Proceedings of the 2025 Winter Simulation Conference
E. Azar, A. Djanatliev, A. Harper, C. Kogler, V. Ramamohan, A. Anagnostou, and S. J. E. Taylor, eds.

EXPLORE THEN CONFIRM: INVESTMENT PORTFOLIOS FOR NEW DRUG THERAPIES

Zaile Li[1], Stephen E. Chick[1], Sam Daems[2,3], Shane G. Henderson[4]

[1]Technology and Operations Management Area, INSEAD, Fontainebleau, FRANCE
[2] Université Libre de Bruxelles, Brussels, BELGIUM
[3] Waterland Private Equity, Antwerp, BELGIUM
[4]School of Operations Research and Information Engineering, Cornell University, Ithaca, NY, USA

ABSTRACT

New medical technologies must pass several risky hurdles, such as multiple phases of clinical trials, before market access and reimbursement. A *portfolio* of technologies pools these risks, reducing the collective financial risk of such development while also improving the chances of identifying a successful technology. We propose a stylized model of a portfolio of technologies, each of which must pass one or two phases of clinical trials before market access is possible. Using ideas from Bayesian sequential optimization, we study the value of running response-adaptive clinical trials to flexibly allocate resources across clinical trials for technologies in a portfolio. We suggest heuristics for the response-adaptive policy and find evidence for their value relative to non-adaptive policies.

1 PORTFOLIOS OF DRUG INNOVATIONS AND SIMULATION OPTIMIZATION

We present preliminary work to assess whether innovative response-adaptive clinical trials, which have largely been applied to individual trials one at a time, can also be of medical and financial benefit for operating portfolios of clinical trials. Our analysis draws upon techniques from Bayesian ranking and selection in the simulation optimization community, among other operations research techniques. We see this work as one step in a larger set of applied research challenges to help lower the financial risk of taking validated health innovations from concept to market. The work is motivated by recent proposals to assess how portfolios of new medical technologies can incentivize the development of new therapies for pediatric oncology relative to the development of each technology on its own (Das et al. 2018; Daems et al. 2023).

New medical technologies usually proceed through a sequence of steps to assess their effectiveness and cost effectiveness in treating a given medical condition. For example, a new drug will typically undergo preclinical, clinical (e.g., Phases I, II and III clinical trials), and market access approval phases, each run only if the preceding phase is successful. Each of these steps may be costly, and the probability that a given technology will pass through all phases and have a market impact may be low. There has been a lot of work in the biostatistics and operations research communities to improve the cost and/or effectiveness of a given clinical trial (Berry 2011; Hampson and Jennison 2013; Williamson et al. 2017; Pallmann et al. 2018; Angus et al. 2019; Chick et al. 2022).

Given the low probability of approval for an individual technology, researchers have also explored whether a *portfolio* of technologies could help advance technologies that, on their own, are too risky to pursue (Fernandez, Stein, and Lo 2012). From the perspective of an investor, for example, it may be valuable to acquire a 5% stake in a portfolio of 20 potential health innovations, each with a 5% probability of becoming a blockbuster drug, rather than taking a 100% stake in just one such potential health innovation. Indeed, if a single blockbuster is sufficient for profitability and the health innovation approvals are independent of each other, then the probability of having a positive revenue stream with the portfolio investment is $1-0.95^{20}=0.64$, as opposed to the 5% probability when investing in a single health innovation. There are also medical and ethical arguments for such portfolios (London and Kimmelman 2019).

Andrew Lo and coauthors (e.g., Fernandez et al. 2012; Fagnan et al. 2014) develop and analyze the idea of a drug megafund portfolio. They study the challenges, and potential value, of investment in a portfolio of drugs that face a sequence of Phase I, II and III trials and a market access step. Daems et al. (2023) builds on this line of work and proposes an investment vehicle to help incentivize investments in pediatric oncology – an area with a high potential impact for children diagnosed with many cancers, but with potentially smaller payoffs for each successful treatment because childhood cancer is, thankfully, rare.

There are at least three problems that a manager in a private equity firm faces when implementing a portfolio of new biotech projects. 1) The portfolio construction problem involves selecting projects to include in the portfolio and raising a suitable amount of money. 2) The portfolio management problem entails the allocation of funds, over time, to run various phases of various projects. 3) The portfolio liquidation problem entails connecting approved projects to revenue streams following market access. We study Problem 2, but our results also have some bearing on Problem 1.

A central assumption in much of the drug portfolio work to date is that the resources (primarily money and time) for a given phase of a project are committed up front once and for all. However, there may be value in dynamically allocating resources *within* a phase across the projects in development. For example, if signals accumulating during a clinical trial for a drug suggest that the phase will not be successful, one might consider stopping that phase of the trial for that drug in order to dynamically reallocate the freed resources to other trials that portend more success.

We present a stylized model to estimate the value of performing *adaptive* clinical trials across a portfolio of drugs under development. Our work exposes new issues, and modeling and analysis questions, associated with portfolios of financial investments in clinical trials, relative to the extant literature on the design of clinical trials using Bayesian optimization in the simulation and broader operations research communities. We do not address the practical challenges of dynamic allocation of resources, which may be considerable.

We draw upon the field of Bayesian multi-arm ranking and selection techniques (Branke et al. 2007; Frazier et al. 2008; Eckman and Henderson 2022), to decide how to allocate patients to one or more treatment arms (an allocation rule), or to determine when the cost of continuing a clinical trial outweighs the expected value of information from continuing the trial (a stopping rule), before picking one treatment alternative as best for a single post-trial population of patients of interest (a selection rule).

We are interested in a question beyond selecting a single best alternative: how should one select a subset of a portfolio of health technologies that are cost effective for the potentially different populations that each is intended to treat? Fairley et al. (2020) explored a variation on this question in health economics, essentially extending to portfolios the one-shot expected value of information (EVI) approach to valuing individual new health technologies (Strong et al. 2015; Fenwick et al. 2020). They used EVI techniques to identify a subset of projects for further analysis and an optimal one-shot sampling budget for the projects that are to be analyzed, with a goal of maximizing expected gain from those projects identified as worth deploying. Variations of this problem include work on sequential subset selection or on thresholding bandits, in which a sequence of observations are made, whereupon a subset of alternatives whose unknown means are believed to be above a threshold is selected, e.g., Xie and Frazier (2013) and Locatelli et al. (2016).

Our context has some differences with these papers. Drugs face a sequence of phases of analysis. Success in a Phase I trial does not imply overall project approval: follow up trials are required, each of which may succeed or fail. Estimates of success probabilities for different phases are available (e.g., BIO et al. 2021). We do not assume success probabilities for an entire phase; rather we assume that the probability of success in a phase is a consequence of a sequence of the outcomes of patients in that phase.

To the best of our knowledge, such a multi-phase analysis has not been considered in the simulation optimization community. We study a simplified setting to focus on the most salient issues. We assume that there are only two phases: an *exploratory* phase and a *confirmatory* phase. The exploratory phase can model, for example, a Phase I or Phase I/II trial, which assesses whether the therapy is nontoxic and has an effect on the population. Success in this phase means it is eligible for further assessment in a confirmatory phase, for example a Phase II or Phase II/III trial. This two-phase setup is relevant for rare diseases (EMA

2025). In a second example, this framework can model a single set of Phase III trials, where the exploratory phase is a pilot study that is used to identify a subset of trials for which a full Phase III trial is to be run.

We assume market access, and hence revenues accrue, for successfully confirmed projects. These simplifications, and others we adopt below, are abstractions, but they enable us to obtain insights on the benefits of adaptively allocating financial resources across projects as the trials proceed, as compared to committing fully to completing each phase of a trial in advance of the phase, and as compared to running the trials in a portfolio as independent projects. They also enable us to propose some heuristics for exploratory phase operations and study their potential value.

Methodologically, we build on one-step look-ahead techniques to sequentially allocate resources to one of several projects in the portfolio during the exploratory phase, using greedy EVI/knowledge gradient (KG) ideas to compute allocation indices (Chick and Inoue 2001; Frazier et al. 2008; Powell and Ryzhov 2012; Chick et al. 2022). As with some related work (Chick and Frazier 2012), we balance the cost of sampling with the EVI of exploratory samples, rather than assuming a fixed budget or fixed precision approach that is sometimes assumed (Hong et al. 2021). The computation of the indices extends work of Xie and Frazier (2013), who studied sequential subset selection without a confirmatory phase, by accounting for the expected value of continuing to a confirmatory phase. That computation involves computing the expected value of doing a confirmatory trial with the remaining budget allocation, building on related work of Fairley et al. (2020) for single-phase learning, and leveraging a greedy algorithm (Martello and Toth 1990, §2.4) to approximately solve a certain knapsack problem.

The primary contributions of this paper are the simplified framework for the portfolio of projects, the development of heuristics for exploratory policies, and experiments suggesting the potential value of adaptive policies and new insights on when they are beneficial relative to non-adaptive policies.

Section 2 gives the model and articulates important assumptions. Section 3 gives some analytical results, and motivates certain heuristics. Section 4 gives illustrative numerical results. Section 5 discusses implications and limitations of our proposal, and outlines potential areas for future work.

2 BASE MODEL FOR EXPLORATORY AND CONFIRMATORY PHASES

We assume that a finite total budget B_0 is used to assess a set $[K] = \{1, 2, \cdots, k\}$ of new health technologies (projects). Each project $x \in [K]$ must successfully pass through two phases, exploratory and confirmatory, to be approved. Projects that are approved following those phases can treat patients and thereby generate revenue for investors in the portfolio. The objective is to manage the two phases in a way that maximizes the expected net value (revenue and health benefits less costs) of the portfolio.

2.1 Assumed Structure for Learning in Exploratory and Confirmatory Phases

Health Economic Outcomes. We assume that health outcomes can be mapped to financial outcomes, so that the EVI of sampling plans to infer health outcomes can be compared with the costs of trials. We therefore assume health outcomes are continuous valued and represent quality adjusted life years (QALYs) that can be converted to money by a parameter λ, e.g., $\lambda = £20,000/\text{QALY}$. Patient outcomes are measured in incremental net monetary benefit (INMB), consisting of QALY outcomes converted to money minus the cost of treatment and cost of complications. We use $X_{x,\ell}$ during the exploratory phase, and $Y_{x,\ell}$ during the confirmatory phase, for the INMB of the ℓth patient for project x in the given phase.

Assume the INMB of all patients are independent Gaussian random variables, i.e., that $(X_{x,l}, Y_{x,m} : x \in [K], l, m = 1, 2, \ldots)$ are independent Gaussian random variables. Importantly, this implies that the treatments and patient subpopulations for different projects differ. Assume the unknown mean INMB $\theta_x = \mathbb{E}[X_x]$ has a Normal$(\mu_{0,x}, \sigma^2_{0,x})$ prior distribution independently across projects, and that the observation variance $\sigma^2_x = \text{Var}[X_x]$ is known. We also assume $\mathbb{E}[X_x] = \mathbb{E}[Y_x]$ and $\text{Var}[X_x] = \text{Var}[Y_x]$. This is a strong assumption for separate Phase I/II and Phase III trials, but is a natural assumption if the portfolio's exploratory phase is a pilot study to inform the selection of projects that merit a full Phase III trial.

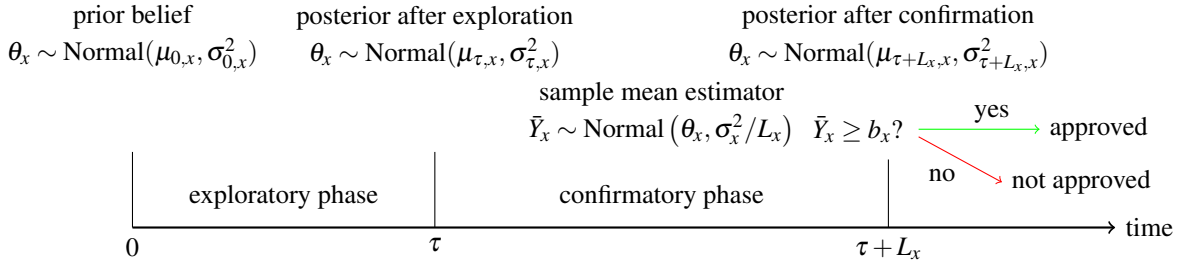

Figure 1: Timeline and Information Flow for Project x

Exploratory phase. We assume that patients can be enrolled and treated in the exploratory phase of any project selected by the portfolio manager. At each time step $t = 0, 1, \ldots, \tau - 1$ in the exploratory phase, exactly one project x_t is activated, i.e., a patient is assigned to the exploratory phase for project $x_t \in [K]$. We assume that the outcome of the treatment can be observed quickly, before the next time step, so that the ensuing assignments of patients to projects is fully response adaptive. We often assume that τ is a prespecified fixed value, but, in general, allow τ to be a stopping time.

As the exploratory phase progresses, posterior distributions for the unknown means are updated. If a new patient is enrolled in project $x = x_t$ at time t, then we update $(\mu_{t,x}, \sigma^2_{t,x})$ to $(\mu_{t+1,x}, \sigma^2_{t+1,x})$ using Bayes rule, and we increment the sample size $n_{t+1,x} = n_{t,x} + 1$. For projects $x' \neq x_t$ not enrolling a new patient at time t, we set $(\mu_{t+1,x'}, \sigma^2_{t+1,x'}) = (\mu_{t,x}, \sigma^2_{t,x})$ and $n_{t+1,x'} = n_{t,x'}$. (This assumes that projects have independent rewards.) We write vectors as $\vec{\mu}_t = (\mu_{t,1}, \mu_{t,2}, \ldots, \mu_{t,k})$ and $\vec{\sigma}_t = (\sigma^2_{t,1}, \ldots, \sigma^2_{t,k})$. At the end of the exploratory phase, the posterior means and sample sizes are $\vec{\mu}_\tau$ and $\vec{n}_\tau$, respectively.

Confirmatory phase. After τ patients have been treated across all projects, the exploratory phase stops and a subset $[K]_\tau$ of projects are selected for a confirmatory phase. The choice of $[K]_\tau$ is made based on the posterior mean rewards of each project through time τ, the portfolio manager's objective, and the residual budget. The confirmatory phase consists of running confirmatory trials for all the projects in $[K]_\tau$.

Although the exploratory phase uses Bayesian inference, in line with regulatory agency practice, we use a frequentist framework for the design of the confirmatory trials that start at time τ. A standard calculation gives the sample size, $L_x = \sigma_x^2 (z_\alpha + z_\beta)^2 / \delta_x^{-2}$, required to achieve specified type 1 and type 2 error levels, α and β, respectively, to detect a minimal acceptable improvement of δ_x. Here z_α and z_β are the appropriate quantiles of a standard normal distribution. Improvement is assessed relative to the expected INMB η_x of the current known best treatment alternative for the patients that are targeted by project x. The confirmatory trial uses a sample mean estimator $\bar{Y}_x \mid \theta_x \sim \text{Normal}(\theta_x, \sigma_x^2/L_x)$ with sample size L_x. The confirmatory trial is said to be positive if the sample mean $\bar{Y}_x$ is at least as big as a threshold b_x that is determined by the confidence and power requirements associated with the minimum sample size, where

$$b_x = \eta_x + z_\alpha \sigma_x / \sqrt{L_x}. \tag{1}$$

Summary of Phases. The exploratory phase samples projects sequentially, one per time step, then the confirmatory phases run in parallel starting at time τ. Figure 1 gives a summary.

The manager's decision variables for a given portfolio are a sequence of projects from which to learn, x_t, for times $t = 0, 1, \ldots, \tau - 1$, a (potentially response adaptive) stopping time τ representing the total number of patients in exploratory trials, and a subset $[K]_\tau$ of projects to take to a confirmatory phase. Together, these decisions constitute a policy π, and we require π to be non-anticipative.

2.2 The Portfolio's Costs and Potential Benefits

We assume that the cost of a single observation for project x is c_x during the exploratory phase, so that $B_{t+1} = B_t - c_{x_t}$. We assume that the fixed costs of the exploratory trial are sunk at time $t = 0$, and that there is a fixed cost C_x, to include $x \in [K]_\tau$ in the confirmatory phase. The fixed cost C_x includes the

per-observation cost required for the sample size L_x in the confirmatory phase. At the time of completion, the overall posterior mean for $x \in [K]_\tau$ from both exploratory and confirmatory phases is denoted $\mu_{\tau+L_x,x}$, the posterior variance of the unknown mean is $\sigma^2_{\tau+L_x,x}$, and the overall number of observations is $n_{\tau,x}+L_x$.

A project x is said to be **approved** if it is both selected for the confirmatory trial, i.e, $x \in [K]_\tau$, *and*, after the completion of a confirmatory trial, its sample mean is at least as large as the target threshold b_x. It is said to be **successful** if it is both approved and its mean INMB exceeds that of its comparator, $\theta_x > \eta_x$.

An approved project x generates revenue $P_x(p_x - m_x) + P_x\nu(\mu_{\tau+L_x,x} - \eta_x)$, where P_x is the number of patients, p_x is the treatment price, and m_x is the manufacturer's cost per treatment associated with project x. The second term rewards better health outcomes with higher prices. It depends on the *bargaining power* $\nu \in [0,1]$ of the portfolio manager in obtaining some of the surplus of the new treatment, measured in INMB, given data from both trial phases. *Value-based prices* arise when $\nu > 0$, while *flat prices* arise when $\nu = 0$. For a discussion of value-based pricing in a related context, see Yapar et al. (2025).

Portfolio value. We take the perspective of a for-profit portfolio manager who invests in risky projects developed by for-profit providers of health technologies. Let $\mathscr{F}_t$ be a sigma field giving the information available to the portfolio manager at time t, prior to decisions made at time t. Rewards are not discounted.

The financial value of the portfolio during exploration, at time $t = 0, 1, \ldots, \tau$, with policy π, is

$$V_{\pi,t}(\vec{\mu}_t, \vec{\sigma}_t, B_t) = B_t + \mathbb{E}\left[-\sum_{t'=t}^{\tau-1} c_{X^\pi_{t'}} + \sum_{x \in [K]^\pi_\tau} (r_\tau(x, \mu_{\tau+L_x,x}, \eta_x, b_x) - C_x) \,\middle|\, \mathscr{F}_t \right], \tag{2}$$

on the event $t \le \tau$, where $r_\tau(x, \mu_{\tau+L_x,x}, \eta, b) = [P_x(p_x - m_x) + P_x\nu(\mu_{\tau+L_x,x} - \eta)] \cdot \mathbb{1}(\bar{Y}_x \ge b)$ is the profit from confirming arm x when its posterior mean is $\mu_{\tau+L_x,x}$, η_x is the mean INMB of the current standard, and b_x is the threshold for approval for the observed treatment effect. At time τ, when τ patients have been treated and their outcomes have been observed in the freshly completed exploratory phase, $r_\tau(\cdot)$ is a random variable, and its realization is observed at time $\tau + L_x$. The value b_x will typically be strictly larger than η_x; see (1). At time $t = \tau$, the first sum in (2) is taken to be 0.

Regulatory and health metrics are also of interest. The expected number of approved projects is $\mathbb{E}A = \mathbb{E}\sum_{x \in [K]_\tau} \mathbb{1}(\bar{Y}_x \ge b_x)$ and the expected number of successful projects is $\mathbb{E}S = \mathbb{E}\sum_{x \in [K]_\tau} \mathbb{1}(\bar{Y}_x \ge b_x, \theta_x > \eta_x)$. The probability that at least one trial is approved (successful) is $\Pr(A \ge 1)$, $(\Pr(S \ge 1))$, respectively.

We further assume that the portfolio's budget B_t is not allowed to go negative (no borrowing) at any time. At time $\tau + \max_{x \in [K]^\pi_\tau} L_x$, all value, including the residual budget, is distributed to investors.

3 STRUCTURAL ANALYSIS AND SOME HEURISTIC SOLUTIONS

How should the portfolio manager maximize the objective function, $V_{\pi,0}(\vec{\theta}_0, \vec{\sigma}_0, B_0)$? We first discuss the selection of $[K]_\tau$ at time τ, when switching from exploration to confirmation. We then discuss the selection of the next project to activate, x_t, while the exploratory phase continues. This yields a Bellman equation for a dynamic program which we cannot solve analytically. Thus we explore heuristics for project selection during the exploratory phase and for the choice of τ. We also propose some comparators, such as the value of no information collection, and the value of having full information.

3.1 Picking Projects for a Confirmatory Phase

At the time of switching to the confirmatory phase, τ, the available budget is $B_\tau = B_0 - \sum_{t=0}^{\tau-1} c_{X^\pi_t}$. The portfolio manager uses this budget to select a subset of projects $[K]_\tau \subseteq \{1, \ldots, K\}$ for confirmation, to obtain a reward that corresponds to the net profit in the second summation in the expectation of (2).

The fixed sample size assumption means that the cost, C_x, of the confirmatory trial for project x is fixed. Thus, we can select $[K]_\tau$ at time τ by solving the deterministic (conditional on $\mathscr{F}_\tau$) knapsack problem $\max B_\tau + \sum_{x=1}^k (R_x - C_x) y_x$, subject to $\sum_{x=1}^k C_x y_x \le B_\tau$ with binary variables y_x, $x = 1, 2, \ldots, k$, where $R_x = \mathbb{E}[r_\tau(x, \mu_{\tau+L_x,x}, \eta_x, b_x) \mid \mathscr{F}_\tau]$. An explicit formula for R_x can be derived as follows; see Li et al.

(2025) for the details. Let $\hat{\sigma}_{\tau,x} = \sqrt{\sigma^2_{\tau,x} + \sigma^2_x / L_x}$, $w_x = \sigma^2_{\tau,x} / \hat{\sigma}^2_{\tau,x}$, and $z_{\tau,x} = (\mu_{\tau,x} - b_x)/\hat{\sigma}_{\tau,x}$. Then,

$$\mathbb{E}[r_\tau(x, \mu_{\tau+L_x,x}, \eta_x, b_x) \mid \mathscr{F}_\tau] = [P_x(p_x - m_x) + P_x \nu(\mu_{\tau,x} - \eta_x)]\Phi(z_{\tau,x}) + P_x \nu w_x \phi(z_{\tau,x})\hat{\sigma}_{\tau,x},$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ denote the pdf and cdf of the standard normal distribution, respectively.

Any slack in the financing constraint, $B_\tau - \sum_{x \in [K]_\tau} C_x \geq 0$, refers to funds available to the portfolio manager that are not used for trials, but form part of the value of the portfolio.

3.2 Bellman Equation for Portfolio Value

For $t < \tau$, i.e., during the exploratory phase, the optimal value function satisfies the Bellman equation

$$V(\vec{\mu}_t, \vec{\sigma}_t, B_t) = \max\left\{V_{\text{stop}}(\vec{\mu}_t, \vec{\sigma}_t, B_t), V_{\text{continue}}(\vec{\mu}_t, \vec{\sigma}_t, B_t)\right\}. \tag{3}$$

Here, the first maximand is the value of stopping the exploratory phase, solving the knapsack problem and proceeding to confirmation for the selected projects, in which case $\tau = t$. The second maximand is the value of continuing to explore at least one project. These terms are

$$V_{\text{stop}}(\vec{\mu}_t, \vec{\sigma}_t, B_t) = V_{\pi,\tau}(\vec{\mu}_t, \vec{\sigma}_t, B_t) \text{ in (2), on the event } \tau = t. \tag{4}$$

$$V_{\text{continue}}(\vec{\mu}_t, \vec{\sigma}_t, B_t) = \max_{x: B_t \geq c_x} \mathbb{E}\left[V(\vec{\mu}_{t+1}, \vec{\sigma}_{t+1}, B_t - c_x) \middle| \vec{\mu}_t, \vec{\sigma}_t, B_t\right]. \tag{5}$$

Note that $V_{\text{stop}}(\cdot)$ is bounded below by B_t, even if there is insufficient budget to explore or confirm any project, because stopping at time $\tau = t$ and selecting no projects for confirmation ($[K]_\tau = \emptyset$) returns the remaining value B_t. If there is insufficient budget to explore any project, then the maximization in (5) has no terms, in which case we define $V_{\text{continue}} = -\infty$ and then it is optimal to stop. If it is optimal to continue exploration, then the maximand in (5) identifies the optimal project to explore next.

3.3 Heuristics for Project Selection in the Exploratory Phase

Solving the Bellman equation is numerically challenging, even if the stopping condition did not require solving a knapsack problem. We therefore consider some heuristic policy classes to simplify computation. We begin with heuristics for project selection during the exploratory phase.

Equal Allocation. Select the project at each time t that has the smallest number of observations so far, $n_{t,x}$, with ties broken so that the project with the smallest index (from $1, 2, \ldots, k$) is selected.

KG Allocations. The knowledge gradient (KG, Frazier et al. 2008) is a useful greedy expected value of information heuristic (Chick and Inoue 2001) that selects the project x at time t that maximizes the expected reward, if a commitment were made to switch to the confirmatory phase after obtaining the observation from the selected project, less the reward of switching to the confirmatory phase without that observation. A generalization is KG_β, in which the figure of merit depends on $\beta \geq 1$ observations rather than just 1 observation.

To be more precise, the KG_β allocation rule at time t of the exploratory phase chooses a project to explore that maximizes the expected value of information in β observations per unit cost of that information:

$$\text{KG}_{x,\beta} = \mathbb{E}\left[V_{\text{stop}}(\vec{\mu}_{t+\beta}, \vec{\sigma}_{t+\beta}, B_t - \beta c_x) - V_{\text{stop}}(\vec{\mu}_t, \vec{\sigma}_t, B_t) \mid \mathscr{F}_t\right] / (\beta c_x). \tag{6}$$

In the case of ties we select the tied project with the lowest index. Even though we allow for $\beta \geq 1$ patients in computing the information, at each time t we pick only one patient and one project for exploration.

If we do not explicitly mention β, its value is presumed to be 1, giving the *KG allocation rule*. The *KG** *allocation rule* samples from the project x that maximizes $\sup_{\beta \geq 1} \text{KG}_{x,\beta}$.

Constrained KG allocations, $\text{KG}^*_{a;b}$. We define the $\text{KG}^*_{a;\cdot}$ allocation as one that first allocates a observations to each project, then uses KG* thereafter (as might be done with an uninformative prior and

a initial stage observations). The $\text{KG}^*_{:,b}$ allocation allocates to the project with the largest KG^*_x among projects with fewer than b observations. Similarly, $\text{KG}^*_{a;b}$ allocates a observations to each project, then allocates with KG^* among projects with fewer than b observations.

Computation. The KG and KG^* allocation rules both require evaluating (6) for at least one β for each of the k projects at each time $t = 0, 1, \ldots, \tau$. However, (6) cannot be evaluated analytically, as it involves two separate instances of V_{stop}, each requiring the solution of a knapsack problem. We therefore use approximations to calculate $\text{KG}_{x,\beta}$. We approximate the expectation in (6) using a Monte Carlo sample average. Each sample corresponds to an independent realization of $\vec{\mu}_{t+\beta}$, conditional on $(\vec{\mu}_t, \vec{\sigma}_t)$, and requires solving a knapsack problem to evaluate each V_{stop}. To speed up computation, we approximately solve each knapsack problem using a greedy algorithm (Martello and Toth 1990, §2.4), which essentially ranks projects in decreasing order of expected return per unit of budget consumed. Besides, when computing KG^* for each project x, instead of solving the maximization over all $\beta \geq 1$, we restrict the search to a discrete set of values $\beta \in \{2^{\ell/2} : \ell = 0, 1, \ldots, 14\}$, an approximation that has proven effective in similar settings. Our experiments below use these approximations. We refer to the indices based on these approximations by $\overline{\text{KG}}_x$ for the KG allocation, by $\overline{\text{KG}}^*_x$ for the KG^* allocation, and by $\overline{\text{KG}}^*_{a;b}$ for the $\text{KG}^*_{a;b}$ allocation.

3.4 Heuristics for Stopping the Exploratory Phase

We also consider several heuristics for stopping the exploratory phase. We define the Fixed Exploration Sample Stopping (**FESS(e_0)**) stopping rule to stop when all projects in the portfolio achieve a specified sample size e_0 for the exploratory phase. And the $\overline{\textbf{KG}}^*$ stopping rule is defined to stop when the estimated value of information from the best β-step lookahead for each project is less than the cost of that information, i.e., stop the first time that $\overline{\text{KG}}^*_x < 0$ for all x.

If an allocation rule permits no further observations, then the stopping rule is assumed to be triggered. The special case $\overline{\text{KG}}^*_{e_0;e_0}$ is equivalent to the Equal allocation with FESS(e_0) stopping.

3.5 An Oracle

It is useful to have a bound on the expected value of the portfolio. To this end, we define an Oracle with perfect information about every arm's true mean INMB. The Oracle's expected value of the portfolio is then the expected reward, averaged over the prior distribution, assuming that at time zero we have the exact INMB, θ_x, of each project x, and thus exploratory trials are not run.

$$\mathbb{E}\left[\max_{[K]_o \subseteq [K]: \sum_{x\in[K]_o} C_x \leq B_0} \left[B_0 + \sum_{x\in[K]_o} (\mathbb{E}[r_\tau(x, \theta_x, \eta_x, b_x)|\theta_x] - C_x)\right] \middle| \mathscr{F}_0\right].$$

We cannot compute this value analytically, so we approximate it using Monte Carlo. Each Monte Carlo replication consists of sampling a realization θ_x and computing the conditional expectation of $r_\tau(x, \theta_x, \eta_x, b_x)$ conditional on θ_x for each x, and solving the knapsack problem to obtain a Monte Carlo approximation of the optimal subset of arms $[K]_o$ and their value. On any replication, if the budget is insufficient to confirm any project, then $[K]_o = \emptyset$, and the Oracle value on that replication is simply the budget, B_0.

4 ILLUSTRATIVE NUMERICAL EXPERIMENTS

We consider a base case portfolio with a set of projects that have independent and identically distributed prior means and jointly independent observations. Parameters are chosen to reflect clinical trials in pediatric oncology where possible (ECORYS et al. 2015; BIO et al. 2021; Oliviero et al. 2022). That said, we need to make assumptions about the values of other parameters. Thus, the results below are illustrative, not prescriptive. We explore the question of whether response adaptation across exploratory phases of projects helps with financial and health metrics, relative to not using response adaptation, and, if so, identify some settings where it is particularly helpful.

4.1 Parameter Settings

We assume a willingness-to-pay of $\lambda = £20,000$/QALY, motivated by UK practice (Claxton et al. 2015). For each project, x, we choose a prior mean of the INMB $\mu_{0,x} = £0$ to not favor the new treatment or comparator. We assume a standard deviation (stdev) of the unknown mean INMB of $\sigma_x = 10\lambda = £200,000$, which is consistent with a stdev of the mean QALY gain of 10 with a known cost of treatment and complications. The effective sample size of the prior for the unknown mean INMB is taken to be $n_0 = \sigma_x^2/\sigma_{0,x}^2 = 25$. For the confirmatory phase, the type 1 and type 2 error levels are assumed to be $\alpha = 0.05$ and $\beta = 0.1$, respectively. The minimal acceptable improvement $\delta_x = £60,000$, so that a treatment must deliver a mean INMB equivalent of 3 or more additional QALYs to be considered meaningful. This yields a required confirmatory sample size $L_x = 95$, similar to sample sizes of 70-100 in pediatric drug development practice.

Regarding the projects' potential economic benefits, we consider an ultra-rare oncology setting with an accessible patient population size of $P_x = 2,000$ for each x. The treatment price p_x of x and the fixed manufacturing cost m_x are set to £80,000 and £20,000, respectively. We assume the bargaining power of the portfolio manager ν is 50%. On the cost side, the variable cost per patient enrolled in both the exploratory and confirmatory phases c_x is £230,000. For the exploratory phase, a typical sample size is 30 patients, giving a total variable cost of £6.9 million. As discussed earlier, we treat the fixed cost of the exploratory phase as sunk. For the confirmatory phase, we assume a 50:50 split between fixed and variable costs. With a required sample size of $L_x = 95$ patients, the total variable cost of each x is £21.9 million, and the total confirmatory cost C_x of each x is £43.8 million.

We then determine a baseline portfolio size and the total budget. For a small-molecule pediatric oncology drug entering clinical trials, the probability of transitioning from the exploratory phase to the confirmatory phase (e.g., from Phase I to Phase II) is typically 40–50%, while the probability of transitioning from the exploratory phase to regulatory approval is estimated to be 10–20%. We target a portfolio that yields 3 drug approvals. This goal translates into $3/20\% = 15$ projects in the portfolio. Given a total variable cost of £6.9 million for each exploratory trial, the total exploration cost of 15 projects is approximately $2C_x$. Furthermore, on average, at least 6 of the 15 projects are expected to transition to the confirmatory phase, resulting in a total confirmatory cost of $6C_x$. Thus, we use a total budget of $8C_x \approx £350.4$ million.

According to the estimated exploratory sample size and exploratory cost of each x, we set $e_0 = 30$ for the FESS stopping heuristic. When computing the $\mathrm{KG}_{x,\beta}$ index for each x in (6), we use 50 independent simulated realizations of $\vec{\mu}_{t+\beta}$ to estimate the expectation. In the experiments, we conduct 300 (3000) macro replications for each adaptive (non-adaptive) policy (a combination of a project selection heuristic and a stopping heuristic) to estimate the performance metrics. Additionally, we include a baseline policy named the **Prior Naive** policy, which selects the drugs to advance to the confirmatory phase based solely on prior information, without conducting any exploration. Thus, Prior Naive resembles the methods for one-shot trials in Fairley et al. (2020) at the confirmatory phase.

4.2 Value of Portfolio-Adaptive Trials

We explore how several combinations of allocation and stopping rules perform as we vary the number of projects or the initial budget, relative to the baseline portfolio (15 projects, $B_0 = 8C_x$). Performance is measured through a financial metric (total expected value), regulatory metrics (the probability at least one trial is approved, the expected number of approved trials), and health economic metrics (the expected number of successful projects approved), among others.

4.2.1 Changing the Number of Projects for a Fixed Budget

Figure 2 shows how several policies perform as the number of projects, k, varies with a fixed budget, B_0.

First, the "Prior Naive" curve is flat, with total expected value at $B_0 = £350.4$ million, independent of k. This is because each of the k projects have a negative NPV when considered alone, so Prior Naive makes no investment, whether in exploratory or confirmatory trials. Policies that correspond to curves

that increase above £350.4 million have the potential, in expectation, to improve value by leveraging the portfolio design and/or adaptive learning. Importantly, this also suggests that a portfolio of projects all at the same Phase can significantly benefit from an initial response adaptive pilot before committing to a subset of projects to complete that Phase. For the Oracle (data not shown), which gives an unachievable upper bound, total expected value increases from £555.8 million to £953.3 million, the probability of an approval increases from 88.0% to 99.9%, and the expected number of approvals increases from 1.9 to 5.2.

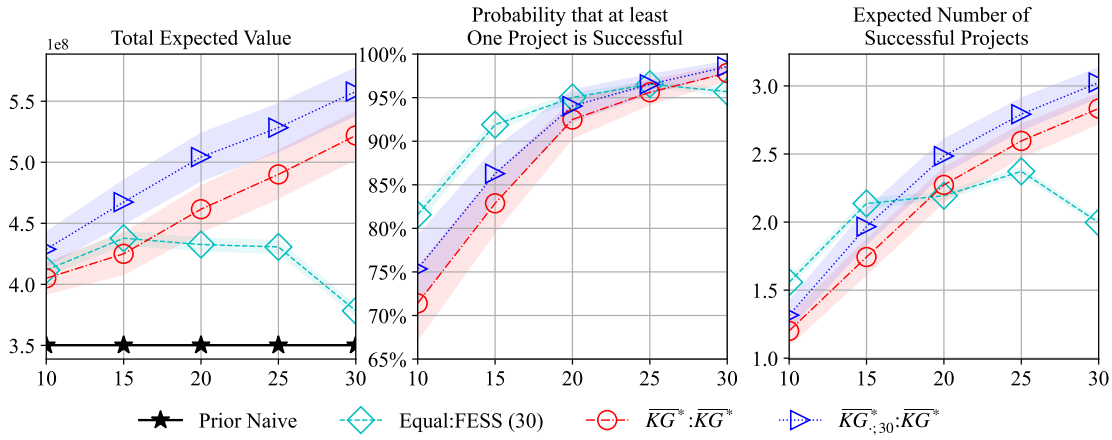

Figure 2: Performance of policies as we vary the number of projects in the portfolio (on the x-axis).

Next, we distinguish two mechanisms to improve upon the NPV of considering projects individually. The first mechanism is the use of a portfolio, which allows the subset of projects that proceed to the confirmatory phase to be selected after seeing the results of the exploratory trials, even without adaptation during the exploratory trials themselves. The value of a standard portfolio can be seen as the difference between the "Equal:FESS(30)" and "Prior Naive" curves in the Total Expected Value graph. These benefits are like those for portfolios in the work of (Fagnan et al. 2014). That difference is greatest with 15 projects in Figure 2, and decreases thereafter. In the graph, the differences decrease for larger numbers of projects, and would become negative if there were more projects to explore than could be paid for with the initial budget. Colored bands describe 95% confidence intervals.

The second mechanism to improve the NPV is to response-adaptively allocate resources across projects in the portfolio during exploration, the value of which depends on which allocation policy is used. The value of response adaptiveness of $\overline{\mathrm{KG}}^*_{\cdot;30} : \overline{\mathrm{KG}}^*$ relative to a standard portfolio is the difference between the $\overline{\mathrm{KG}}^*_{\cdot;30} : \overline{\mathrm{KG}}^*$ and the Equal:FESS(30) curves. This mechanism is new to the literature on biotech portfolios.

The fully adaptive $\overline{\mathrm{KG}}^* : \overline{\mathrm{KG}}^*$ policy underperforms relative to the non-adaptive Equal:FESS(30) policy for all metrics when the number of projects is small (e.g., $k \leq 15$). Sample size analysis suggests that this is because $\overline{\mathrm{KG}}^* : \overline{\mathrm{KG}}^*$ explores too much, having sample sizes far above the typical exploration sample size of 30 for projects that are ultimately selected for confirmation. On the other hand, $\overline{\mathrm{KG}}^* : \overline{\mathrm{KG}}^*$ tends not to over-explore projects that are not, ultimately, confirmed. As the number of projects grows, $\overline{\mathrm{KG}}^* : \overline{\mathrm{KG}}^*$ outperforms Equal:FESS(30). This happens because Equal:FESS(30) consumes so much budget during exploration that little remains for confirmation.

Compared to Equal:FESS(30), the $\overline{\mathrm{KG}}^*_{\cdot;30} : \overline{\mathrm{KG}}^*$ policy dominates in terms of total expected value and produces more successful projects as more projects are included in the portfolio (e.g., when $k \geq 20$); there is positive additional value for this adaptation policy for all values tested in the graph. However, for $k \leq 15$, Equal:FESS(30) performs modestly better on the non-financial metrics. This is not surprising: KG-type allocations optimize total reward, rather than the number of successful projects.

The $\overline{\text{KG}}^*_{\cdot;30} : \overline{\text{KG}}^*$ policy also outperforms $\overline{\text{KG}}^* : \overline{\text{KG}}^*$ across all metrics. The cap at 30 effectively reduces over-exploration of projects in the confirmatory set under the $\overline{\text{KG}}^* : \overline{\text{KG}}^*$ policy, while the exploratory sample sizes for confirmed projects rarely drop much below 30. This reduction in exploration enables more projects to be confirmed compared to $\overline{\text{KG}}^* : \overline{\text{KG}}^*$, leading to improvements overall.

4.2.2 Changing the Budget for a Fixed Number of Projects

We also explored performance measures when the budget varies from $B_0/C_x = 8$ times the cost of one confirmatory trial, keeping the number of projects fixed at $k = 15$. See Figure 3.

The total expected value achieved by the Equal:FESS(30) policy for the baseline setting with $k = 15$ and $B_0/C_x = 8$, can be achieved by the $\overline{\text{KG}}^*_{\cdot;30} : \overline{\text{KG}}^*$ and $\overline{\text{KG}}^* : \overline{\text{KG}}^*$ policies with about 12.5% less initial capital, i.e, with $B_0/C_x \approx 7$, corresponding to £43.8 million less initial capital. However, for B_0/C_x at either 7 or 8, this improved financial metric comes at the expense of other metrics, i.e., those relating to the number of approved and/or successful projects, which are slightly lower than those of Equal:FESS(30).

Exploring this observation further, $\overline{\text{KG}}^*_{\cdot;30} : \overline{\text{KG}}^*$ achieves an 7.6% higher total expected value than Equal:FESS(30) with a budget of $8C_x$, thus achieving its goal of optimizing financial outcomes. A side effect is a drop of 8.3% in the number of successful projects. That is perhaps undesirable, but it is not the metric we optimize in this case study. With bargaining power $\nu = 0$ and no budget and confirmatory cost terms in (2), the objective is maximizing the number of successful projects. Then, $\overline{\text{KG}}^*_{5;\cdot} : \overline{\text{KG}}^*$ outperforms Equal:FESS(30) for the expected number of successful projects in our case study (data not shown).

For smaller budgets with $B_0/C_x \leq 6$, the adaptive $\overline{\text{KG}}^*_{\cdot;30} : \overline{\text{KG}}^*$ policy attains both higher total expected value and more successful projects than does Equal:FESS(30). This highlights the value of response adaptation in a limited budget setting, which is typical for pediatric oncology.

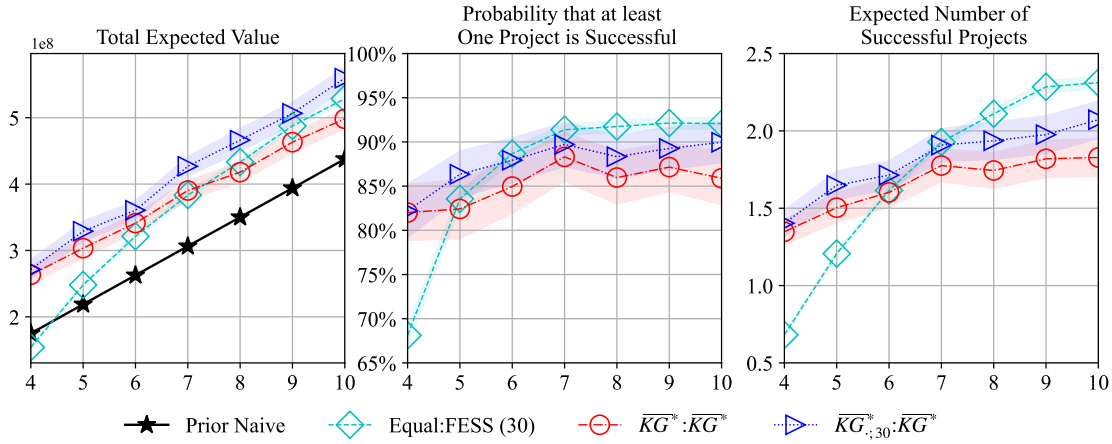

Figure 3: Performance of policies as the budget varies. The x-axis is B_0/C_x.

5 DISCUSSION AND CONCLUSION

We explored the question of how two important concepts, portfolios of projects and response-adaptive clinical trials, might be combined to improve financial return to investors (total expected value) and health-economic benefits (the number of successful projects) in drug development for pediatric oncology. The results indicate value from both concepts: the value of a portfolio of projects is demonstrated by the fact that the Equal:FESS(30) policy that exploits a portfolio greatly outperforms Prior Naive, which does not. Moreover, the additional value of adaptation is clear from the performance of the adaptive policies relative to Equal:FESS(30). We find that capital requirements can be reduced when exploiting these two concepts, particularly in settings like pediatric oncology or rare diseases where capital may be limited.

Our study has limitations. For example, we did not address the portfolio construction (asset identification) and liquidation (philanthropic and equity tranche management) problems, nor the organizational overhead to align projects across research groups. We ignored the fixed costs of including projects. And outputs from different phases of a given project may differ in practice. Financial incentives may also have a bearing beyond the bargaining power parameter.

A number of research avenues have clear value. They include discounted rewards, delays in observing outcomes, and running exploratory trials in parallel. Work with delays in single-trial contexts (e.g., Yapar et al. 2025) suggests an approach that involves incurring the costs to kick-start trials, then triggering the stopping time for exploration based on data observed so far. Modeling of the end-points of the different phases would have practical relevance, as would response-adaptation during the confirmatory phase, too. We assumed independence of projects in both patient outcomes and markets, yet in reality the portfolio may include multiple potential treatments for a given indication and/or multiple potential indications for a given treatment. We could use simulation to generate power curves for confirmatory phases, and to allow for more than 2 trial phases. And, as always, the specification of prior distributions is important.

We believe the area of portfolio-adaptive trials is worth further exploration with simulation optimization tools. Our initial results show that portfolio-adaptive trials may help to de-risk portfolios of investments in biotechnology, and may particularly benefit areas where financial resources are constrained.

ACKNOWLEDGMENTS

The authors thank David Goldberg and Qingyuan Chen for helpful discussions. Chick and Li thank Dr Simba Gill and Sabi Dau for their support of INSEAD's Healthcare Management Initiative. This work was partially supported by National Science Foundation Grants DMS-2230023 and OAC-2410950.

REFERENCES

Angus, D. C., B. Alexander, S. Berry *et al*. 2019. "Adaptive Platform Trials: Definition, Design, Conduct and Reporting Considerations". *Nature Reviews Drug Discovery* 18:797–807.

Berry, D. A. 2011. "Adaptive Clinical Trials in Oncology". *Nature Reviews Clinical Oncology* 9(4):199–207.

BIO, QLS Advisors, and Informa UK. 2021. *Clinical Development Success Rates and Contributing Factors, 2011-2020*. Biotechnology Innovation Organization. Available at https://www.bio.org/clinical-development-success-rates-and-contributing-factors-2011-2020, Accessed 24 Mar 2025.

Branke, J., S. Chick, and C. Schmidt. 2007. "Selecting a Selection Procedure". *Management Science* 53(12):1916–1932.

Chick, S. E., and P. I. Frazier. 2012. "Sequential Sampling for Selection with Economics of Selection Procedures". *Management Science* 58(3):550–569.

Chick, S. E., N. Gans, and O. Yapar. 2022. "Bayesian Sequential Learning for Clinical Trials of Multiple Correlated Medical Interventions". *Management Science* 68(7):4919–4938.

Chick, S. E., and K. Inoue. 2001. "New Two-Stage and Sequential Procedures for Selecting the Best Simulated System". *Operations Research* 49(5):732–743.

Claxton, K., S. Martin, M. Soares, N. Rice, E. Spackman, S. Hinde *et al*. 2015. "Methods for the Estimation of the National Institute for Health and Care Excellence Cost-Effectiveness Threshold". *Health Technology Assessment* 19(14):1–503, v–vi.

Daems, S., H. Stevens, M. Dewatripont, H.-G. Eichler, and M. Goldman. 2023. "A Novel Approach to Boost Drug Development in Paediatric Oncology". *Nature Reviews Drug Discovery* 22(10):769–770.

Das, S., R. Rousseau, P. C. Adamson, and A. W. Lo. 2018. "New Business Models to Accelerate Innovation in Pediatric Oncology Therapeutics". *JAMA Oncology* 4(9):1274–1280.

Eckman, D. J., and S. G. Henderson. 2022. "Posterior-Based Stopping Rules for Bayesian Ranking-and-Selection Procedures". *INFORMS Journal on Computing* 34(3):1711–1728.

ECORYS, Empirica, Technopolis, S. K *et al*. 2015. *Study on the Economic Impact of the Paediatric Regulation, including Its Rewards and Incentives - Final Report (Redacted Version)*. Publications Office of the European Union. Available at https://data.europa.eu/doi/10.2875/662696, Accessed 13 April 2025.

EMA 2025. "Authorization of Medicines". European Medicines Agency, Available at https://www.ema.europa.eu/en/about-us/what-we-do/authorisation-medicines, Accessed 3 April 2025.

Fagnan, D. E., A. A. Gromatzky, R. M. Stein, J.-M. Fernandez, and A. W. Lo. 2014. "Financing Drug Discovery for Orphan Diseases". *Drug Discovery Today* 19(5):533–538.

Fairley, M., L. Cipriano, and J. Goldhaber-Fiebert. 2020. "Optimal Allocation of Research Funds under a Budget Constraint". *Medical Decision Making* 40(6):797–814.

Fenwick, E. *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.

Fernandez, J., R. Stein, and A. Lo. 2012. "Commercializing Biomedical Research through Securitization Techniques". *Nature Biotechnology* 30:964–975.

Frazier, P. I., W. B. Powell, and S. Dayanik. 2008. "A Knowledge-Gradient Policy for Sequential Information Collection". *SIAM Journal on Control and Optimization* 47(5):2410–2439.

Hampson, L. V., and C. Jennison. 2013. "Group Sequential Tests for Delayed Responses (with Discussion)". *Journal of the Royal Statistical Society Series B: Statistical Methodology* 75(1):3–54.

Hong, L. J., W. Fan, and J. Luo. 2021. "Review on Ranking and Selection: A New Perspective". *Frontiers in Engineering Management* 8:321–343.

Li, Z., S. E. Chick, S. Daems, and S. G. Henderson. 2025. "Adaptation Across Trials in a Portfolio of Drug Development Projects". Work in progress.

Locatelli, A., M. Gutzeit, and A. Carpentier. 2016. "An Optimal Algorithm for the Thresholding Bandit Problem". In *Proceedings of The 33rd International Conference on Machine Learning*, edited by M. F. Balcan and K. Q. Weinberger, 1690–1698. New York: PMLR.

London, A. J., and J. Kimmelman. 2019. "Clinical Trial Portfolios: A Critical Oversight in Human Research Ethics, Drug Regulation, And Policy". *Hastings Center Report* 49(4):31–41.

Martello, S., and P. Toth. 1990. *Knapsack Problems: Algorithms and Computer Implementations*. John Wiley & Sons, Inc.

Oliviero, E., G. Kourkopoulos, and J. Kimmelman. 2022. "Success Rates for US and Canadian Anticancer Drug Development Efforts in Pediatric Oncology". *Pediatric Blood & Cancer* 69(8):e29534.

Pallmann, P., A. W. Bedding, B. Choodari-Oskooei, M. Dimairo, L. Flight, L. V. Hampson, *et al*. 2018. "Adaptive Designs in Clinical Trials: Why Use Them, and How to Run and Report Them". *BMC Medicine* 16:1–15.

Powell, W. B., and I. O. Ryzhov. 2012. *Optimal Learning*. John Wiley & Sons.

Strong, M., J. E. Oakley, A. Brennan, and P. Breeze. 2015. "Estimating the Expected Value of Sample Information Using the Probabilistic Sensitivity Analysis Sample: A Fast, Nonparametric Regression-Based Method". *Medical Decision Making* 35(5):570–583.

Williamson, S. F., P. Jacko, S. S. Villar, and T. Jaki. 2017. "A Bayesian Adaptive Design for Clinical Trials in Rare Diseases". *Computational Statistics and Data Analysis* 113:136–153.

Xie, J., and P. I. Frazier. 2013. "Sequential Bayes-Optimal Policies for Multiple Comparisons with a Known Standard". *Operations Research* 61(5):1174–1189.

Yapar, O., S. E. Chick, and N. Gans. 2025. "Conditional Approval and Value-Based Pricing for New Health Technologies". *Management Science* 71(0):in press.

AUTHOR BIOGRAPHIES

ZAILE LI is a post-doctoral researcher at the INSEAD Healthcare Management Initiative. His research explores adaptive selection algorithms, particularly in applications such as clinical trial design. His email address is zaile.li@insead.edu and his website is https://sites.google.com/view/zaileli/.

STEPHEN E. CHICK holds the Novartis Chair of Healthcare Management in the Technology and Operations Management and is the academic director of INSEAD's Healthcare Management Initiative. His research interests include stochastic simulation, Bayesian inference, sequential learning, and applications in health care innovation and public health. His email address is stephen.chick@insead.edu and his website is https://www.insead.edu/faculty/stephen-e-chick.

SAM DAEMS is researcher at the Institute for Interdisciplinary Innovation in Healthcare, at the Université Libre de Bruxelles. His research explores innovative funding models for pediatric oncology. Additionally, he is Portfolio Director at Waterland Private Equity in Antwerp, Belgium. His email address is daems@waterland.be and his websites are https://www.ulb.be/fr/sam-daems and https://www.waterlandpe.com/people/sam-daems/.

SHANE G. HENDERSON holds the Charles W. Lake, Jr. Chair in Productivity in the School of Operations Research and Information Engineering at Cornell University. His research interests center on simulation and its applications in emergency services, transportation and, since the COVID-19 pandemic, intervention design for epidemics. He is a co-creator of SimOpt, a testbed of simulation optimization problems and solvers. He is an INFORMS Fellow. His email address is sgh9@cornell.edu and his website is http://people.orie.cornell.edu/shane.