T (or r and Q) in a group sequential trial as data accumulates.

Here, we explore option (a), which fully adapts the duration as each outcome is observed and allows recruitment rate to be optimized, with a limit of not adapting the rate. Options (b) and (c) adapt the recruitment rate itself, which adds complexity. Kouvelis et al. (2017) adapt recruitment rate through time and Rojas-Cordova and Bish (2018) adapt sample size per batch, for 0-1 trials, under a somewhat different set of assumptions (e.g., for-profit setting, or with 0-1 outcomes). Combining group rate-adaptive and value-adaptive trials is left for future work.

Optimally adapting trial duration for optimal fixed recruitment rate. For four special cases of the general model in this paper, Appendix F shows that a modest adaptation of Chick et al. (2017), Alban et al. (2018) allows us to (i) identify the optimal stopping time policy π_r^* for a fully adaptive trial with fixed recruitment rate r , (ii) compute its expected net gain, $V(\pi_r^*, r)$, and (iii) optimize over r to obtain a trial design $(r^*, \pi_{r^*}^*)$ to optimize expected net gain $V(\pi_r, r)$ over the class of all fixed recruitment rate trials with fully response-adaptive sample sizes. This allows us to assess the benefits of optimizing recruitment rate for a one-stage trial, as compared to optimizing the fixed recruitment rate of a fully adaptive trial, at least for four special cases.


The four special cases each assume fixed patient pool ($P(T) = P$) and are: 1) current practice is one treatment only ($p_N = 0$); 2) rewards are undiscounted and one of the two technologies must be adopted ($\tilde{\rho} = 0$ and $\mathcal{D} \in \{N, S\}$); 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 ($\tilde{\rho} \geq 0$, $I_N = I_S = \mu_0 = 0$). These settings cover the two scenarios noted in section 3.5.3.

Application to ProFHER trial. We use the context of the ProFHER trial to illustrate the relative benefits of optimizing r for a fixed trial and for a fully sequential trial. We assume base case parameter settings with fixed patient pool, which matches special case 4) above.

Figure 8 compares the expected net gain and expected sample size of the sequential and one-shot designs as a function of the fixed recruitment rate and delay. The left panel shows that the optimal rate of recruitment for the sequential trial, $r_P^{*,seq} \approx 7$ is lower than the optimal rate, $r_P^* \approx 9.3/\text{month}$, for a one-shot trial. If the recruitment rate is set to 7/month, the expected net gain of the rate-optimal fully sequential trial is £0.4Mill ($= 92.5 - 92.1$). Increasing the recruitment rate to 9.3/month, the optimal rate for a one-shot trial increases the expected net gain by £0.2Mill ($= 92.3 - 92.1$). Comparing the sequential and one-shot designs each with their respective optimal recruitment rates, we see a difference of £0.2Mill ($= 92.5 - 92.3$). At high recruitment rates, there is no benefit to running a response-adaptive trial over an optimal one-shot trial in this case (center panel), because the number of pipeline patients is too large to obtain a benefit from adaptation.

The right panel of Figure 8 illustrates the expected net gain for optimal sequential and one-shot trials with fixed recruitment rate as a function of the delay in observing patient outcomes (the delay was 12 months in the left and center panels as in the baseline parameters of the application). The gain from sequential, above one-shot, is largest for short delays, as there is more time to adapt the length of the trial. The difference becomes zero for particularly long delays in observing outcomes because the optimal sequential and optimal one-shot have the same expected net gain when delays exceed the length of the optimal one-shot trial.

Figure 8 Comparison of the sequential and one-shot designs as a function of the fixed recruitment rate in terms of expected net gain (left) and expected sample size (center), and as a function of delay for outcomes (right).



6. Discussion

Consistent with existing literature on value-based designs as noted in section 1, our results challenge some traditional perceptions of the efficiency of clinical trials and their design. We confirm the clear trade-off between value and statistical power argued previously, and contribute to this stream of literature in several ways. Firstly, we show that it is important to model the recruitment rate and duration, not just the sample size, because most trials arguably should be considered as a Case IV trial in Table 1: trials with more patients are more expensive, and either patent expiration is a concern and/or future cash flows for benefits from technology adoption are discounted. Secondly, we consider these decisions in a wide range of regulatory contexts. Modeling the context appropriately (especially discounted or undiscounted, fixed patient pool or fixed horizon) and cost structure for obtaining a rate of data acquisition are important. Thirdly, we address mixed clinical practice (when two HTs are used to treat similar patients in ongoing medical practice), a scenario commonly encountered in clinical research. Fourthly, we provide comparative statics, asymptotic analysis, and a numerical analysis based on the ProFHER trial and HT adoption decision. We now outline our findings and recommendations for the stakeholders identified in the introduction.

Our model can support clinical trial managers tasked with maximizing value by informing the optimal recruitment rate and duration. In section 4.3.2, we illustrate that the marginal cost of increasing the recruitment rate (e.g., using techniques like those mentioned in section 4.2) should be balanced against the marginal benefits of a more rapid rate of information gain. In addition, we recommend that trial managers consider that pragmatic trials may have both HTs used in clinical practice: the fraction using each HT can significantly impact both the optimal value and sample size when the switching costs are large. Furthermore, our asymptotic analysis provides easy to follow guidance for large populations. For undiscounted rewards, we find that the optimal sample size is proportional to the square root of the post-adoption population. For discounted rewards,

we show that fixed horizon and fixed patient pool are equivalent in the limit of an infinite horizon. For trials in a clinical area with lower rates of innovation, determining the population size with an infinite horizon discounted reward model makes sense. In clinical areas with rapid innovation, the model of [Bravo et al. \(2018\)](#) may be useful to set the population size. Importantly, the optimal sample size is proportional to the square root of the incidence (for a large incidence).

Funders of trials can use our approach to assess the value of different proposed trials under the current regulatory context. Thus, funders of trials can use our model to identify trials worth running and allocate the optimal amount of resources for the chosen trials. Moreover, our investigation on the costs of increasing the recruitment rate in section 4.2 reveals the current challenges in estimating the cost of setting up a trial with a target recruitment rate. This suggests a role for funders in developing tools to match the supply of conditions tested in trials with demand from geographies having relevant disease burden (e.g., [Lock 2019](#)). Our results also underscore the importance of improving recruitment rates and sharing best practice ([Bower et al. 2014](#), [Walters et al. 2017](#), [Lipman et al. 2019](#)) as an area for dialog between funders and trial managers.

For public sector policymakers and regulators, we have presented results about the impact of the regulatory context on the value of the trial and the trial design. We find that value increases with increasing post-adoption population, decreasing discount rate, and increasing monetary valuation of effectiveness in the definition of INMB. From our application, we observe trends that are likely to occur in practical settings. Firstly, the fixed patient pool scenario obtains higher value, recruits more patients, and its value is less sensitive to the trial design. Secondly and most importantly, [Figure 6](#) illustrates a large potential loss in expected net gain if a trial is designed without accounting for variable trial costs and discounting, if discounting is used in the ensuing health technology assessment decisions (i.e., if assuming Case I when Case IV is relevant).

While this paper focused on clinical trials and HT adoption, the results associated with the different regulatory contexts may find application in non-health sectors, particularly for A/B tests where observations 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 have lower exploitation value). Sequential tests are enabled with techniques in section 5, even when both A and B are in current usage.

7. Conclusion

Several initiatives across the United States and Europe call for more innovative and efficient trial designs ([NIHR 2018, 2020](#), [US FDA 2019](#), [European Union 2021](#)). In part, these initiatives are designed to increase technical efficiencies in the way trials operate under the existing paradigm, focusing on the clinical effectiveness of HTs. However, echoing authors mentioned in section 1, the current drive for more value-based health care also raises questions about whether there exists a role

for more value-based clinical trials in the HT assessments of the future. We have presented a model that establishes the logical implications of implementing value-based trials, including the influence of the regulatory context. In this model, trials are mechanisms to potentially improve health value, rather than activities whose costs are to be minimized subject to statistical constraints.

We briefly consider future research. Firstly, the value-based model may lead to a different sample size than that of a classical design. A change in sample size, higher or lower, may lead to a different degree of adoption of the HT by practitioners (who differ from HT appraisers). The volume of uptake as a function of the strength of evidence and degree of health improvement is an interesting empirical question. Secondly, we focused on value in expectation in a risk-neutral system. Other options include risk-averse utility functions or a robust optimization approach. Thirdly, one might also wish to adapt the recruitment rate, going beyond finding the optimal fixed recruitment rate for the response adaptive duration of section 5. Fourthly, we assumed a homogeneous patient population. Identifying subpopulations for which treatments are particularly effective or ineffective would be a useful addition. Fifthly, an assessment with respect to randomness in delays of outcomes and of potentially missing data is an area for further work. Sixthly, although most jurisdictions discount QALYs and cash flows the same (as do 27 of 30 countries studied by Sharma et al. 2021), allowing different discount rates for health and money would be an applicable generalization. Finally, it would be interesting to pursue further analytical investigation of the benefits of co-optimizing recruitment rate and duration, to complement the evidence that we present in our application in section 4.

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). Forster acknowledges funding from the Research Infrastructure Support Fund of the Department of Economics and Related Studies, University of York. We thank seminar attendees at the University of Sheffield, MIT, INSEAD, UCL, LBS, and the University of Lausanne for feedback. We thank the referees for their constructive feedback.

References

- Achilleas T, Farrokhyar F, McKnight L, Bhandari M (2010) How to optimize patient recruitment. *Canadian Jour Surgery* 53(3):205–210, URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC2878987/>.
- 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., *Proc. 2018 Winter Simulation Conference*, 2459–2470 (Piscataway, NJ: IEEE, Inc.).

-
- Anderer A, Bastani H, Silberholz J (2019) Adaptive clinical trial designs with surrogates: When should we bother? Available at SSRN: <https://ssrn.com/abstract=3397464>.
- Angus D, Musthafa A, Clermont G, et al. (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 cesarean section rates: global, regional and national estimates: 1990–2014. *PLoS One* 11(2):e0148343.
- Bower P, Brueton V, Gamble C, et al. (2014) Interventions to improve recruitment and retention in clinical trials: a survey and workshop to assess current practice and future priorities. *Trials* 15(399).
- Branke J, Chick S, Schmidt C (2007) Selecting a selection procedure. *Management Science* 53(12):1916–1932.
- Bravo F, Corcoran T, Long E (2018) Flexible drug approval policies. <https://ssrn.com/abstract=3027263>.
- 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, Jaggi 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 (2021) Bayesian sequential learning for clinical trials of multiple correlated medical interventions. *Management Science* 67:accepted to appear.
- 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, 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. *JRSS Series B* 75(1):48.
- Drummond MF, Sculpher MJ, Claxton K, Stoddart GL, Torrance GW (2015) *Methods For the Economic Evaluation of Health Care Programmes* (New York: Oxford), 4th edition.
- Eckermann S, Willan AR (2007) Expected value of information and decision making in HTA. *Health Economics* 16:195–209.
- Emanuel EJ, Zhang C, Glickman A, et al. (2020) Drug Reimbursement Regulation in 6 Peer Countries. *JAMA Internal Medicine* 180(11):1510–1517, URL <https://doi.org/10.1001/jamainternmed.2020.4793>.
- European Union (2021) Clinical trial regulation EU no. 536/2014. <https://www.ema.europa.eu/en/human-regulatory/research-development/clinical-trials/clinical-trial-regulation>, European Medicines Agency, Update 21 April 2021.
- Fenwick E, Steuten L, Knies S, et al. (2020) Value of information analysis for research decisions – an introduction: Report 1 of the ISPOR value of information analysis emerging good practices task force. *Value in Health* 23(2):139–150.
- 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.
- Flight L, Arshad F, Barnsley R, Patel K, Julious S, Brennan A, Todd S (2019) A review of clinical trials with an adaptive design and health economic analysis. *Value in Health* 22:391–398.
- Forster M, Brealey S, Chick S, et al. (2021) Cost-effective clinical trial design: Application of a Bayesian sequential stopping rule to the ProFHER pragmatic trial. *Clinical Trials* 18(6):647–656, URL <http://dx.doi.org/10.1177/17407745211032909>.
- 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.
- Gittins J, Pezeshk H (2000) A behavioural Bayes method for determining the size of a clinical trial. *Drug Information Journal* 34:355–363.
- 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.
- Hampson L, Jennison C (2013) Group sequential tests for delayed responses. *Journal of the Royal Statistical Society, Series B* 75:3–54.
- 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-effectiveness 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.
- Inoue K, Chick SE (1998) Comparison of Bayesian and frequentist assessments of uncertainty for selecting the best system. Medeiros DJ, Watson EJ, Manivannan M, Carson J, eds., *Proc. 1998 Winter Simulation Conference*, 727–734 (Piscataway, NJ: IEEE, Inc.).
- 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.
- Kaplan EH, Brandeau ML, eds. (1994) *Modeling the AIDS epidemic: Planning, policy, and prediction* (Raven Press).
- 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.
- Kunst NR, Alarid-Escudero F, Paltiel AD, Wang SY (2019) A value of information analysis of research on the 21-gene assay for breast cancer management. *Value in Health* 22(10):1102–1110.
- Lachin JM (1981) Introduction to sample size determination and power analysis for clinical trials. *Control. Clin. Trials* 2:93–113.
- Lexchin J (2021) Time to Marketing of Generic Drugs After Patent Expiration in Canada. *JAMA Network Open* 4(3):e211143–e211143, URL <http://dx.doi.org/10.1001/jamanetworkopen.2021.1143>.
- Lipman P, Druzak L, Stoney C (2019) Is this study feasible? facilitating management of pragmatic trial planning milestones under a phased award funding mechanism. *Trials* 20(2):241–51.
- Lock S (2019) Matching clinical trials with unmet clinical need. *BioScience Today* 16(spring):24–25, <https://issuu.com/distinctivepublishing/docs/bst16/24>.
- Long EF, Vaidya N, Brandeau ML (2008) Controlling co-epidemics: Analysis of HIV and tuberculosis infection dynamics. *Operations Research* 56(6):1366–1381.
- Mitchell JM, Patterson JA (2020) The inclusion of economic endpoints as outcomes in clinical trials reported to clinicaltrials.gov. *Journal of Managed Care and Specialty Pharmacy* 26(4):386–393, URL <http://dx.doi.org/10.18553/jmcp.2020.26.4.386>.
- Montazerhodjat V, Chaudhuri SE, Sargent DJ, Lo AW (2017) Use of Bayesian decision analysis to minimize harm in patient-centered randomized clinical trials in oncology. *JAMA Oncology* 3(9):e170123.
- NICE (2013) Guide to the methods of technology appraisal. <https://www.nice.org.uk/process/pmg9/chapter/foreword> (Accessed 19th January, 2021).
- NICE (2014) Developing NICE guidelines: The manual. UK National Institute for Health and Care Excellence, <https://www.nice.org.uk/process/pmg20/chapter/incorporating-economic-evaluation>.

-
- NICE (2018) Guide to the processes of technology appraisal. <https://www.nice.org.uk/Media/Default/About/what-we-do/NICE-guidance/NICE-technology-appraisals/technology-appraisal-processes-guide-apr-2018.pdf>.
- NIHR (2018) NIHR Efficacy and mechanism evaluation. <https://www.nihr.ac.uk/explore-nihr/funding-programmes/efficacy-and-mechanism-evaluation.htm>.
- NIHR (2019) NIHR Clinical Research Network: Impact and value assessment. https://www.nihr.ac.uk/documents/impact-and-value-report/21427#Executive_summary.
- NIHR (2020) Annual efficient studies funding calls for CTU projects. UK National Institute for Health Research, <https://www.nihr.ac.uk/documents/ad-hoc-funding-calls-for-ctu-projects/20141>.
- Pallmann P, et al. (2018) Adaptive designs in clinical trials: why use them, and how to run and report them. *BMC Medicine* 16(29), accessed 14 July 2021, <https://doi.org/10.1186/s12916-018-1017-7>.
- Panteli D, Eckhardt H, Nolting A, et al. (2015) From market access to patient access: overview of evidence-based approaches for the reimbursement and pricing of pharmaceuticals in 36 European countries. *Health Research Policy and Systems* 13(39), URL <https://doi.org/10.1186/s12961-015-0028-5>.
- Paul SM, Mytelka DS, et 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 fractures of the proximal humerus: the PROFHER randomized clinical trial. *Journal of the American Medical Association* 313(10):1037–1047.
- Robert CP (2007) *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation* (New York: Springer), 2nd edition.
- Rojas-Cordova A, Bish EK (2018) Optimal patient enrollment in sequential adaptive clinical trials with binary response. Available at SSRN, <https://ssrn.com/abstract=3234590>.
- Rome B, Lee C, Kesselheim A (2021) Market exclusivity length for drugs with new generic or biosimilar competition, 2012–2018. *Clinical Pharmacology & Therapeutics* URL <http://dx.doi.org/https://doi.org/10.1002/cpt.1983>.
- Rothery C, Strong M, et al. (2020) Value of information analytical methods: Report 2 of the ISPOR value of information analysis emerging good practices task force. *Value in Health* 23(3):277–286.
- 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.

-
- Sharma D, Aggarwal A, Downey L, et al. (2021) National healthcare economic evaluation guidelines: A cross-country comparison. *PharmacoEconomics Open* <https://doi.org/10.1007/s41669-020-00250-7>.
- US FDA (2015) Patents and exclusivity. US FDA, <https://www.fda.gov/downloads/drugs/developmentapprovalprocess/smallbusinessassistance/ucm447307.pdf>.
- US FDA (2019) Interacting with the FDA on complex innovative trial designs for drugs and biological products. US FDA, <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/interacting-fda-complex-innovative-trial-designs-drugs-and-biological-products>.
- US HHS (2014) *Examination of Clinical Trial Costs and Barriers for Drug Development*. US Dept. Health and Human Services, Office of the Assistant Secretary for Planning and Evaluation.
- Villar SS, Bowden J, Wason J (2018) Response-adaptive designs for binary responses: How to offer patient benefit while being robust to time trends? *Pharmaceutical statistics* 17(2):182–197.
- Villar SS, Rosenberger WF (2018) Covariate-adjusted response-adaptive randomization for multi-arm clinical trials using a modified forward looking Gittins index rule. *Biometrics* 74(1):49–57.
- Walters SJ, et al. (2017) Recruitment and retention of participants in randomised controlled trials: A review of trials funded and published by the United Kingdom Health Technology Assessment programme. *BMJ Open* 7(3):e015276, <https://bmjopen.bmj.com/content/7/3/e015276.long>.
- Willan AR (2008) Optimal sample size determination from an industry perspective based on the expected value of information. *Clin. Trials* 5:587–594.
- 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.
- World Trade Organization (2006) Pharmaceutical patents and the TRIPS agreement. https://www.wto.org/english/tratop_e/trips_e/pharma_ato186_e.htm.
- Xie J, Frazier PI, Chick SE (2016) Bayesian optimization via simulation with pairwise sampling and correlated prior beliefs. *Operations Research* 64(2):542–559.

Appendices

Appendix A summarizes 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 presents numerical comparative statics analysis. Appendix E provides formalism for the probabilistic performance measures in the experiments of section 4. Appendix F shows how the one-shot trial design of the main paper can be transformed to a response adaptive trial. Appendix G displays supplemental figures for the application to the ProFHER trial. We provide Matlab code to reproduce the results at <https://github.com/andres-alban/ValueBasedTrials>.

Appendix A: Table of principal notation

Table EC.1 Table of principal notation.	
<i>Parameter</i>	<i>Description</i>
$p_N \in [0, 1/2]$	Fraction of patients treated with HT N under the current practice
$E_N, E_S \in \mathbb{R}$	Effectiveness of technologies N and S, respectively
$C_N, C_S \in \mathbb{R}_{\geq 0}$	Patient-level costs of using technologies N and S, respectively
$\lambda \in \mathbb{R}_{> 0}$	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
$\zeta \in \mathbb{R}_{> 0}$	Incidence rate of the condition in the population
$r_{\max} \in [0, \zeta]$	Capacity of rate of recruitment
$Q_{\max} \in \mathbb{N}$	Maximum sample size for the trial
$T \in [0, T_{\max}]$	Recruitment period duration (decision variable)
$r \in [0, r_{\max}]$	Average recruitment rate to the trial (decision variable)
$Q = Tr/2$	Sample size of the trial
$\mathcal{D} \in \{M, N, S\}$	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 patients to be observed
$\Delta \in [0, H)$	Delay in observing realization of pairwise allocation (in time units)
$P(T) \in \mathbb{R}_{\geq 0}$	Number of patients to receive implemented technology once adoption decision is made at $T + \Delta$
$H \in \mathbb{R}_{> 0}$	Maximum time horizon for decision, in the case of fixed horizon
$P \in \mathbb{R}_{> 0}$	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, $\rho = \ln(1 + \rho_{yr})$
$\tilde{T}_\rho(T)$	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
$c \in \mathbb{R}_{\geq 0}$	Variable per-patient cost per patient completing the trial
$c_{\text{cap}}(r)$	Fixed setup cost of a trial to achieve recruitment rate r
$c_{\text{fix}}, c_r, \theta$	Parameters of setup cost function, e.g., $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r^\theta$
$I_{\mathcal{D}} \in \mathbb{R}_{\geq 0}$	Fixed cost of switching to technology \mathcal{D} from standard technology
$\alpha_N(T), \alpha_S(T) \in \mathbb{R}$	Expected costs per patient if technology N or S is adopted
$\Psi(z)$	Standard normal loss function, $\Psi(z) \equiv \mathbb{E}[(Z - z)^+] = \phi(z) - z(1 - \Phi(z))$

Appendix B: Derivation of partial derivatives of V

The first order conditions for interior solutions of (5), 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:

$$\begin{aligned} \frac{\partial \alpha_N(T)}{\partial T} &= -\frac{\alpha_N(T)}{P_\rho(T)} \frac{dP_\rho(T)}{dT}, & \frac{\partial \alpha_S(T)}{\partial T} &= -\frac{\alpha_S(T)}{P_\rho(T)} \frac{dP_\rho(T)}{dT}, \\ \frac{\partial \sigma_Z}{\partial T} &= \sqrt{\frac{n_0 \sigma_X^2 r}{(2n_0 + Tr)^3 T}}, & \frac{\partial \sigma_Z}{\partial r} &= \sqrt{\frac{n_0 \sigma_X^2 T}{(2n_0 + Tr)^3 r}}. \end{aligned}$$

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

$$\begin{aligned} \frac{\partial((\alpha_N(T))/\sigma_Z)}{\partial T} &= -\frac{\alpha_N(T)}{\sigma_Z} \left(\frac{1}{P_\rho(T)} \frac{dP_\rho(T)}{dT} + \frac{1}{\sigma_Z} \frac{\partial \sigma_Z}{\partial T} \right) \\ \frac{\partial((\alpha_S(T))/\sigma_Z)}{\partial T} &= -\frac{\alpha_S(T)}{\sigma_Z} \left(\frac{1}{P_\rho(T)} \frac{dP_\rho(T)}{dT} + \frac{1}{\sigma_Z} \frac{\partial \sigma_Z}{\partial T} \right) \\ \frac{\partial \sigma_Z}{\partial T} &= \frac{\sigma_X^2 r}{\sigma_Z (2n_0 + Tr)^2} = \frac{n_0 \sigma_Z}{(2n_0 + Tr)T} \\ \frac{\partial \sigma_Z}{\partial r} &= \frac{\sigma_X^2 T}{\sigma_Z (2n_0 + Tr)^2} = \frac{n_0 \sigma_Z}{(2n_0 + Tr)r}. \end{aligned}$$

Consider first

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

Similarly,

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

Thus, we get

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

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:

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

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 $f: S \rightarrow \mathbb{R}$ is upper semi-continuous at x_0 if $\limsup_{x \rightarrow x_0} f(x) \leq f(x_0)$, or equivalently, if for every $\epsilon > 0$ there is a neighborhood S' around x_0 such that $f(x) \leq f(x_0) + \epsilon$ for all $x \in S'$.

The domain of $V(T, r)$ is $D = \{(T, r) \mid 0 \leq T \leq T_{\text{max}}, 0 \leq r \leq r_{\text{max}}\}$, 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 $\{(T, r) \mid 0 < T \leq T_{\text{max}}, 0 < r \leq r_{\text{max}}\}$ 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 (9). Now, let $\delta_T > 0$ and consider the following inequality using (8):

$$V(\delta_T, r + \delta_r) = -(c_{\text{cap}}(r + \delta_r) + c(r + \delta_r) \tilde{T}_\rho(\delta_T)) + \delta_{\text{on}} [(r + \delta_r) \tilde{T}_\rho(\delta_T) / 2] (1 - 2p_{\text{N}}) \mu_0$$

$$\begin{aligned}
& + e^{-\rho(\delta_T+\Delta)} \mathbb{E} \left[(P_\rho(\delta_T)(1-p_N)Z_{\delta_T(r+\delta_r)} - I_N)^+ + (-P_\rho(\delta_T)p_N Z_{\delta_T(r+\delta_r)} - I_S)^+ \right] \\
\leq & \delta_{\text{on}}[(r+\delta_r)\tilde{T}_\rho(\delta_T)/2](1-2p_N)\mu_0 \\
& + e^{-\rho\delta_T} e^{-\rho\Delta} \mathbb{E} \left[(P_\rho(0)(1-p_N)Z_{\delta_T(r+\delta_r)} - I_N)^+ + (-P_\rho(0)p_N Z_{\delta_T(r+\delta_r)} - I_S)^+ \right].
\end{aligned}$$

The inequality holds because $c_{\text{cap}}(r) \geq 0$ and $P_\rho(T)$ is non-increasing, so $P_\rho(\delta_T) \leq P_\rho(0)$. The right-hand side of the inequality is further continuous in δ_T and is such that it converges to $V(0, r)$ as $\delta_T \rightarrow 0$. Thus, we can conclude that $V(\delta_T, r + \delta_r) \leq V(0, r) + \epsilon$ for any $\epsilon > 0$ if we make δ_T small enough, and that $V(T, r)$ 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 (4a). Now, let $\delta_r > 0$ and consider the following inequality that follows from $c_{\text{cap}}(r) \geq 0$ and $P(T)$ being non-increasing:

$$\begin{aligned}
V(T + \delta_T, \delta_r) & = -(c_{\text{cap}}(\delta_r) + c\delta_r\tilde{T}_\rho(T + \delta_T)) + \delta_{\text{on}}[(\delta_r)\tilde{T}_\rho(T + \delta_T)/2](1-2p_N)\mu_0 \\
& + e^{-\rho(T+\delta_T+\Delta)} \mathbb{E}[(P_\rho(T + \delta_T)(1-p_N)Z_{(T+\delta_T)\delta_r} - I_N)^+ \\
& + (-P_\rho(T + \delta_T)p_N Z_{(T+\delta_T)\delta_r} - I_S)^+] \\
& \leq \delta_{\text{on}}[\delta_r\tilde{T}_\rho(\delta_T)/2](1-2p_N)\mu_0 \\
& + e^{-\rho\delta_T} e^{-\rho(T+\Delta)} \mathbb{E}[(P_\rho(0)(1-p_N)Z_{(T+\delta_T)\delta_r} - I_N)^+ \\
& + (-P_\rho(0)p_N Z_{(T+\delta_T)\delta_r} - I_S)^+].
\end{aligned}$$

The right-hand side of the inequality is continuous in δ_T and δ_r and is such that it converges to $V(T, 0)$ as $\delta_T \rightarrow 0$ and $\delta_r \rightarrow 0$. Therefore, we can conclude that $V(T + \delta_T, \delta_r) \leq V(T, 0) + \epsilon$ for any $\epsilon > 0$ if we make $|\delta_T|$ and δ_r small enough, and that $V(T, r)$ is upper semi-continuous at $r = 0$. \square

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 set $\mathcal{S} = \{(T, r) \in [0, T_{\text{max}}] \times [0, r_{\text{max}}] : Tr/2 = Q\}$. If $Q = 0$, then $V(T, r) = V(0, 0)$ for all members \mathcal{S} . This follows directly from the definition of V in (4a). 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:

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

Proof of Prop. 3. Let (T^*, r^*) be an optimal solution to (5). First, observe that if $T^* = 0$ or $r^* = 0$, then $(0, r_{\text{max}})$ is also an optimal solution by the definition of V in (4a).

Now, consider the case $T^* > 0$ and $r^* > 0$ such that $Q^* = T^*r^*/2 > 0$. We will show that any trial design with the same sample size Q^* but with a larger recruitment rate ($r \geq r^*$) and shorter duration ($T = 2Q^*/r \leq T^*$) obtains equal or larger expected net gain.

The posterior mean variance, σ_Z^2 , is the same for any design that has the same sample size. The non-increasing assumption of $P(T)$ implies that $P_\rho(T)$ is non-increasing, and, hence, $P_\rho(T^*) \leq P_\rho(2Q^*/r)$, $\alpha_N(T^*) \geq \alpha_N(2Q^*/r)$, and $\alpha_S(T^*) \geq \alpha_S(2Q^*/r)$. Because $\Psi(\cdot)$ is decreasing, we can show that the post-trial rewards, R^* , under the optimal design are smaller than the post-trial rewards obtained when the recruitment rate r is larger than r^* while maintaining the same sample size Q^* :

$$\begin{aligned} R^* &= P_\rho(T^*)\sigma_Z \left[(1-p_N)\Psi\left(\frac{\alpha_N(T^*)-\mu_0}{\sigma_Z}\right) + p_N\Psi\left(\frac{\alpha_S(T^*)+\mu_0}{\sigma_Z}\right) \right] \\ &\leq P_\rho(2Q^*/r)\sigma_Z \left[(1-p_N)\Psi\left(\frac{\alpha_N(2Q^*/r)-\mu_0}{\sigma_Z}\right) + p_N\Psi\left(\frac{\alpha_S(2Q^*/r)+\mu_0}{\sigma_Z}\right) \right]. \end{aligned}$$

The cost and online rewards may however be larger with a larger recruitment rate. Therefore, it is not clear that increasing the recruitment rate obtains larger expected net gain, which is what we will now prove.

To simplify notation, let $C = c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2$ be the per patient cost net of online rewards, such that the optimal expected net gain is given by

$$V^* = -c_{\text{cap}} - Cr^*\tilde{T}_\rho(T^*) + e^{-\rho(T^*+\Delta)}R^*.$$

Because $Q^* > 0$, we know that the optimal expected net gain is positive: $c_{\text{cap}} + Cr^*\tilde{T}_\rho(T^*) \leq e^{-\rho(T^*+\Delta)}R^*$. The expected net gain for the design that maintains the same sample size but uses a larger recruitment rate can be bounded below by the function $f(r)$:

$$V(Q^*/r, r) \geq f(r) = -c_{\text{cap}} - Cr\tilde{T}_\rho(2Q^*/r) + e^{-\rho(2Q^*/r+\Delta)}R^*$$

because we have shown that the post-trial rewards are larger when the recruitment rate is larger.

We show that the derivative of $f(r)$ is non-negative for $r \geq r^*$, which will show that expected net gain is non-decreasing as the recruitment rate increases. Basic calculus gives us

$$f'(r) = -\frac{C}{\rho} \left(1 - e^{-2\rho Q^*/r} - \frac{2\rho Q^*}{r} e^{-2\rho Q^*/r} \right) + \frac{2\rho Q^*}{r^2} e^{-2\rho Q^*/r} e^{-\rho\Delta} R^*,$$

and because the optimal expected net gain is positive, we can use $e^{-\rho\Delta}R^* \geq Cr^*(1 - e^{-2\rho Q^*/r^*})/(\rho e^{-2\rho Q^*/r^*})$ to obtain

$$\begin{aligned} f'(r) &\geq -\frac{C}{\rho} \left(1 - e^{-2\rho Q^*/r} - \frac{2\rho Q^*}{r} e^{-2\rho Q^*/r} \right) + \frac{2\rho Q^*}{r^2} e^{-2\rho Q^*/r} \frac{Cr^*(1 - e^{-2\rho Q^*/r^*})}{\rho e^{-2\rho Q^*/r^*}} \\ &= -\frac{C}{\rho} \left(1 - e^{-\rho T} - \rho T e^{-\rho T} - \frac{\rho T^2}{T^*} e^{-\rho T} \frac{1 - e^{-\rho T^*}}{e^{-\rho T^*}} \right). \end{aligned}$$

The second line inputs $T = 2Q^*/r$ and $T^* = 2Q^*/r^*$ and is a function of r through T . To show that $f'(r) \geq 0$ for $r \geq r^*$, it is sufficient to show that

$$g(T, T^*) = \left(1 - e^{-\rho T} - \rho T e^{-\rho T} - \frac{\rho T^2}{T^*} e^{-\rho T} \frac{1 - e^{-\rho T^*}}{e^{-\rho T^*}} \right) \leq 0$$

for any $T \leq T^*$.

It can be shown that $g(T, T^*)$ is increasing in T^* by taking partial derivatives with respect to T^* . Thus, it is sufficient to show that $g(T, T) = 1 - e^{-\rho T} - \rho T \leq 0$, which can be easily shown by taking derivatives with respect to T . We can conclude that $f(r)$ is nondecreasing.

We have shown that the expected net gain is nondecreasing when the recruitment rate is increased and duration shortened to maintain the same sample size. Thus, we have shown that the design given $T = T^* r^* / r_{\max}$ and $r = r_{\max}$ obtains at least the same expected net gain as the optimal design: $V(T^* r^* / r_{\max}, r_{\max}) \geq V(T^*, r^*)$. Because (T^*, r^*) is optimal, we have established that $V(T^* r^* / r_{\max}, r_{\max}) = V(T^*, r^*)$ and that an optimal solution with $r = r_{\max}$ exists. \square

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

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 (5) presented in section 3.5. For simplicity, we present the results with undiscounted rewards for Cases I and II ($c_{\text{cap}}(r) = c_{\text{cap}}$), i.e., it is sufficient to optimize only T for a fixed r ($V^* = \max_{T \in [0, T_{\max}]} V(T, r)$). The results are, however, as general as presented in section 3.5.

We aim to find the sign of the derivatives $dV^*/db = \partial V(T^*, r_{\max})/\partial b$ and $dV(T^*, r_{\max})/db = -(\partial^2 V(T^*, r_{\max})/\partial T \partial b)/(\partial^2 V(T^*, r_{\max})/\partial T^2)$. The right-hand side of the second equation only holds when $\partial^2 V(T^*, r_{\max})/\partial T^2 < 0$. While it is guaranteed to be non-positive at an interior solution by the second order necessary conditions (Andreasson et al. 2007, Theorem 4.17), we assume that it is strictly negative to exclude the rare cases when it is zero. Thus, the algebra we show in this section only shows the sign of the numerator $\partial^2 V(T^*, r)/\partial T \partial b$.

C.3.1. Post-adoption population. For fixed patient pool, the optimal expected net gain is increasing in P :

$$\begin{aligned} \frac{\partial V(T, r)}{\partial P} &= (1 - p_N) \sigma_Z \Psi \left(\frac{\alpha_N - \mu_0}{\sigma_Z} \right) + p_N \sigma_Z \Psi \left(\frac{\alpha_S + \mu_0}{\sigma_Z} \right) \\ &\quad + (1 - p_N) \alpha_N \left(1 - \Phi \left(\frac{\alpha_N - \mu_0}{\sigma_Z} \right) \right) + p_N \alpha_S \left(1 - \Phi \left(\frac{\alpha_S + \mu_0}{\sigma_Z} \right) \right) > 0. \end{aligned}$$

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

$$\begin{aligned} \frac{\partial^2 V(T, r)}{\partial T \partial P} &= \frac{\partial \sigma_Z}{\partial T} \left[(1 - p_N) \phi \left(\frac{\alpha_N - \mu_0}{\sigma_Z} \right) + p_N \phi \left(\frac{\alpha_S + \mu_0}{\sigma_Z} \right) \right. \\ &\quad \left. + (1 - p_N) \frac{\alpha_N (\alpha_N - \mu_0)}{\sigma_Z^2} \phi \left(\frac{\alpha_N - \mu_0}{\sigma_Z} \right) + p_N \frac{\alpha_S (\alpha_S + \mu_0)}{\sigma_Z^2} \phi \left(\frac{\alpha_S + \mu_0}{\sigma_Z} \right) \right]. \end{aligned}$$

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 $\alpha_N < \mu_0$ there is a negative contribution from the third term but positive from all other terms. Similarly, when $\alpha_S < -\mu_0$ 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 $-\alpha_S \leq \mu_0 \leq \alpha_N$, 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 the time horizon H . The results are very similar to fixed patient pool. The expected net gain is again increasing in H :

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

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

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

Comparing this expression to that for $\partial^2 V(T, r) / \partial T \partial P$, 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.

C.3.2. Discount rate ρ . To show the effect of the discount rate, we need to assume here discounted rewards. Instead, to make the algebraic expressions more manageable, we assume that $p_N = 0$ noting that the results hold equally for any p_N . We obtain

$$\begin{aligned} \frac{\partial V(T, r)}{\partial \rho} = & (-cr + \delta_{\text{on}} r (1 - 2p_N) \mu_0 / 2) \frac{\partial \tilde{T}_\rho(T)}{\partial \rho} - (T + \Delta) e^{-\rho(T+\Delta)} P_\rho(T) \sigma_Z \Psi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) \\ & + e^{-\rho(T+\Delta)} \frac{\partial P_\rho(T)}{\partial \rho} \sigma_Z \Psi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) + e^{-\rho(T+\Delta)} \left(1 - \Phi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) \right) \alpha_N(T) \frac{\partial P_\rho(T)}{\partial \rho} \end{aligned}$$

The second line above is negative so we only need to show that the first line is negative:

$$\begin{aligned} (-cr + \delta_{\text{on}} r (1 - 2p_N) \mu_0 / 2) \frac{\partial \tilde{T}_\rho(T)}{\partial \rho} & \leq (T + \Delta) e^{-\rho(T+\Delta)} P_\rho(T) \sigma_Z \Psi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) \\ (cr - \delta_{\text{on}} r (1 - 2p_N) \mu_0 / 2) \tilde{T}_\rho(T) \left(-\frac{1}{\tilde{T}_\rho(T)} \frac{\partial \tilde{T}_\rho(T)}{\partial \rho} \right) & \leq (T + \Delta) e^{-\rho(T+\Delta)} P_\rho(T) \sigma_Z \Psi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) \end{aligned}$$

To show that the first line is negative at the optimal design, we need to make use of the non-negativity of the optimal expected net gain, $V(T^*, r^*) \geq 0$, equivalent to

$$\begin{aligned} (c_{\text{cap}} + cr^* - \delta_{\text{on}}r^*(1 - 2p_{\text{N}})\mu_0/2)\tilde{T}_\rho(T^*) &\leq e^{-\rho(T^*+\Delta)}P_\rho(T^*)\sigma_Z^*\Psi\left(\frac{\alpha_{\text{N}}(T^*) - \mu_0}{\sigma_Z^*}\right) \\ (cr^* - \delta_{\text{on}}r^*(1 - 2p_{\text{N}})\mu_0/2)\tilde{T}_\rho(T^*) &\leq e^{-\rho(T^*+\Delta)}P_\rho(T^*)\sigma_Z^*\Psi\left(\frac{\alpha_{\text{N}}(T^*) - \mu_0}{\sigma_Z^*}\right) \end{aligned}$$

Thus, the problem simplifies to showing that

$$\begin{aligned} -\frac{1}{\tilde{T}_\rho(T)}\frac{\partial\tilde{T}_\rho(T)}{\partial\rho} &\leq T + \Delta \\ \frac{1 - e^{-\rho T}(\rho T + 1)}{\rho(1 - e^{-\rho T})} &\leq T + \Delta, \end{aligned}$$

which requires some simple additional algebra to show that it indeed holds.

The comparative statics for T^* require the calculation of $\partial^2V(T, r)/\partial T\partial\rho$. While it is straightforward to calculate, it results in a long expression that does not elucidate any interesting results and we do not show it here.

C.3.3. Effective sample size of the prior distribution. To analyze the sensitivity of V to n_0 , we compute

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

where $\partial\sigma_Z/\partial n_0 = -\sigma_Z^2(2n_0 + Tr)/(2n_0(n_0 + Tr)) < 0$. 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 analyze the sensitivity on the optimal recruitment duration, we obtain the following derivative:

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

The terms in the second and third line 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 $n_0 < Tr/4$ because $\partial^2\sigma_Z/\partial T\partial n_0 = \sigma_Z(-4n_0 + Tr)/4(2n_0 + Tr)^2T$. Therefore, T^* often increases with n_0 when n_0 is small, as we observe in the numerical results of Appendix D.

C.3.4. Sampling standard deviation σ_X The analysis here is similar to the analysis for n_0 :

$$\frac{\partial V(T, r)}{\partial \sigma_X} = P \frac{\partial \sigma_Z}{\partial \sigma_X} \left[(1 - p_N) \phi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) + p_N \phi \left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z} \right) \right] > 0,$$

where $\partial \sigma_Z / \partial \sigma_X = \sqrt{Tr / (2n_0(n_0 + Tr/2))} > 0$.

$$\begin{aligned} \frac{\partial^2 V(T, r)}{\partial T \partial \sigma_X} &= P \frac{\partial^2 \sigma_Z}{\partial T \partial \sigma_X} \left[(1 - p_N) \phi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) + p_N \phi \left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z} \right) \right] \\ &\quad + (1 - p_N) P \phi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) \frac{(\alpha_N(T) - \mu_0)^2}{\sigma_Z^3} \frac{\partial \sigma_Z}{\partial T} \frac{\partial \sigma_Z}{\partial \sigma_X} \\ &\quad + p_N P \phi \left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z} \right) \frac{(\alpha_S(T) + \mu_0)^2}{\sigma_Z^3} \frac{\partial \sigma_Z}{\partial T} \frac{\partial \sigma_Z}{\partial \sigma_X} > 0. \end{aligned}$$

Here we obtain a definitive sign because $\partial^2 \sigma_Z / \partial T \partial \sigma_X = \sqrt{n_0 r / ((2n_0 + Tr)^3 T)} > 0$.

C.3.5. Per-patient cost c . We can easily obtain

$$\begin{aligned} \frac{\partial V}{\partial c} &= -Tr < 0 \\ \frac{\partial^2 V}{\partial T \partial c} &= -r < 0, \end{aligned}$$

which confirms the intuitive result that the optimal expected net gain and optimal enrolled patients in Cases I–III is decreasing with variable per-patient cost.

C.3.6. Delay in observing outcomes Δ . We obtain

$$\frac{\partial V}{\partial \Delta} = -\rho e^{-\rho(T+\Delta)} \sigma_Z P_p \left[(1 - p_N) \Psi \left(\frac{\alpha_N - \mu_0}{\sigma_Z} \right) + p_N \Psi \left(\frac{\alpha_S + \mu_0}{\sigma_Z} \right) \right] \leq 0$$

C.3.7. Cost of recruitment c_r . We obtain

$$\frac{\partial V}{\partial c_r} = -r^\theta \leq 0$$

and for Case III trials (only recruitment rate is optimized), we obtain

$$\frac{\partial^2 V}{\partial r \partial c_r} = -\theta r^{\theta-1} \leq 0.$$

C.3.8. Fraction of patients in new technology p_N . It is straightforward to check that

$$\begin{aligned} \frac{\partial V(T, r)}{\partial p_N} &= -2\delta_{\text{on}} Tr \mu_0 + P(T) \sigma_Z \left(\Psi \left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z} \right) - \Psi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) \right) \\ &\quad + P(T) \left[\alpha_S(T) \left(1 - \Phi \left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z} \right) \right) - \alpha_N(T) \left(1 - \Phi \left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z} \right) \right) \right]. \end{aligned}$$

There are three summands in this expression. The first is the online learning effect. The second is the effect due to 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 $P(T)\alpha_N(T)$ and $P(T)\alpha_S(T)$. When $P(T)\sigma_Z$ dominates, the effect depends on the expected gains of adopting technologies N or S; if $\Psi((\alpha_S(T) + \mu_0)/\sigma_Z) > \Psi((\alpha_N(T) - \mu_0)/\sigma_Z)$, 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 $\partial^2 V(T, r)/\partial T \partial p_N$ is a long expression that does not elucidate any interesting insights. We present a special scenario, $I_N = I_S = 0$, with some interesting results. Note that $\alpha_N(T) = \alpha_S(T) = 0$ and we obtain

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

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. The optimal expected net gain is independent of p_N under the former, but not the latter, condition.

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 $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 > 0$ and that there is no discounting, $\rho = 0$. The proofs further assume that $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r$ 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 $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 \leq 0$ 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, analyze the more interesting setting where $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 > 0$. Because Q^* grows unbounded with P (shown below), we let $r_{\text{max}} = \infty$ and $T_{\text{max}} = \infty$, i.e., there is no upper bound on the decision variables.

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

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

$$\lim_{P \rightarrow \infty} \frac{V^*}{P} = (1 - p_N)\sigma_0 \Psi(-\mu_0/\sigma_0) + p_N \sigma_0 \Psi(\mu_0/\sigma_0).$$

Proof of Lemma EC.1. For functions f and g , we use the notation $f(P) \sim g(P)$ to denote $\lim_{P \rightarrow \infty} f(P)/g(P) = 1$. 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

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

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 . The expression $\partial V(T, r)/\partial T = 0$ was derived in Appendix B. By rearranging terms, we obtain

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

$$= \sqrt{\frac{(2n_0 + rT^*(P))^3 T^*(P)}{n_0 \sigma_X^2 r}} (cr - \delta_{\text{on}} r (1 - 2p_N) \mu_0 / 2).$$

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

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

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

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

Combining results, we obtain

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

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 $\lim_{P \rightarrow \infty} T^* = \infty$, we know that $\lim_{P \rightarrow \infty} \sigma_Z^* = \sigma_X / \sqrt{n_0} = \sigma_0$. Then,

$$V^* = -c_{\text{cap}} - crT^* + \delta_{\text{on}} (1 - 2p_N) \mu_0 r T^* / 2 + P \sigma_Z^* \left((1 - p_N) \Psi \left(\frac{\alpha_N - \mu_0}{\sigma_Z^*} \right) + p_N \Psi \left(\frac{\alpha_S + \mu_0}{\sigma_Z^*} \right) \right)$$

$$V^* \sim -crT^* + \delta_{\text{on}} (1 - 2p_N) \mu_0 r T^* / 2 + P \sigma_Z^* \left((1 - p_N) \Psi \left(\frac{-\mu_0}{\sigma_Z^*} \right) + p_N \Psi \left(\frac{\mu_0}{\sigma_Z^*} \right) \right)$$

$$V^* \sim -crT^* + \delta_{\text{on}} (1 - 2p_N) \mu_0 r T^* / 2 + P \sigma_0 \left((1 - p_N) \Psi \left(\frac{-\mu_0}{\sigma_0} \right) + p_N \Psi \left(\frac{\mu_0}{\sigma_0} \right) \right)$$

$$V^* \sim (-c + \delta_{\text{on}} (1 - 2p_N) \mu_0 / 2) \left(\frac{\sqrt{n_0 \sigma_X^2} \phi(\sqrt{n_0} \mu_0 / \sigma_X)}{(c - \delta_{\text{on}} (1 - 2p_N) \mu_0 / 2) r^2} \right)^{1/2} \sqrt{P} + P \sigma_0 \left((1 - p_N) \Psi \left(\frac{-\mu_0}{\sigma_0} \right) + p_N \Psi \left(\frac{\mu_0}{\sigma_0} \right) \right)$$

$$V^* \sim P \sigma_0 \left((1 - p_N) \Psi \left(\frac{-\mu_0}{\sigma_0} \right) + p_N \Psi \left(\frac{\mu_0}{\sigma_0} \right) \right),$$

and the last line is the statement in the lemma. \square

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

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

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

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

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

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

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

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

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 $\lim_{H \rightarrow \infty} T^*(H) = \infty$, $\lim_{H \rightarrow \infty} (H - T^*(H) - \Delta) = \infty$, 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:

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

By rearranging and some additional asymptotic approximations, we obtain

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

The desired result is obtained by rearranging the terms and setting $r = r_{\text{max}}$, 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^* :

$$\begin{aligned} V^* &= -c_{\text{cap}}(r) - crT^* + \delta_{\text{on}}(1 - 2p_N)\mu_0rT^*/2 + P(T^*)\sigma_Z^* \left((1 - p_N)\Psi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) + p_N\Psi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right) \\ V^* &\sim -crT^* + \delta_{\text{on}}(1 - 2p_N)\mu_0rT^*/2 + \zeta(H - T^* - \Delta)\sigma_0 \left((1 - p_N)\Psi \left(\frac{-\mu_0}{\sigma_0} \right) + p_N\Psi \left(\frac{\mu_0}{\sigma_0} \right) \right) \\ V^* &\sim \zeta(H - \Delta)\sigma_0 \left((1 - p_N)\Psi \left(\frac{-\mu_0}{\sigma_0} \right) + p_N\Psi \left(\frac{\mu_0}{\sigma_0} \right) \right). \quad \square \end{aligned}$$

C.4.3. Case III: affine setup costs, no discounting, fixed patient pool. If $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 \leq 0$, then the first order condition is never satisfied and it is optimal to sample infinitely. In that scenario, a solution to (5) does not exist unless r is constrained and $r^* = r_{\text{max}}$. We therefore analyze the setting where $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 > 0$. Because the optimal recruitment rate grows unboundedly with P (as shown below), we assume that $r_{\text{max}} = \infty$, i.e., no upper bound on r . To obtain asymptotic results on Q^* , we only need to multiply the results of r^* with $T_{\text{max}}/2$, as Prop. 4 shows that $T = T_{\text{max}}$ is optimal.

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

$$\begin{aligned} \lim_{P \rightarrow \infty} \frac{r^*}{\sqrt{P}} &= \left(\frac{\sqrt{n_0}\sigma_X^2\phi(\mu_0/\sigma_0)}{[cT_{\text{max}} + c_r - \delta_{\text{on}}(1 - 2p_N)\mu_0T_{\text{max}}/2]} \right)^{1/2}, \quad (\text{EC.4}) \\ \lim_{P \rightarrow \infty} \frac{V^*}{P} &= \sigma_0[\Psi(\mu_0/\sigma_0) + (1 - p_N)\mu_0/\sigma_0]. \end{aligned}$$

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

$$\begin{aligned} P \left(\frac{1 - p_N}{\sqrt{2\pi}} e^{-\frac{n_0(2n_0 + r^*(P)T_{\text{max}})}{2\sigma_X^2 r^*(P)T_{\text{max}}}(\alpha_N - \mu_0)^2} + \frac{p_N}{\sqrt{2\pi}} e^{-\frac{n_0(2n_0 + r^*(P)T_{\text{max}})}{2\sigma_X^2 r^*(P)T_{\text{max}}}(\alpha_S + \mu_0)^2} \right) \\ = \sqrt{\frac{(2n_0 + r^*(P)T_{\text{max}})^3 r^*(P)}{n_0\sigma_X^2 T_{\text{max}}}} (c_r + cT_{\text{max}} - \delta_{\text{on}}T_{\text{max}}(1 - 2p_N)\mu_0/2). \end{aligned}$$

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. \square

C.4.4. Case IV: affine setup costs, no discounting, fixed horizon. If $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 < 0$ in this case, the effective cost of sampling is less than zero, and, given the trial is run, it is optimal to sample infinitely for large H . We, therefore, analyze the more interesting setting where $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 \geq 0$. Notice that, unlike the previous cases, $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 = 0$ is considered here and gives interesting results that differ substantially from the scenario when $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 > 0$. 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 $\rho = 0$, $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r$, where $c_r > 0$, $P(T) = \zeta(H - T - \mathbf{1}_{T > 0}\Delta)$, $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 \geq 0$, and $T_{\max} = \infty$ and $r_{\max} = \infty$, then*

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

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

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

if $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 = 0$, and

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

if $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 = 0$, 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 $T^* r^* > 0$ 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

$$\begin{aligned} & (c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2)r^* + (1 - p_N) \left[\zeta \alpha_N(T^*) \left(1 - \Phi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) \right) + \sigma_Z^* \zeta \Psi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) \right] \\ & + p_N \left[\zeta \alpha_S(T^*) \left(1 - \Phi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right) + \sigma_Z^* \zeta \Psi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right] \\ & = \zeta(H - T^* - \Delta) \frac{\partial \sigma_Z^*}{\partial T} \left[(1 - p_N) \phi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) + p_N \phi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right] \end{aligned} \quad (\text{EC.5})$$

and

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

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:

$$(1 - p_N) \left[\zeta \alpha_N(T^*) \left(1 - \Phi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) \right) + \sigma_Z^* \zeta \Psi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) \right] \\ + p_N \left[\zeta \alpha_S(T^*) \left(1 - \Phi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right) + \sigma_Z^* \zeta \Psi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right] = c_r \frac{r^*}{T^*} \quad (\text{EC.7})$$

We first show that $\lim_{H \rightarrow \infty} H - T^* - \Delta = \infty$ by means of contradiction. Assuming that $\lim_{H \rightarrow \infty} H - T^* - \Delta < \infty$ implies that $\lim_{H \rightarrow \infty} T^* = \infty$. Then, the right-hand side of (EC.5) approaches zero as $H \rightarrow \infty$. 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 $\lim_{H \rightarrow \infty} T^* = \infty$ by contradiction. Assuming that $\lim_{H \rightarrow \infty} T^* < \infty$, (EC.5) implies that $\lim_{H \rightarrow \infty} r^* = \infty$. 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 $\lim_{H \rightarrow \infty} T^* = \infty$, (EC.7) implies $\lim_{H \rightarrow \infty} r^* = \infty$. In fact, (EC.7) implies

$$r^* = \left[\zeta (1 - p_N) \left[\alpha_N(T^*) \left(1 - \Phi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) \right) + \sigma_Z^* \Psi \left(\frac{\alpha_N(T^*) - \mu_0}{\sigma_Z^*} \right) \right] \right. \\ \left. + \zeta p_N \left[\alpha_S(T^*) \left(1 - \Phi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right) + \sigma_Z^* \Psi \left(\frac{\alpha_S(T^*) + \mu_0}{\sigma_Z^*} \right) \right] \right] \frac{T^*}{c_r} \\ \sim K T^*, \quad (\text{EC.8})$$

where $K = \sigma_X \zeta (\Psi(\mu_0/\sigma_0) + (1 - p_N)\mu_0/\sigma_0) / (\sqrt{n_0} c_r)$. The asymptotic equivalence follows due to $\lim_{H \rightarrow \infty} T^* = \infty$ and $\lim_{H \rightarrow \infty} r^* = \infty$.

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

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

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

and with further simplifications

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

From (EC.8) it follows

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

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

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

$$Kc_r \sim \zeta(H - T^* - \Delta) \sqrt{\frac{n_0\sigma_X^2 K}{(2n_0 + K(T^*)^2)^3}} \phi(\sqrt{n_0}\mu_0/\sigma_X),$$

and with further simplifications

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

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 $c - \delta_{\text{on}}(1 - 2p_N)\mu_0/2 = 0$.

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

$$V^* \sim \zeta(H - \Delta) ((1 - p_N)\sigma_0\Psi(-\mu_0/\sigma_0) + p_N\sigma_0\Psi(\mu_0/\sigma_0)). \quad \square$$

C.5. Proofs for asymptotic results of Prop. 6 in section 3.6.2

Proof of Prop. 6 It is easy to check that $\lim_{P \rightarrow \infty} P_\rho = \zeta/\rho$ with fixed patient pool and that $\lim_{H \rightarrow \infty} P_\rho(T) = \zeta/\rho$ with fixed horizon, where convergence is uniform on the bounded domain $0 \leq T \leq T_{\text{max}}$. Similarly, $\alpha'_N = I_N\rho/((1 - p_N)\zeta) = \lim_{P \rightarrow \infty} \alpha_N = \lim_{H \rightarrow \infty} \alpha_N(T)$, and $\alpha'_S = I_S\rho/(p_N\zeta) = \lim_{P \rightarrow \infty} \alpha_S = \lim_{H \rightarrow \infty} \alpha_S(T)$. Consider first $Tr > 0$:

$$\begin{aligned} |V_\infty(T, r) - V(T, r)| &= e^{-\rho(T+\Delta)}\sigma_Z \left| \frac{\zeta}{\rho} - P_\rho(T) \right| \left((1 - p_N) \left| \Psi\left(\frac{\alpha'_N - \mu_0}{\sigma_Z}\right) - \Psi\left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z}\right) \right| \right. \\ &\quad \left. + p_N \left| \Psi\left(\frac{\alpha'_S + \mu_0}{\sigma_Z}\right) - \Psi\left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z}\right) \right| \right) \\ &\leq \sigma_0 \left| \frac{\zeta}{\rho} - P_\rho(T) \right| \left((1 - p_N) \left| \Psi\left(\frac{\alpha'_N - \mu_0}{\sigma_Z}\right) - \Psi\left(\frac{\alpha_N(T) - \mu_0}{\sigma_Z}\right) \right| \right. \\ &\quad \left. + p_N \left| \Psi\left(\frac{\alpha'_S + \mu_0}{\sigma_Z}\right) - \Psi\left(\frac{\alpha_S(T) + \mu_0}{\sigma_Z}\right) \right| \right), \end{aligned}$$

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

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

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

$$\begin{aligned} |V_\infty(T, r) - V(T, r)| &= \left| \frac{\zeta}{\rho} \max\{0, (1 - p_N)\mu_0 - I_N\rho/\zeta, -p_N\mu_0 - I_S\rho/\zeta\} \right. \\ &\quad \left. - P_\rho(T) \max\{0, (1 - p_N)\mu_0 - I_N/P_\rho(T), -p_N\mu_0 - I_S/P_\rho(T)\} \right| \\ &\leq \left| \frac{\zeta}{\rho} - P_\rho(T) \right| \max\{0, (1 - p_N)\mu_0 - I_N\rho/\zeta, -p_N\mu_0 - I_S\rho/\zeta\}. \end{aligned}$$

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

Due to uniform convergence, it follows that $\lim_{P \rightarrow \infty} V_P^* = \lim_{H \rightarrow \infty} V_H^* = \max_{T,r} V_\infty(T, r)$. We cannot guarantee a unique maximizer of $V_\infty(T, r)$. However, uniform convergence also guarantees that, if a unique maximizer $(T_\infty, r_\infty) = \arg \max_{T,r} V_\infty(T, r)$ exists, then $\lim_{P \rightarrow \infty} T_P^* = \lim_{H \rightarrow \infty} T_H^* = T_\infty$ and $\lim_{P \rightarrow \infty} r_P^* = \lim_{H \rightarrow \infty} r_H^* = r_\infty$. Thus, $\lim_{P \rightarrow \infty} Q_P^* = \lim_{H \rightarrow \infty} Q_H^* = r_\infty T_\infty / 2$. \square

Appendix D: Numerical comparative statics

This appendix provides the technical details for the numerical investigation of the sign of the following derivatives with respect to parameter b : dV^*/db , dT^*/db , dr^*/db , and dQ^*/db . The results are based on exploring the parameter space by evaluating the optimal design and expected net gain at representative points in the parameter space and recording the finite difference computation of the derivatives at each point using the envelope and implicit function theorems.

We consider only offline learning ($\delta_{\text{on}} = 0$) because it is equivalent to a shift in per-patient cost as discussed in section 3.5. We consider both fixed patient pool and fixed horizon, and let $P = \zeta H$ for fixed patient pool. To explore different shapes of setup costs, we consider polynomial functions $c_{\text{cap}}(r) = c_r r^\theta$ and explore values for $c_r \geq 0$ and $\theta \geq 1$. We assume no fixed costs because those only affect the decision whether to run the trial or not, while we are interested in the effects when the trial is run. We use $\theta \geq 1$ to consider the more realistic scenario of convex setup costs because recruitment becomes harder as the trial size increases.

To explore the parameter space, we split the procedure in the four cases described in section 3.4 which are defined based on the setup cost function $c_{\text{cap}}(r)$, the post-adoption population, either fixed patient pool or fixed horizon, and discount rate ρ .

We consider the range of parameters in Table EC.2 chosen with the following reasoning. We need to set a scale for value/money and time. We set the scale for value and money by fixing the sampling cost $c = 1,000$ and sampling standard deviation $\sigma_X = 1,000$; this is without loss of

generality because it is easy to check that a change in these parameters is equivalent to a scaled problem with appropriate compensation in other parameters¹³:

$$V(T, r; \beta c, c_r, \zeta) = \beta V(T, r; c, c_r/\beta, \zeta/\beta)$$

$$V(T, r; \gamma \sigma_X, \zeta, \mu_0) = V(T, r; \sigma_X, \gamma \zeta, \mu_0/\gamma)$$

$$V(T, r; \beta c, \gamma \sigma_X, c_r, \zeta, \mu_0, I_N, I_S) = \beta V(T, r; c, \sigma_X, c_r/\beta, \zeta \gamma/\beta, \mu_0/\gamma, I_N/\beta, I_S/\beta)$$

for any $\beta > 0$ and $\gamma > 0$. We choose 1,000 as the fixed value to keep the units in the same scale as the application in section 4 for ease of interpretation.

We set the scale of time by fixing the time horizon at $H = 200$. We choose 200 again to keep the units in the same scale as the application in section 4, where each unit is a month.

Thus, we set all parameters in terms of c , σ_X , and H , using base case values similar in magnitude with the application in section 4, and exploring values one or two orders of magnitudes smaller or larger than the base case values.

Table EC.2 includes the parameters for all four cases but we present the results for each case independently because some parameters are only relevant for some cases (e.g., the discount rate is not relevant for Cases I and II because the discount rate needs to be zero) and to identify trends that differ in different cases. There are 10 degrees of freedom in these 13 parameters, after normalization of the values of c , σ_X , and H .

Parameter	Range of values
$c_{\text{cap}}(r) = c_r r^\theta$	$c_r = 0; 1,000; 100; 10,000, \theta = 1, 2, 3$
H ($P = \zeta H$)	200
Post-adoption	fixed patient pool, fixed horizon
ρ	$0, 0.5, 0.1, 0.01/H$
σ_X	1,000
n_0	1, 0.1, 10
μ_0	$0; -1,000; 1,000$
c	1,000
I_N, I_S	$0, 0.01, 0.1 \times P \sigma_X / \sqrt{n_0}$
p_N	0, 0.2, 0.4
Δ	$0.1, 0.01, 0.3 \times H$
ζ	$100,000; 10,000; 1,000,000/H$
$T_{\text{max}}, r_{\text{max}}$	$H - \Delta, 200$

We report the fraction of instances where the derivatives are larger than zero. We exclude instances where the optimal design is found at a boundary, i.e., $Q^* = 0$, or $T^* = T_{\text{max}}$ (for Cases I,

¹³The notation $V(T, r; b)$ makes the dependence of V on the parameters b explicit.

II, and IV), or $r^* = r_{\max}$ (for Cases I, III, and IV). We exclude these instances because the derivatives of the design variables are zero at these points (small changes cannot push the design out of extreme points) and we do not want the fractions to be biased by these instances.

Two computational issues are present in this analysis. First, the optimization problem to find T^* and r^* is nonconvex. We find the optimal design by initializing local optimization algorithms¹⁴ at ten random starting points. To validate that errors in the optimization algorithm do not significantly bias the results, we compared the optimal design for 10,000 randomly selected instances using 100 vs. 10 random starting points and found a discrepancy in the optimal value in less than 0.2% instances. Second, we use finite centered differences to compute the derivatives at the optimal points. These computational issues explain a slight inconsistency between a few of the theoretical results in Appendix C.3 above and three of the entries in the numerical results of Table EC.4 and EC.6 below. However, the errors are not large enough to alter the insights of this analysis.

In Table EC.3, we summarize the results of the numerical comparative statics for Case IV trials. The “number of instances” column represents the number of instances with an optimal design found in the interior of the domain. The number of instances for P and H are smaller than the rest because they are only valid for fixed patient pool and fixed horizon, respectively. The other derivatives pool the results for fixed patient pool and fixed horizon together. The number of parameter combinations tested for Case IV were 137,781 and we found an optimal design in the interior of the domain for 56,272 (41%). We discuss these results in section 3.5.

Table EC.3 Numerical comparative statics for Case IV trials

Parameter b	Number of instances	$dV^*/db > 0$	$dT^*/db > 0$	$dr^*/db > 0$	$dQ^*/db > 0$
P	28,224	100%	1.3%	99.3%	98.7%
H	28,048	100%	100%	100%	100%
ρ	56,272	0.0%	2.0%	54.9%	0.0%
Δ	56,272	0.0%	49.8%	0.0%	0.0%
n_0	56,272	0.0%	99.9%	11.7%	96.2%
σ_X	56,272	100%	21.8%	100%	97.6%
c	56,272	0.0%	0.0%	0.0%	0.0%
c_r	56,272	0.0%	99.8%	0.0%	0.2%

In Table EC.4, we present the results for Case I trials. Because Δ has no impact on the expected net gain in such trials, we only considered $\Delta = 0$. The number of parameter combinations tested was 729 of which 648 (89%) were in the interior of the domain. We see the presence of numerical error due to the optimization and numerical computation of derivatives: $dV^*/d\sigma_X$ and $dQ^*/d\sigma_X$ are theoretically greater than zero but we find 1.1% and 1.9%, respectively, of instances where this does not hold and dV^*/dn_0 is theoretically smaller than zero but we find 0.2% instances where

¹⁴We use the MATLAB[®] `fmincon` function with the SQP algorithm.

this does not hold. This provides further evidence that the errors induced by the analysis exist but are small enough not to affect the insights of the analysis.

The results yield similar insights to those yielded by the discussion of Case IV trials with some notable differences. First, 92.0% of instances have $dQ^*/dP > 0$ compared to 98.7% in Case IV trials. While both have the same direction in a majority of instances, the proportion is significantly different. Second, 72.7% of instances have $dQ^*/dn_0 > 0$ compared with 96.2% in Case IV trials. The difference between Case I and IV may be due to the optimal sample size for Case I (mean=1,630) being much larger than the that for Case IV (mean=210).

Table EC.4 Numerical comparative statics for Case I trials

Parameter b	Number of instances	$dV^*/db > 0$	$dQ^*/db > 0$
P	648	100%	92.0%
n_0	648	0.2%	72.7%
σ_X	648	98.9%	98.1%
c	648	0.0%	0.0%

In Table EC.5, we present the results for Case II trials. Of 15,309 parameter combinations, we find interior solutions for 10,957 (72%). We see small errors in the expected net gain derivative with respect to σ_X . The most notable difference with the results for Case IV is that $dT^*/d\sigma_X > 0$ for 73.2% of instances compared to 97.6% of Case IV in which $dQ^*/d\sigma_X > 0$.

Table EC.5 Numerical comparative statics for Case II trials

Parameter b	Number of instances	$dV^*/db > 0$	$dT^*/db > 0$
P	5,093	100%	63.2%
H	5,864	100%	100%
ρ	10,957	0.0%	0.0%
Δ	10,957	0.0%	0.0%
n_0	10,957	0.0%	97.7%
σ_X	10,957	100%	73.2%
c	10,957	0.0%	0.0%

Finally, Table EC.6 presents the results for Case III trials. Of 6,561 parameter combinations, we find interior solutions for 5,624 (86%). We have small numerical errors that do not match the theoretical results for the derivatives with respect to σ_X and n_0 . The most relevant result is that $dr^*/dn_0 > 0$ for 75.2% of instances compared to 96.2% of Case IV trials in which $dQ^*/dn_0 > 0$.

Appendix E: Computation of CPCS and power

In section 4, we defined the CPCS as the conditional probability of correct selection of a technology for adoption, given $W = w$:

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

Table EC.6 Numerical comparative statics for Case III trials

Parameter b	Number of instances	$dV^*/db > 0$	$dr^*/db > 0$
P	5,624	100%	95.7%
n_0	5,624	0.1%	75.2%
σ_X	5,624	99.6%	99.3%
c	5,624	0.0%	0.0%
c_r	5,624	0.0%	0.0%

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:

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

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

$$\text{CPCS}(w) = \Pr(Z_{Tr} > \alpha_N(T) | W = w).$$

We can then compute $\text{CPCS}(w) = 1 - \Phi(U_N)$, where

$$U_N = \frac{2n_0(\alpha_N(T) - \mu_0) + Tr(\alpha_N(T) - w)}{\sigma_X \sqrt{2Tr}}.$$

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

$$U_S = \frac{2n_0(\alpha_S(T) + \mu_0) + Tr(\alpha_S(T) + w)}{\sigma_X \sqrt{2Tr}}.$$

Similarly, if $-\alpha_S(T) \leq w \leq \alpha_N(T)$, the oracle adoption decision is M, and we obtain $\text{CPCS}(w) = \Pr(-\alpha_S(T) \leq Z_{Tr} \leq \alpha_N(T) | W = w) = \Phi(U_N) + \Phi(U_S) - 1$.

The final closed-form expression of CPCS is then

$$\text{CPCS}(w) = \begin{cases} 1 - \Phi(U_N), & w > \alpha_N(T) \\ 1 - \Phi(U_S), & w < -\alpha_S(T) \\ \Phi(U_N) + \Phi(U_S) - 1, & -\alpha_S(T) \leq w \leq \alpha_N(T) \end{cases}$$

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 $\bar{X} = 2/(Tr) \sum_{i=1}^{Tr/2} X_i$ is larger than $q_{\alpha/2} \sigma_X / \sqrt{Tr/2}$, or smaller than $-q_{\alpha/2} \sigma_X / \sqrt{Tr/2}$. The power is the probability of rejection given $W = w$:

$$\begin{aligned} \text{power}(w) &= \Pr \left(\bar{X} > \frac{\sigma_X}{\sqrt{Tr/2}} q_{\alpha/2} \cup \bar{X} < -\frac{\sigma_X}{\sqrt{Tr/2}} q_{\alpha/2} \mid W = w \right) \\ &= \Pr \left(\bar{X} > \frac{\sigma_X q_{\alpha/2}}{\sqrt{Tr/2}} \mid W = w \right) + \Pr \left(\bar{X} < -\frac{\sigma_X q_{\alpha/2}}{\sqrt{Tr/2}} \mid W = w \right) \\ &= 1 - \Phi \left(q_{\alpha/2} - \sqrt{\frac{Tr}{2\sigma_X^2}} w \right) + 1 - \Phi \left(q_{\alpha/2} + \sqrt{\frac{Tr}{2\sigma_X^2}} w \right). \end{aligned}$$

Appendix F: Some special parameter settings that allow fully sequential trials

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. The sequential version of the one-stage optimal trial design problem in section 2.5, with fixed rate of recruitment r , is written in Appendix F.1 below. If we can show that this problem is equivalent (after suitable reparameterization) to an existing fully sequential trial model, then 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. The fixed recruitment rate can then be optimized. Appendix F.2 affirms that this can be done for four special parameter settings to match the model of Chick et al. (2017a). For these special settings, 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.

This section assumes a fixed patient pool ($P(T) = P$) and that the recruitment rate r is fixed. The results of this section are used in numerical experiments in section 5 to assess the benefit of optimizing recruitment rate in a one-shot trial, of shifting from a one-stage trial to a sequential trial that adapts trial duration, or of optimizing the recruitment rate of a sequential trial that adapts trial duration (if one of the four special parameter settings holds).

F.1. Optimize fixed recruitment, with optimal adaptive stopping time.

The model in section 2 assumes that T and r are continuous, whereas the decision epochs are discrete. We first recall a related discrete-time sequential model to introduce notation and fix ideas. We then adapt the continuous-time model of section 2 to a corresponding discrete-time model.

Chick et al. (2017a) provided a fully sequential two-armed trial with optimal adaptive stopping time, and which essentially assumes $p_N = 0$ and fixed patient pool in terms of our model. Alban et al. (2018) extended that model to allow for pragmatic trials with more general p_N . Following that model, we define decision epochs $t = 0, 1, \dots, Q_{\max} - 1$. At each decision epoch, measurements of health outcome and treatment cost are observed (with appropriate delay) for a pair of patients randomized to each arm of the trial. An action $a_t \in \{0, 1\}$ is made to continue sampling ($a_t = 1$) or stop sampling ($a_t = 0$) before enrolling each patient pair (hence the term fully sequential). Once recruitment stops, at stopping time $Q = \min\{t : a_t = 0\}$, we observe the outcomes of patients in the *pipeline* – those who have been randomized but whose outcomes have not been observed yet – before making the adoption decision. A policy, π , maps available knowledge (prior information, plus acquired data) to an action at each decision epoch. Let $\mathbb{E}_\pi[\cdot | K_0]$ denote the expectation induced by policy π given the prior information set $K_0 = (\mu_0, n_0)$. If $Q = 0$, the trial does not run, and the adoption decision is made immediately, based on the prior information alone.

We now adapt the model in section 2 to discrete-time. Let $\tau = \lceil r\Delta/2 \rceil$ be the number of patient pairs enrolled in the trial during the follow-up period of length Δ and let $\tilde{\rho} = e^{2\rho/r} - 1$ be the

discrete discount rate per patient pair. We write the expected net gain for a discrete time trial with recruitment rate r and policy π , which sequentially determines $Q \in \{0, 1, \dots, Q_{\max}\}$, as:

$$V(\pi, r) = -\mathbf{1}_{Q>0}c_{\text{cap}}(r) + \mathbb{E}_{\pi} \left[\sum_{t=0}^{Q-1} \frac{\delta_{\text{on}}(1 - 2p_{\text{N}})X_{t+1} - 2c}{(1 + \tilde{\rho})^t} \middle| K_0 \right] + \mathbb{E}_{\pi} \left[\frac{\mathbf{1}_{\mathcal{D}=\text{N}}((1 - p_{\text{N}})P_{\rho}(2Q/r)W - I_{\text{N}}) + \mathbf{1}_{\mathcal{D}=\text{S}}(-p_{\text{N}}P_{\rho}(2Q/r)W - I_{\text{S}})}{(1 + \tilde{\rho})^{\mathbf{1}_{Q>0}(Q+\tau)}} \middle| K_0 \right]. \quad (\text{EC.9})$$

A policy π_r^* that maximizes $V(\pi, r)$ over non-anticipative policies π is said to optimize the trial duration for a given fixed r . An optimum can be found by solving the univariate optimization $\max_r V(\pi_r^*, r)$, assuming the computability of π_r^* and $V(\pi_r^*, r)$, for $r \in [0, r_{\max}]$.

F.2. Equivalence for four special cases

If we assume a fixed patient pool ($P(T) = P$), we may identify four special parameter settings that permit the computation of the optimal sequential policy π_r^* and expected net gain $V(\pi_r^*, r)$ in (EC.9) for any given r using the methods of [Chick et al. \(2017a,b\)](#):

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

$$V(\pi) = \mathbb{E}_{\pi} \left[\left\{ \sum_{t=0}^{Q-1} \frac{-c + \delta_{\text{on}}X_{t+1}}{(1 + \tilde{\rho})^t} \right\} + \frac{\mathbf{1}_{\mathcal{D}=\text{N}}(P_{\rho}W - I)}{(1 + \tilde{\rho})^{\mathbf{1}_{Q>0}(Q+\tau)}} \right]. \quad (\text{EC.10})$$

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 any observations seen by time t , for $t = 0, 1, 2, \dots, Q$. After stopping and observing all the data, the new alternative is selected as best if $P_{\rho}W - I \geq 0$.

We first explore the differences between (EC.9) and (EC.10) related to the fixed costs of the trial. The fixed cost $c_{\text{cap}}(r)$ in (EC.9) 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 patient pair is to be observed, for some parameter settings.

F.2.1. Current practice is standard treatment. When $p_{\text{N}} = 0$, (EC.9) becomes

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

which is equivalent to (EC.10) by letting $c' = 2c$ and $I' = I_{\text{N}}$.

F.2.2. Forcing the decision to adopt one of the two HTs, undiscounted rewards. When the adoption decision is forced to be N or S with undiscounted rewards¹⁵ ($\tilde{\rho} = 0$), we can use $\mathbf{1}_{\mathcal{D}=\text{S}} = 1 - \mathbf{1}_{\mathcal{D}=\text{N}}$ such that (EC.9) becomes

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

If $p_{\text{N}} < 1/2$, then we can divide through by $1 - 2p_{\text{N}}$ and add the constant $Pp_{\text{N}}\mu_0 - I_{\text{S}}$ without changing the optimal π . Thus, the optimal π maximizes:

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

which is in the form of (EC.10) when we let $c' = 2c/(1 - 2p_{\text{N}})$, $P' = P/(1 - 2p_{\text{N}})$, and $I' = (I_{\text{N}} - I_{\text{S}})/(1 - 2p_{\text{N}})$.

If $p_{\text{N}} = 1/2$, the optimal π maximizes

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

which is in the form of (EC.10) when $\delta'_{\text{on}} = 0$, $c' = 2c$, $P' = P$, and $I' = I_{\text{N}} - I_{\text{S}}$.

F.2.3. No switching costs and undiscounted rewards. When $I_{\text{N}} = I_{\text{S}} = 0$, then it is optimal to adopt N or S. When rewards are not discounted ($\tilde{\rho} = 0$) the transformation in Appendix F.2.2 is also valid in this setting.

F.2.4. No switching costs and $\mu_0 = 0$. When $I_{\text{N}} = I_{\text{S}} = \mu_0 = 0$, 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 $\mathbf{1}_{\mathcal{D}=\text{S}} = 1 - \mathbf{1}_{\mathcal{D}=\text{N}}$ in (EC.9) to obtain

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

If $p_{\text{N}} < 1/2$, then we can divide through by $1 - 2p_{\text{N}}$ and obtain the optimal π by solving (EC.10) with $c' = 2c/(1 - 2p_{\text{N}})$ and $P'_{\tilde{\rho}} = P_{\tilde{\rho}}/(1 - 2p_{\text{N}})$. If $p_{\text{N}} = 1/2$ we obtain the optimal π by solving (EC.10) with $\delta'_{\text{on}} = 0$, $c' = 2c$, and $P'_{\tilde{\rho}} = P_{\tilde{\rho}}$.

Appendix G: Supplementary figures for the application to the ProFHER trial

This section provides supplementary numerical analysis for the application to the ProFHER trial in section 4. Appendix G.1 gives sensitivity analysis for the expected net gain of a trial and optimal trial as a function of several key parameters that describe the value of a trial. Appendix G.2 illustrates numerically the quality of the asymptotic results for the various scenarios. Appendix G.1.4 explores online versus offline learning.

¹⁵ Alban et al. (2018) incorrectly claim that this result is also valid with discounted rewards.

G.1. Additional sensitivity analysis

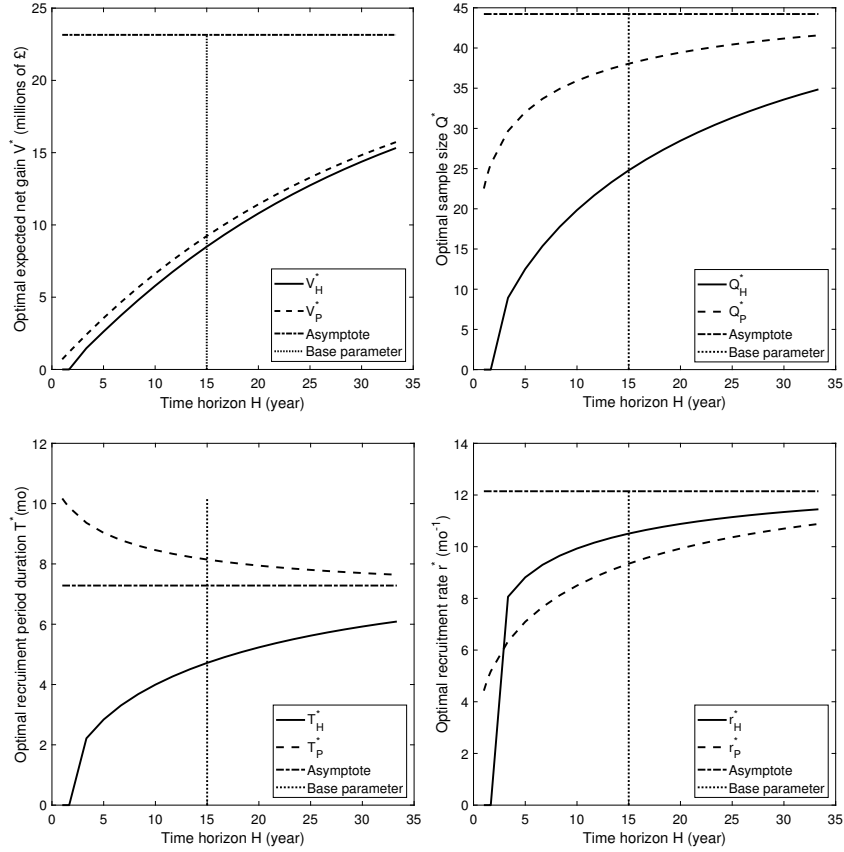
Appendix G.1.1 gives sensitivity analysis to assumptions to the size of the post-adoption population. Appendix G.1.2 does so for the discount rate. Appendix G.1.3 gives sensitivity analysis for the fixed costs of changing technology, and the fraction of patients on the new health technology.

G.1.1. Sensitivity to the post-adoption population. For ease of exposition, we fix $P = \zeta H$ and analyze the effect of a change in H for both fixed horizon and fixed patient pool, while keeping the remaining base case parameters constant. Figure EC.1 presents the results for the optimal expected net gain, sample size, recruitment duration, and recruitment rate, along with the asymptotes derived in section 3.6.2. The optimal sample size is increasing in H . The more interesting result is how the additional patients are obtained in terms of recruitment duration and rate. The optimal recruitment rate is increasing under 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. In numerical experiments for fixed patient pool, a longer trial penalizes the expected net gain more for larger P than for smaller P due to greater discounting of a larger value. In contrast, for fixed horizon, the penalty of discounting is moderated by a decreasing value of the post-adoption population for longer trials. The optimal expected net gain V^* is increasing in H as shown in section 3.5. Due to the high incidence rate in this application, it is optimal to run the trial even for a short time horizon. However, a time horizon shorter than 20 months does not justify the trial under a fixed horizon.

While the asymptotic approximations are accurate for Cases I–III, particularly in undiscounted scenarios (see Appendix G.2 below), the asymptotes are not good approximations in the range of interest of our application. However, the optimal expected net gain and trial design do approach the asymptotes for larger time horizons.

G.1.2. Sensitivity to the discount rate. Figure EC.2 shows the optimal expected net gain, sample size, recruitment duration, and recruitment rate as a function of the discount rate, keeping all other parameters constant. We note two important observations. First, the optimal expected net gain and design for a fixed patient pool are much more sensitive to the discount rate than for the fixed horizon: the recruitment duration increases steeply as the discount rate goes to zero. In section 3.4 we showed that in a scenario with fixed patient pool and no discounting, it is optimal to set the recruitment duration to T_{\max} , where T_{\max} can be arbitrarily large. If $T_{\max} = 10$ years, the optimal sample size under the undiscounted fixed patient pool scenario is 170, an example of a situation where the value-based design recruits more patients than the classical design.

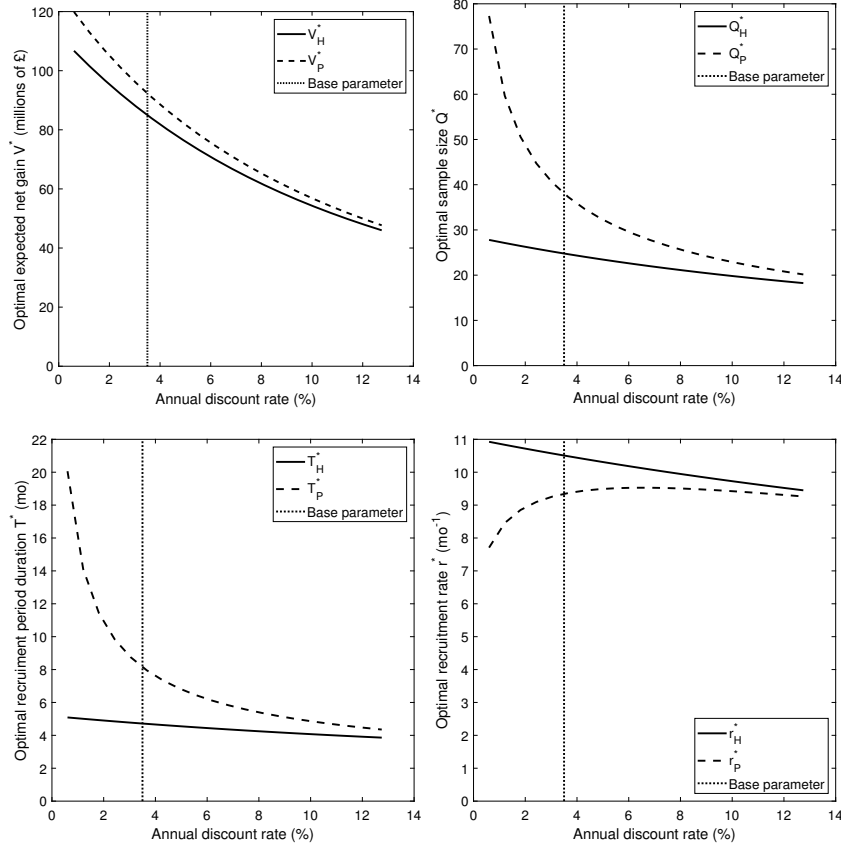
Figure EC.1 Optimal expected net gain, sample size, recruitment duration, and recruitment rate as a function of the time horizon.



Second, with a fixed horizon, both recruitment duration and rate are decreasing with the discount rate. With a fixed patient pool, we observe a decrease in recruitment duration but an inverse U-shaped optimal recruitment rate, increasing for an annual discount rate below 7%. The U-shaped behavior illustrates two opposing effects: (i) a larger discount rate prompts quicker action through a larger recruitment rate; (ii) a larger discount rate reduces the potential gains by reducing the effective post-adoption population, prompting reduced spending and lower recruitment rates.

G.1.3. Sensitivity to the fraction of patients on the new HT, p_N and fixed costs of changing technology. Our base case parameters include $I_N = I_S = \mu_0 = 0$, so there are no fixed costs associated with changing health technology, which imply that p_N has no effect on the expected net gain or optimal trial length (see section 3.5). Here, we consider scenarios where $I_N = I_S = \text{£}10\text{Mill}$ and $I_N = I_S = \text{£}100\text{Mill}$ and a range of values for p_N while keeping the remaining base case parameters fixed. Figure EC.3 shows the optimal expected net gain and sample size as a function of p_N . For $I_N = I_S = \text{£}10\text{Mill}$, p_N has a negligible effect on Q^* and a small effect on the expected net gain. For $I_N = I_S = \text{£}100\text{Mill}$, there is a substantial effect on expected net gain and a noticeable difference in the sample size for the fixed patient pool scenario.

Figure EC.2 Optimal expected net gain, sample size, recruitment duration, and recruitment rate as a function of the discount rate.



We observe that I_N , I_S , and p_N are only relevant to the optimal design when the switching costs are large. By inspecting (8), we see that the switching costs only appear in $\alpha_N(T)$ and $\alpha_S(T)$, which are further divided by σ_Z . If $\alpha_N(T)$ or $\alpha_S(T)$ are similar in magnitude to σ_Z , then the design will have a significant dependency on I_N , I_S , and p_N . As an example, for a fixed horizon with $p_N = 0.39$, the ratio $\alpha_N(T_H^*)/\sigma_Z$ is 0.072 for $I_N = \mathcal{L}10\text{Mill}$ and 0.723 for $I_N = \mathcal{L}100\text{Mill}$, respectively, explaining why we observe a noticeable difference only for the latter scenario.

G.1.4. Online vs. offline learning. Consider the base case parameters introduced in section 4. Because online and offline learning are equivalent when $\mu_0 = 0$, we evaluate the difference between online and offline learning for a range of values taken by μ_0 . We do not see any difference in Figure EC.4 as the lines overlap. We do not observe much larger differences by reducing the post-adoption population (P for fixed patient pool and H for fixed horizon).

G.2. Asymptotic analysis

This section provides numerical results for variations on the ProFHER trial, in order to give insights as to the asymptotic behavior for the four cases in the taxonomy of value-based designs in Table 1 of the main paper.

Figure EC.3 Optimal expected net gain (left) and sample size (right) as a function of p_N .

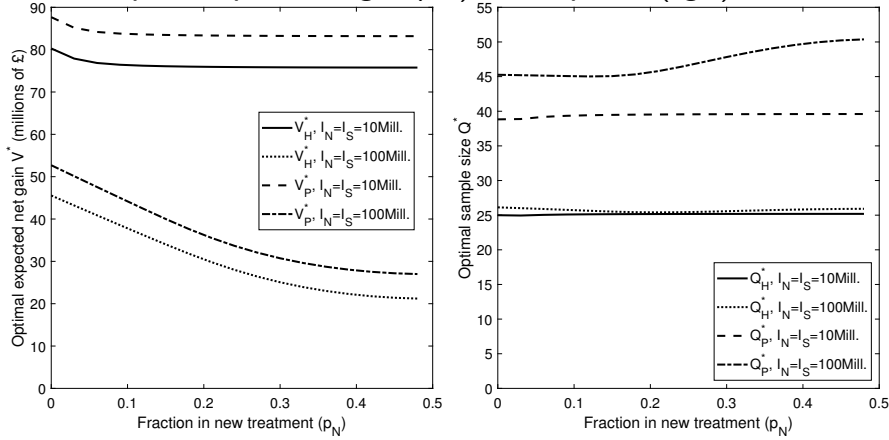



Figure EC.4 Optimal expected net gain and sample size as a function of μ_0 and the type of learning.



G.2.1. Asymptotics ($P(T) \rightarrow \infty$) for Cases I and II. Consider the base case parameters introduced in section 4. To force Cases I and II, we let $c_{\text{cap}}(r) = \mathcal{L}960,000$ independent of the recruitment rate, and let $r_{\text{max}} = \tilde{r} = 7.8/\text{month}$. Figure EC.5 shows how the optimal sample size increases as a function of H (left subfigure, for the fixed horizon case) and P (right subfigure, for the fixed patient pool case) when ρ is as in the base case parameters and $\rho = 0$. Similarly, Figure EC.6 plots the optimal expected net gain. Plotted with dotted and dash-dotted lines are the asymptotic approximations that were derived in section 3.6. Horizontal dotted lines correspond to the approximations with discounted rewards, increasing dash-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 (the lines are overlapping in the figures). This is due to $I_N = I_S = \mu_0 = 0$, and $P(T^*)\sigma_0 \gg c\tilde{r}T^*$. The approximations for discounted rewards are accurate for the optimal sample size under a fixed patient pool, but, for the range of values under consideration, the discounted approximations are less accurate.

Figure EC.5 The optimal sample size as a function of H for fixed horizon (left) and P for fixed patient pool (right) for Cases I and II.





Figure EC.6 The optimal expected net gain as a function of H for fixed horizon (left) and P for fixed patient pool (right) for Cases I and II.



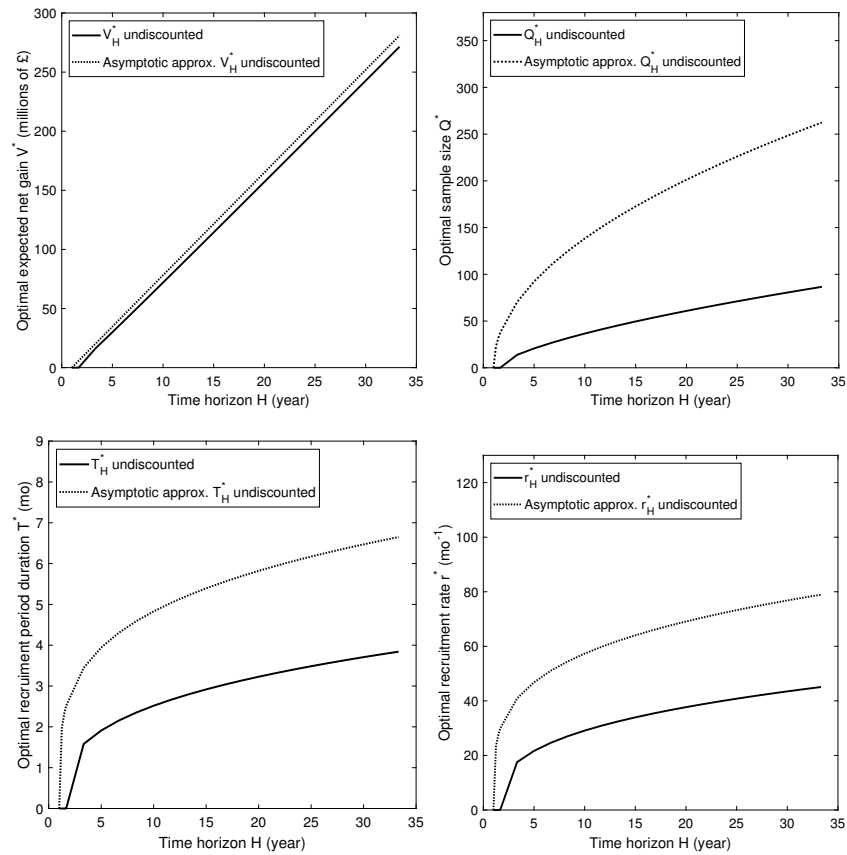
G.2.2. Asymptotics $P(T) \rightarrow \infty$ for Case III. We can force Case III by letting $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r$, $\rho = 0$ and a fixed patient pool. The results are similar to Cases I and II (data not shown).

G.2.3. Asymptotics ($P(T) \rightarrow \infty$) for Case IV. In Appendix G.1.1, we showed the asymptotic approximations for Case IV with $\rho > 0$. Here, we explore the asymptotic approximations with undiscounted rewards ($\rho = 0$). We assume a setup cost function of the form: $c_{\text{cap}}(r) = c_{\text{fix}} + c_r r$. We assume that half of the setup costs of the trial were fixed (as in section 4) and the other half represent marginal costs, obtaining the following estimates: $c_{\text{fix}} = c_{\text{cap}}(\tilde{r})/2 = \text{£}480,000$; $c_r = c_{\text{cap}}(\tilde{r})/2/\tilde{r} = \text{£}5,080$ per year. Figure EC.7 shows the results for the optimal expected net gain and design variables. We find that the approximations for the optimal design variables are poor but for the expected net gain are reasonably close.

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., *Proc. 2018 Winter Simulation Conference*, 2459–2470 (Piscataway, NJ: IEEE, Inc.).

Figure EC.7 Optimal expected net gain, sample size, recruitment duration, and recruitment rate as a function of the time horizon.



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

Chick SE, Forster M, Pertile P (2017a) 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, Forster M, Pertile P (2017b) htadelay package. <https://github.com/sechick/htadelay>.