Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

SEQUENTIAL, VALUE-BASED DESIGNS FOR CERTAIN CLINICAL TRIALS WITH MULTIPLE ARMS HAVING CORRELATED REWARDS

Stephen E. Chick

INSEAD – Europe Campus
Technology and Operations Management Area
Boulevard de Constance
Fontainebleau, 77300, FRANCE

Noah Gans
Ozge Yapar

OID Department – Wharton School
University of Pennsylvania
3730 Walnut Street, Suite 500
Philadelphia, PA 19104-6340, USA

ABSTRACT

The time and money required to run clinical trials, as well as the cost effectiveness of technologies emerging from these trials are receiving increasing scrutiny. This paper explores the use of techniques inspired from fully sequential simulation optimization, based on Bayesian expected value of sampling information arguments, in the context of highly-sequential multi-arm trials. New allocation rules are shown to be useful for selecting which technology to assign a patient in a trial. They are based on clinical cost-benefit tradeoffs and the size of the population who benefits from the technology adoption decision.

1 INTRODUCTION

This paper applies a Bayesian decision-theoretic framework that has been developed in the context of simulation optimization (Chick and Gans 2009; Chick and Frazier 2012) to the problem of designing certain clinical trials that account for health technology adoption (UK NICE 2014). Clinical trials are typically used to assess whether or not a new technology (drug, device, process change, etc.) achieves a given threshold of clinical improvement, and they tend to be designed using a frequentist approach, using a minimal number of observations in one or a small number of stages of sampling, with the goal of reaching statistical confidence and power criteria (Jennison and Turnbull 1999). Even with a positive result from the clinical trial, there is still the question of whether the technology will be adopted, will be eligible for reimbursement, and will be implemented for patients. This health technology assessment can use cost effectiveness analysis (HAS 2012; UK NICE 2014). While clinical trials have traditionally been designed as hypothesis tests of clinical efficacy, the idea of considering cost effectiveness in clinical-trial design has recently gained attention.

Regulatory bodies are encouraging the development of novel methodologies for clinical trial design (EMA 2007; FDA 2010; FDA 2016). Three important trends have arisen to address the high financial costs, potentially long delays, and health-benefit concerns of clinical trials: (a) highly adaptive designs, including sequential designs, that allow allocation of patients to treatments to vary as a function of data observed so far, with the hope of reducing the time and number of patients required to assess new treatments (Berry 2012; Pallmann et al. 2018), (b) value-based designs, which use the financial costs, health benefits, and size of the population meant to benefit from the health technology decision, to guide the amount of sampling required, as opposed to relying on type I and type II error criteria (Draper 2013; Chick et al. 2017); (c) multi-arm trials, which compare multiple treatments to a control, rather than requiring each of several new treatments to have its own control group, thereby reducing the number of patients required to identify the best of several treatments (EMA 2017; Boeree et al. 2017). The mixing of multiple treatments in various combinations and/or dose levels can result in correlation in the mean rewards of arms: we thus explore

the extension of so-called correlated knowledge gradient techniques (cKG from Frazier et al. (2009)) with diffusion approximation techniques (Chick and Gans 2009; Chick and Frazier 2012) for a richer assessment of the expected value of sample information.

This paper extends a parallel project Chick et al. (2019) with correlated arms. Chick et al. (2019) focuses on seamless Phase II/III trials which attempt to optimize treatment doses, we focus here on factorial trial designs, where the multiple arms represent various multiple treatments in a Phase III trial are tested in various combinations with each other. This implies different correlation structures, different prior distributions, and different nuances regarding the benefits of modeling correlation across arms.

2 FULLY SEQUENTIAL VALUE-BASED TRIALS WITH CORRELATED ARMS

We recall a simplified version of the basic framework of Chick et al. (2019), which models highly sequential trials with many arms whose mean responses may be correlated. That paper explores dose response trials, while this paper applies that framework for different applications in which each arm represents combinations of several therapies. The trial is designed to assess $M \geq 2$ intervention alternatives, or *arms*, one of which may be a control or a current standard of care. The arm selected for implementation at the end of the trial will be used to treat P patients.

2.1 Model for Fully Sequential Value-Based Trials with Correlated Arms

Before selecting an arm to implement, we can sample from it to obtain information. Sampling from an arm requires effort and money, and there is a cost c_i per sample for arm $i \in \mathscr{M}$. We require here that all $c_i > 0$ so that we will not sample costlessly over an infinite horizon.

Sampling is modelled as occurring sequentially at equally spaced times in discrete stages indexed by $t = 0, 1, \ldots$. The index t also represents the total number of patients whose samples have been observed. We denote by T the (potentially random) time at which we stop the trial and select an arm to implement.

At time $t = 0, 1, \ldots, T-1$, we decide which arm patient $t+1$ will receive and can use the data observed from the first t patients in doing so. Let Y_i^t be the random variable whose realization is the incremental net monetary benefit of arm $i \in \mathscr{M}$ for patient $t = 1, 2, \ldots, T$. The value of Y_i^t is intended to represent the combined measure of health gained in monetary value, less the costs of treating with arm i. This might be measured by using QALY and willingness to pay parameters (Gold et al. 1996; UK NICE 2014).

Each Y_i^t has an unknown mean θ_i and a known sampling variance λ_i. Here, θ_i can be interpreted as the mean arm effect for the population, and λ_i captures random differences in how individual patients respond to arm i. We assume that observations are independent and normally distributed, so that

$$Y_i^t | \theta_i \sim \mathscr{N}(\theta_i, \lambda_i)$$

for $t = 1, 2, \ldots$ and $i \in \mathscr{M}$. Let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_M)^\mathsf{T}$ be the vector of unknown means. We assume for now the λ_is are known. One can use plug-in estimators in practice.

A prior distribution for $\boldsymbol{\theta}$ describes our initial uncertainty about the M arms' mean effectiveness. This initial belief about $\boldsymbol{\theta}$ is distributed according to a multivariate normal prior, with $\boldsymbol{\theta} \sim \mathscr{N}(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0)$. We assume that $\boldsymbol{\Sigma}^0$ is positive semidefinite. It can be non-diagonal, which would imply that the initial beliefs about the means are correlated. We discuss the specification of $\boldsymbol{\Sigma}^0$ for different types of trials in Section 3.3.

At each time, t, we observe the outcome of the tth patient's intervention, and we use this observation to update our beliefs about the mean arm effects. Then, we choose either to stop the trial or to continue and include one more patient. A decision to stop is followed by the selection of an arm for implementation, and a continuation decision requires the choice of an arm to allocate to the next patient. If delays in responses are long, one reasonable approximation might be to base continuation decisions on data observed to time t.

To track our choices, we define a number of variables. At each time, t, an action u^t is chosen from the set of available actions $\mathscr{U} = \{i | i = 1, 2, \ldots, 2M\}$, with M actions for continuation and M actions for stopping. We let $u^t \in \{i | i = 1, 2, \ldots, M\}$ denote allocating arm u^t to a patient and we let $u^t \in \{i | i = M+1, M+2, \ldots, 2M\}$

denote stopping the trial and selecting arm $\mathscr{D} = u^t - M$ to implement. After an action $u^t \in \{i | i = 1,2,\ldots,M\}$ is chosen, the observation $Y_{u^t}^{t+1}$ is realized before the next period's decision. The belief about the unknown mean at time $t+1$, $(\boldsymbol{\mu}^{t+1}, \boldsymbol{\Sigma}^{t+1})$, is then computed using $(\boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t)$, the observation $Y_{u^t}^{t+1}$, and Bayes' rule. At the first occurrence of $u^t \in \{i | i = M+1, M+2, \ldots, 2M\}$, the trial stops, so that the stopping time is $T = t$.

A policy π gives a sequence of arms $\{u^0, u^1, \ldots\}$ to test in the trial, a stopping time T, and an arm $\mathscr{D}$ to implement. We define Π to be the set of all nonanticipating policies, whose choices at time $t = 0, 1, \ldots$ depend only on the history up to t: $\{\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0, Y_{u^0}^1, \boldsymbol{\mu}^1, \boldsymbol{\Sigma}^1, Y_{u^1}^2, \ldots, \boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t\}$.

Given the prior distribution $\mathscr{N}(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0)$ and a policy $\pi \in \Pi$, the expected discounted value of the future stream of rewards is the sum of the cost of sampling, and the health benefit to the P patients treated by arm $\mathscr{D}$ upon stopping the trial:

$$V(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0) \;=\; \mathbb{E}_\pi \left[\sum_{t=0}^{T-1} -c_{u^t} \;+\; P\, \mathbb{E}\left[Y_{\mathscr{D}}^{T+1} \middle| \boldsymbol{\mu}^T, \boldsymbol{\Sigma}^T \right] \;\middle|\; \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0 \right]. \tag{1}$$

We focus on the problem of choosing a policy $\pi^* \in \Pi$ that maximizes the expected net reward. This problem can be considered a stoppable version of the multi-armed bandit with correlated mean rewards.

$$V^*(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0) \;=\; \sup_{\pi \in \Pi} V(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0). \tag{2}$$

2.2 Comments on Analysis of The Model

We will define index-based heuristics for several potentially good, albeit suboptimal, allocation and stopping rules for Eq. (2) that involve the expected value of sample information (EVSI) over different sampling plans. Policies $\pi_i \in \Pi_i$ sample one or more times in a potentially adaptive, non-anticipative way, from a single arm i, before stopping at time T_i and selecting any arm for implementation. Different sets of sampling policies Π_i will lead to different indices, as discussed below in Section 3.1. The EVSI equals Eq. (1) evaluated with $\pi = \pi_i$, minus the expected reward of immediately stopping and implementing the best alternative without sampling. Because indices are recomputed after each sample is observed, we analyse Eq. (1) assuming sampling starts at time t rather than at time 0.

Although we will allocate arms to one patient at a time, we also consider Bayes' rule with $\tau \geq 1$ observations $Y_i^{t+1}, Y_i^{t+2}, \ldots, Y_i^{t+\tau}$ from arm i, starting at time t. The mean of these observations is $\overline{Y}_i^\tau = \sum_{r=t+1}^{t+\tau} Y_i^r / \tau \sim \mathscr{N}(\theta_i, \lambda_i/\tau)$, and Bayes' rule provides the following posterior mean and covariance (Frazier et al. 2009), where $\mathbf{e_i}$ is a column vector of 0s with M elements except for a 1 in the ith row:

$$\boldsymbol{\mu}^{t+\tau} = \boldsymbol{\mu}^t + \left(\overline{Y}_i^\tau - \mu_i^t \right) \boldsymbol{\Sigma}^t \mathbf{e_i} / (\lambda_i/\tau + \Sigma_{i,i}^t), \tag{3}$$

$$\boldsymbol{\Sigma}^{t+\tau} = \boldsymbol{\Sigma}^t - \boldsymbol{\Sigma}^t \mathbf{e_i} \mathbf{e_i}^\mathsf{T} \boldsymbol{\Sigma}^t / (\lambda_i/\tau + \Sigma_{i,i}^t). \tag{4}$$

We define $Z_i^\tau \equiv \mu_i^{t+\tau} - \mu_i^t$ to express the change in the mean belief of arm i,

$$\boldsymbol{\mu}^{t+\tau} = \boldsymbol{\mu}^t + Z_i^\tau \boldsymbol{\Sigma}^t \mathbf{e_i} / \Sigma_{i,i}^t. \tag{5}$$

To simplify notation, we denote the effective number of samples for arm i at time t by $n_i^t = \lambda_i / \Sigma_{i,i}^t$. The distribution of Z_i^τ for a given τ and information at time t is then (Chick and Gans 2009)

$$Z_i^\tau \sim \mathscr{N}\left(0, \sigma_{Z_i^\tau}^2\right) \text{ where } \sigma_{Z_i^\tau}^2 = \frac{\lambda_i \tau}{n_i^t \left(n_i^t + \tau\right)}. \tag{6}$$

We now use Eq. (5) to adapt Eq. (1) so that posterior means at time T_i are expressed as a function of $Z_i^{T_i}$:

$$V_i^*(\boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t) \;=\; P \times \sup_{\pi_i} \mathbb{E}_{\pi_i} \left[\sum_{r=0}^{T_i - 1} -c_i/P \;+\; \max_j \left\{ \mu_j^t + \Sigma_{i,j}^t Z_i^{T_i} / \Sigma_{i,i}^t \right\} \;\middle|\; \boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t \right]. \tag{7}$$

The second term in the expectation in Eq. (7) involves the maximum of multiple linear functions of a single, normally-distributed random variable. We use techniques of Frazier et al. (2009), originally developed for sampling plans restricted to one additional sample, to transform this maximization into a summation when there is an arbitrary, path-dependent, non-anticipative sampling plan, as long as that plan samples only from a single alternative (Chick et al. 2019). Following Frazier et al. (2009), we define

$$a_j = \mu_j^t \text{ and } b_j = \Sigma_{i,j}^t / \Sigma_{i,i}^t. \tag{8}$$

Let $g(z) = \max\{\text{argmax}_j\{a_j + b_j z\}\}$ be the function that returns the index with the largest $a_j + b_j z$ value when evaluated at $z \in \mathbb{R}$, breaking ties by choosing the largest index. Let $S = \{j \,|\, \exists z \in \mathbb{R}\, s.t.\, g(z) = j\}$ be the set of arms which might possibly be selected as best after sampling only from arm i, and let $M' = |S|$ be the number of these arms. Let (l) denote the arm in the set S that has the lth lowest slope, so that $b_{(1)} < \ldots < b_{(l)} < b_{(l+1)} < \ldots < b_{(M')}$. Notice that $k \notin S$ if $b_k = b_j$ and $a_k \leq a_j$. Therefore the ordering is strict. The new ordering implies that $\max\{z \,|\, g(z) = (l)\} = \min\{z \,|\, g(z) = (l+1)\}$, and we can express the intersection point, in z, between arms (l) and $(l+1)$ as

$$d_{(l)} = (a_{(l)} - a_{(l+1)})/(b_{(l+1)} - b_{(l)}). \tag{9}$$

3 ALGORITHM FOR SEQUENTIAL TRIALS

Each policy $\pi \in \Pi$ leads to a sampling and selection algorithm for the clinical trial, with expected reward in Eq. (1). In algorithmic form, each π specifies: a *stopping rule*, which decides at each time t whether to stop or to continue sampling; an *allocation rule*, which specifies which arm to allocate to the next patient at each time t when sampling continues; and a *selection rule*, which specifies which arm to select for implementation when sampling is stopped. In addition, the parameters of the *prior distribution* for the correlated unknown mean reward of each arm must be specified. This section discusses choices for all of these features of our algorithm for fully sequential, correlated, multi-arm trials.

When stopping at time $T = t$, standard techniques can be used to show that the optimal *selection rule* is to select the arm with largest posterior mean for implementation, arm $\mathscr{D} = \text{argmax}_j \mu_j^T$.

The allocation rule and stopping rule of an optimal policy π^* for Eq. (2) are difficult to characterize and compute. Our problem is an undiscounted, stoppable bandit problem with correlated mean rewards, whereas typical Gittin's index results rely on independent arms and discounted rewards. However, a number of heuristic indices have proven to be useful for defining 'good' allocation rules and stopping rules in related contexts, and we apply or extend them for our analysis here. All of our indices for arm i are derived by actively obtaining sample measurements only from arm i. Even so, correlation among our beliefs regarding the population means of the arms implies that, by sampling from i, we gain information regarding other arms' mean performance as well. See Chick et al. (2019) for more details.

3.1 Allocation Rules for Sampling with Correlated Mean Rewards

We use several indices based on the EVSI of different sampling plans, as mentioned in Section 2.2. In particular, we denote by *cKG1* the special case of using the EVSI from *one* additional sample from a given arm, to determine the allocation index for that arm, as proposed for the *correlated knowledge gradient* of Frazier et al. (2009). We extend cKG1 to *cKG** to define the index of arm i to the be supremum of EVSI per sample in τ samples, the supremum over values of $\tau \geq 1$. These are myopic rules.

We also consider allocation rules which better approximate the EVSI of sampling from a single arm before selecting an alternative, in that they compute the EVSI of adaptive sampling plans from a given arm, rather than the more restrictive KG approach which considers the EVSI of fixed-duration sampling plans. We approximate these adaptive plans with an optimal stopping problem for a diffusion model, a free-boundary partial differential equation, hence the "PDE" in their names. The *cPDELower* and *cPDEUpper* allocation rules approximate the EVSI of non-anticipative sampling policies that restrict sampling to arm i while accounting for changes in beliefs about all arms, not just arm i, due to correlation. The terms Lower and

Upper refer to the fact that these allocation rules use lower and upper bounds on the EVSI, the upper based on a Bonferroni-type inequality (Chick et al. 2019). We also consider the *ESPb* and *ESPB* allocation rules (Chick and Frazier 2012), which consider arbitrary non-anticipative sampling plans from a given arm when computing its index, but which do not consider the correlation in mean rewards across arms. The ESPb index favors alternatives whose posterior means are 'furthest inside', in some sense, the continuation region for sampling, and the ESPB rule more directly models the EVSI when arms are independent.

Computations for the above indices consider the EVSI of potentially multiple samples, but only one patient sample is allocated per time step, after which all indices are recomputed.

Finally, we also study two non-index based allocation rules: *Equal* allocation, which allocates patients to the arm with the smallest number of samples, and *Variance* allocation, which allocates patients to the arm whose posterior variance for the unknown mean is largest.

3.2 Stopping Rules for Sampling with Correlated Mean Rewards

We consider a *Fixed* stopping rule, which stops after a predetermined number of samples, as well as adaptive stopping rules. For each allocation rule above, one can define a corresponding stopping rule that directs sampling to continue if at least one arm has an EVSI (or approximation to EVSI) that is greater than 0. Otherwise, sampling stops. This generates cKG1, cKG*, cPDELower, cPDEUpper, ESPb, and ESPB stopping rules. It is possible to match an allocation rule with one EVSI approximation with a stopping rule based on another approximation (for example, the cKG1 allocation rule can be matched with the cPDEUpper stopping rule). In numerical experiments, we often found a benefit to matching the EVSI approximation of the allocation rule with the same EVSI approximation of the stopping rule.

3.3 Prior Distribution for Unknown Mean Rewards

In general, prior means, μ_i^0, and variances, $\Sigma_{i,i}^0$, can be elicited from medical experts or from the results of earlier phases of a trial. Alternatively, one can begin with a non-informative prior distribution, collect an initial set of samples, and use the resulting posterior distribution as a prior distribution. Sampling variances can be estimated from prior experimentation (e.g., from phase 2 trial data for use in phase 3 trial experimentation, or from past performance, or by use of plug=in estimators as experimentation occurs when new combinations of therapies are tested).

Here, we focus on converting knowledge about the relationship among arms' means to a prior covariance matrix, $\boldsymbol{\Sigma}^0$. We do so for the special case of a factorial trial design which compares a number of interventions and their combinations. This case is applied below in Section 5. A related discussion for dose-finding trials that compare multiple dose levels of the same drug is found in Chick et al. (2019).

We consider a clinical trial with $M = 4$ arms: a placebo, intervention A, intervention B and a combination of interventions A and B. We formulate a linear function that describes the average patient response to these arms and serves as a framework for defining a prior distribution for the mean intervention effects, $\boldsymbol{\theta}$.

The observation we obtain from tth patient for arm $i \in \mathscr{M}$ can be expressed as $Y_i^t = \theta_i + \varepsilon_i$, where $\varepsilon_i \sim \mathscr{N}(0, \lambda_i)$ is the error term. Alternatively, we can express Y_i^t as the following linear function in which $\mathbb{1}_j$ denotes the indicator variable specifying whether intervention $j = A, B$ is administered, along with a vector of coefficients, $\boldsymbol{\xi} = (\xi_0\ \xi_1\ \xi_2\ \xi_3)^\mathsf{T}$,

$$Y_i^t = \xi_0 + \xi_1 \mathbb{1}_A + \xi_2 \mathbb{1}_B + \xi_3 \mathbb{1}_A \mathbb{1}_B + \varepsilon_{\mathbb{1}_A, \mathbb{1}_B}. \tag{10}$$

The baseline response, with the placebo, is captured by ξ_0. We denote the average incremental effects of interventions A and B by ξ_1 and ξ_2, respectively. Finally, the average additional effect of combining the two interventions is ξ_3.

While incremental intervention effects, $\boldsymbol{\xi}$, can explicitly model the relationship among arms, we are ultimately interested in defining the relationship among the mean intervention effects, $\boldsymbol{\theta}$, and we can convert a prior distribution on $\boldsymbol{\xi}$, denoted by $\mathscr{N}(\tilde{\boldsymbol{\mu}}^0, \tilde{\boldsymbol{\Sigma}}^0)$, to a prior distribution on $\boldsymbol{\theta}$.

Here, we construct an example prior covariance matrix structure among the elements of $\boldsymbol{\xi}$. We separately decide on a vector of prior variances of ξ's, diag($\tilde{\boldsymbol{\Sigma}}^0$), and a prior correlation matrix among ξ's, $\tilde{\boldsymbol{\rho}}^0$, and use these two to form the prior covariance matrix $\tilde{\boldsymbol{\Sigma}}^0$.

We assume that the baseline response from the placebo is independent of the incremental intervention effects so that the correlation between ξ_0 and ξ_i is zero for $i = 1,2,3$. We denote the correlation between ξ_1 and ξ_2 by η^0. We assume that the correlations between ξ_3 and ξ_i for $i = 1,2$ are equal, and denote them by ζ^0. Thus, the prior correlation matrix of incremental effects is

$$\tilde{\boldsymbol{\rho}}^0 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & \eta^0 & \zeta^0 \\ 0 & \eta^0 & 1 & \zeta^0 \\ 0 & \zeta^0 & \zeta^0 & 1 \end{bmatrix}. \tag{11}$$

Using different values of η^0 and ζ^0, we can model a range of potential relationships among interventions. If interventions A and B work through the same mechanism, then we would assume that $\eta^0 > 0$ and $\zeta^0 < 0$, because the positive response to one intervention implies a positive response for the other, while a combination of the two interventions would not double the effect. If A and B work through different mechanisms, however, and the two interventions boost each other's effect when combined, then we assume $\eta^0 = 0$ and $\zeta^0 > 0$. If they work via different mechanisms but impair each other, we let $\eta^0 = 0$ and $\zeta^0 < 0$.

4 INFLUENCE OF CORRELATION ON OPTIMAL STOPPING BOUNDARIES

Before assessing the performance of the heuristic indices on clinical trials, we first explore the effect of correlation on the optimal stopping boundaries when two arms have correlated means. Consider the problem that starts at time $t = 0$ and samples from one of two arms $i \in \{1,2\}$. We denote by $-i$ the arm from which we are not sampling. At each time $t \geq 0$, based on our belief, $(\boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t)$, we decide whether to continue sampling from arm i or to stop and implement arm $\text{argmax}_{j=1,2}\, \mu_j^t$.

For notational simplicity, we continue this section with a domain that is based on the effective number of samples, $n_i^t = \lambda_i/\Sigma_{i,i}^t = n_i^0 + t$, in place of the time $t \geq 0$. Because we are sampling only from arm i, the posterior of arm i, (μ_i^t, n_i^t), is sufficient to calculate the posterior distribution for all arms, $(\boldsymbol{\mu}^t, \boldsymbol{\Sigma}^t)$. Therefore, our stopping decision at time t is based on the value of (μ_i^t, n_i^t). If it is optimal to continue sampling, we say that (μ_i^t, n_i^t) lies inside the continuation region. Otherwise, it is optimal to stop, and (μ_i^t, n_i^t) falls outside the continuation region.

If the prior distribution assumes that the arms are independent, μ_{-i}^t stays constant at μ_{-i}^0 as we sample from arm i. In this case, Chick and Frazier (2012) show that the optimal stopping boundary can be described by scaling that of a standardized problem, $\partial_{CF}(s)$. In this case the optimal stopping boundaries in (μ_i^t, n_i^t) coordinates can be approximated by

$$\mu_{-i}^0 \pm (c_i/P)^{1/3} \lambda_i^{1/3} \partial_{CF}\left(\lambda_i^{1/3} (c_i/P)^{-2/3} / n_i^t\right). \tag{12}$$

Chick and Frazier (2012) provide a convenient approximation to $\partial_{CF}(s)$. For points above the upper or below the lower boundary it is optimal to stop and select arm i or $-i$, respectively. For points between the boundaries, it is optimal to continue sampling.

If the arms are assumed to be correlated, μ_{-i}^t is no longer constant, however. By the linearity of conditioning for joint Gaussian random vectors, μ_{-i}^t is a linear function of μ_{-i}^0, μ_{-i}^t and μ_i^t. Appendix A shows this fact can be used to characterize the optimal stopping boundaries for the correlated problem.

Proposition 1. *Suppose that $M' = 2$ and that samples are taken only from arm i. Then it is optimal to continue sampling when (μ_i^t, n_i^t) falls within the following region*

$$d_{(1)} + \mu_i^0 \pm \left((c_i/P)/(b_{(2)} - b_{(1)})\right)^{1/3} \lambda_i^{1/3} \partial_{CF}\left(\lambda_i^{1/3} \left((c_i/P)/(b_{(2)} - b_{(1)})\right)^{-2/3} / n_i^t\right). \tag{13}$$

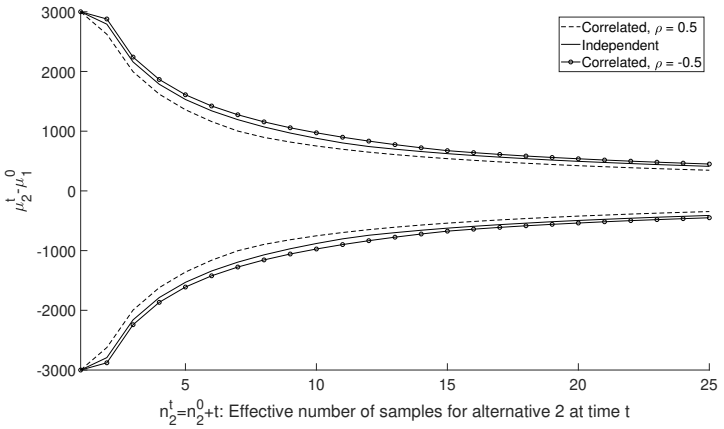

Figure 1: Optimal stopping boundaries for a problem with $M' = 2$, and samples only from the arm $i = 2$.

Comparing Eq. (12) and Eq. (13), we see how correlation between two arms impacts both the center and shape of the continuation region. As the difference between the prior means increases, the center of the region for correlated arms shifts up or down relative to the center for independent arms.

Figure 1 plots stopping boundaries for a sample problem with prior parameters $\mu_1^0 = \mu_2^0 = 0$ and $\Sigma_{1,1}^0 = \Sigma_{2,2}^0 = 10^6$. The figure assumes that $i = 2$, $c_j = 0.1$ and $n_j^0 = 1$ ($\lambda_j = 10^6$) for $j = 1,2$, and checks several values for the correlation $\rho = \eta^0 \in \{-0.5, 0, 0.5\}$. The horizontal axis depicts the effective number of samples from arm 2 at time t. The vertical axis shows the change in the posterior mean of arm 2 after t samples have been observed, relative to the prior mean of arm 1. The pair of solid lines depicts the stopping boundary when arms are independent. The pair of dashed lines represents the stopping boundary when the correlation between arms is positive. The pair with circles represents the boundary when the correlation is negative. Above each upper line of a pair and below each lower line, it is optimal to stop.

We observe in Figure 1 that the continuation region is wider when arms are negatively correlated. This is consistent with what we expect from theory when the prior variances are equal. When $b_{(2)} - b_{(1)} > 1$, the variance of the incremental benefit is larger than that of the independent case, and given greater variability in the benefit that will be obtained if stopped, it is optimal to sample more to reduce this variance. Thus, the continuation region is wider than that in the independent case. We observe $b_{(2)} - b_{(1)} > 1$ either when arms are negatively correlated or when arms are highly positively correlated and the prior variance of $-i$ is much smaller than the prior variance of i.

5 APPLICATION TO A CLINICAL TRIAL FOR RHINITIS TREATMENT

We now examine how different rule pairs perform when they guide allocation and stopping decisions in a Phase III trial that uses the factorial design described in Section 3.3. We generate experiments whose parameters are informed by data from clinical trials. The trials investigate the efficacy of combining azelastine hydrochloride (A), an antihistamine, and fluticasone propionate (F), a corticosteroid, to treat patients with moderate-to-severe seasonal allergic rhinitis (hay fever) and compare four arms: a placebo, A, F, and a combination of A and F. The first trial, labeled MP4001, is a proof-of-concept exploratory study. The second is a confirmatory trial, labeled MP4002. Thus, we apply our usage framework to an example motivated by real clinical trials (CADTH. 2016). We do not make medical recommendations.

5.1 Framework for Numerical Assessment with Rhinitis Treatment Data

We now introduce the experiments we use to test the allocation and stopping rules summarized in Section 3. We call the probability distribution from which the unknown mean vector, $\boldsymbol{\theta}$, is sampled Nature's distribution and denote it by $\mathcal{N}(\boldsymbol{\mu}(N), \boldsymbol{\Sigma}(N))$, differentiated by letter N. We refer to any particular $\boldsymbol{\theta}$ vector, which is a

realization of Nature's distribution, as a random problem instance. We denote by $\mathcal{N}\left(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right)$ the so-called "Modeler's" prior distribution, which guides allocation and stopping decisions. If the Modeler's belief is consistent with Nature's distribution then $\boldsymbol{\mu}^0 = \boldsymbol{\mu}(N)$ and $\boldsymbol{\Sigma}^0 = \boldsymbol{\Sigma}(N)$, but it need not be so in general.

We use the results of MP4002 to construct Nature's distribution. We use the results of the exploratory study, MP4001, to construct the Modeler's prior. The idea is that the Modeler's prior emerges from the information available from the preliminary study, although it may not be perfectly consistent with Nature's distribution. Thus, we test the performance of rule pairs when Nature's distribution is consistent with results from a published multi-arm Phase III trial, while the Modeler's prior variances are higher than Nature's variances, and therefore the prior is relatively non-informative. We generate many random problem instances, $\boldsymbol{\theta} \sim \mathcal{N}\left(\boldsymbol{\mu}(N), \boldsymbol{\Sigma}(N)\right)$, test our policies on each instance, and calculate each rule pair's average performance across the sampled $\boldsymbol{\theta}$'s.

We use Monte Carlo simulation to test the policies' performance. Each replication of the simulation is a sample path within which allocation and stopping decisions are based on the rule pairs and the Modeler's prior. To avoid a long run times, we slightly modify the stopping rules to stop either at period $t = 3000$ or when the stopping condition is satisfied, whichever comes first. No sample paths reached period 3000 in experiments reported here.

Three metrics are used to measure the performance of rule pairs: sampling cost, opportunity cost and total cost. The expected sampling cost is the sum of the sampling costs before stopping, $\mathbb{E}\left[\sum_{t=0}^{T-1} c_{u^t} \,\middle|\, \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right]$. The expected opportunity cost is $\mathbb{E}\left[\max_j \left\{P\theta_j\right\} \; - \; P\theta_{\mathscr{D}} \,|\, \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right]$ and measures the accuracy of the selection decision $\mathscr{D}$. The expected total cost is the sum of these two costs.

Expected values are estimated by calculating the average and standard error over 10^3 sample-path replications. We use common random numbers to ensure that results are roughly comparable across different allocation-stopping rule pairs and to reduce the standard error of pairwise comparisons. For notational simplicity in figures, we denote the sample average of simulation replications, which estimates the expectation, by E$(\cdot)$.

Since ESPb and ESPB assume that the arms are independent, any rule pair that includes ESPb or ESPB as a stopping or allocation rule starts with a version of the Modeler's prior in which arms are independent. Namely, we use the Modeler's variances, $\Sigma^0_{i,i}$ for all i, but set all covariances, $\Sigma^0_{i,j}$ for $i \neq j$, to zero. With all other rule pairs, we use the Modeler's prior as it is.

5.2 Parameter Values Motivated by Rhinitis Treatment Data

CADTH. (2016) reports hay-fever-specific quality of life scores for subjects in MP4001 and MP4002, which it labels RQLQs (Rhinoconjunctivitis Quality of Life Questionnaire scores). We assume that a one point decrease in the RQLQ score corresponds to a 0.05 increase in QALYs (Petersen et al. 2013) and that one QALY is valued at US$20,000. In this case, a one point reduction in the RQLQ score corresponds to a gain of US$1,000 per patient, and we present the results in U.S. dollar values.

We use the results of the trial MP4002 to construct Nature's distribution. The vector of sampling variances is calculated as $\boldsymbol{\lambda} = \left[1071^2, 1127^2, 1209^2, 1392^2\right]$, and the means of Nature's distribution are calculated as $\boldsymbol{\mu}(N) = [850, 1360, 1630, 1640]$. We assume that the results from MP4002 are a reasonable representation of reality, set the effective number of samples for Nature's distribution to the number of samples taken from each arm in MP4002, $n_i(N) = 832/4 = 208$ (the MP4002 trial had 207 to 210 patients on each arm), and calculate the variance of Nature's distribution for arm i as $\Sigma(N)_{i,i} = \lambda_i / n_i(N)$.

Using the approach described in Section 3.3, we first construct Nature's correlation matrix for $\boldsymbol{\xi}$, $\tilde{\boldsymbol{\rho}}(N)$. Since interventions A and F follow different mechanisms of action and appear to impair each other, we set $\eta(N) = 0$ and $\zeta(N) = -0.5$. Then, we use $\tilde{\boldsymbol{\rho}}(N)$ together with $\Sigma(N)_{i,i}$ for all i to construct the covariance matrix of the Nature's distribution, $\boldsymbol{\Sigma}(N)$.

Using the results of the trial MP4001, we set the Modeler's prior mean to be $\boldsymbol{\mu}^0 = [1010, 1170, 1430, 1580]$. To model the fact that the Modeler has less information as compared to Nature, we use MP4001's sample size of 607, assume $n_i^0 = 10^{-2} \times 607/4 \approx 1.51$, and set the prior variances as $\Sigma^0_{i,i} = \lambda_i / n_i^0$. We also assume

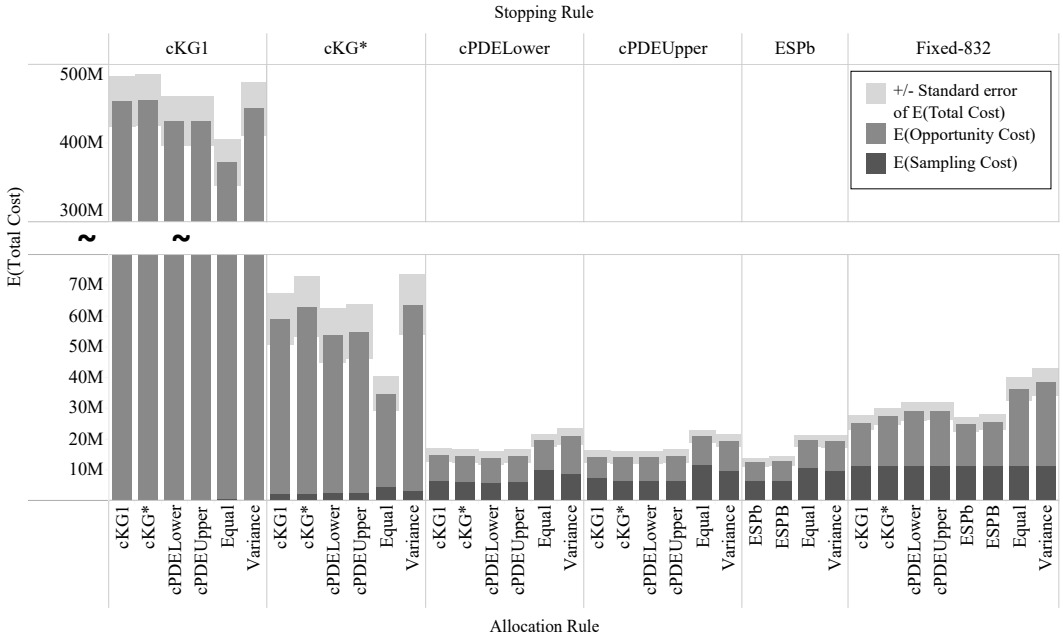

Figure 2: Average total, opportunity and sampling costs of each rule pair over 10^3 replications of MP4002.

that the Modeler correctly recognizes that interventions follow different mechanisms and that the Modeler's prior correlation values are consistent with Nature's correlation values, so $\eta^0 = 0$ and $\zeta^0 = -0.5$.

There are 60 million people in the US who suffer from hay fever (Settipane 2001), and we assume that $P = 6 \times 10^6$. For example, this figure corresponds to 10% of the population adopting the new intervention for one year or 1% of the population adopting the intervention for 10 years.

We use MP4002's sample size of 832 samples to back out the implicit sampling cost, the sampling cost that would have resulted in a sample size of 832 under the assumption that the study begins with the Modeler's prior and that the sample size is chosen to minimize the average total cost, and we find $c_i =$US\$13,500 for all i. If the actual sampling cost were lower than US\$13,500 per patient, this analysis would suggest that, given the Modeler's prior, there were too few subjects enrolled in the MP4002 trial.

5.3 Results Using Rhinitis Treatment Data

Figure 2 presents the sampling, opportunity and total costs of all rule pairs, including the Fixed stopping rule, which matches MP4002 and stops after 832 samples. The horizontal axis displays allocation-stopping rule pairs, which are grouped by stopping rule. Stopping rules appear on the upper horizontal axis, and the allocation rules paired with each stopping rule are shown on the lower horizontal axis of the graph. The dark-colored portions of the bars show the average sampling cost, the light-colored portions display the average opportunity cost, and the full bar lengths represent the average total cost. Light gray boxes at the tops of the bars range from plus one to minus one standard error of the average total cost.

First, we observe that the average total cost is lower for stopping rules based on dynamic stopping times: stopping rules cPDELower, cPDEUpper and ESPb outperform stopping rules that assume fixed stopping times, cKG1 and cKG*. In addition, the cKG* stopping rule, which tests multiple fixed stopping times, outperforms the cKG1 stopping rule, which tests only one. We note that the cPDELower, cPDEUpper and ESPb stopping rules sample the most and obtain the lowest average opportunity cost, while the cKG1 stopping rule undersamples and results in the highest average opportunity cost. The Fixed stopping rule obtains higher average opportunity and total costs and a lower sampling cost compared to cPDELower and cPDEUpper. Thus, given the implicit sampling cost, the total cost of MP4002 could have been lower if the sample size had been higher.

Second, we see that the choice of stopping rule impacts performance greatly, while the choice of allocation rule matters less. For any given stopping rule, there is no significant difference among the cKG1, cKG*, cPDELower and cPDEUpper allocation rules, all of which use varying methods to estimate the expected incremental value obtained from sampling. While the performance of the Equal and Variance allocation rules depends on the type of the stopping rule, generally these two rules result in the highest average sampling, opportunity and total costs. In particular, the Equal allocation rule does not prioritize arms that have the highest expected benefit of sampling and therefore samples more. Similarly, the Variance allocation rule prioritizes an arm only based on the relative lack of current information about its mean, the posterior variance, rather than on prospective information to be gained from sampling, and it appears to sample excessively from some arms to reduce their variances.

Third, in this experiment we do not detect a significant improvement coming from the incorporation of correlation among arms: the performance of all stopping and allocation rules that consider dynamic stopping times (cPDELower, cPDEUpper, ESPB and ESPb) are similar, regardless of whether or not arms are assumed to be independent. Nevertheless, we show that the incorporation of correlation provides a considerable benefit in the experiments in dose-response models (Chick et al. 2019).

We have also run an additional series of experiments to examine if the results above are robust to changes in the Modeler's and Nature's correlation matrices (data not shown). In one set of experiments, we assume that the Modeler's prior correlation is consistent with Nature's correlation values, and we vary $\eta(N)$ and $\zeta(N)$ to represent the three types relationships among interventions introduced in Section 3.3. In another set, we assume that the Modeler's prior correlation values are not consistent with Nature's correlation values, so η^0 and ζ^0 differ from $\eta(N)$ and $\zeta(N)$, and we again test each relationship type. In these experiments, the performance measures we track – average sampling, opportunity and total costs – are not impacted when the Modeler's prior correlation values are not fully consistent with Nature's, and the performance measures are robust to the different correlation-matrix structures we tested.

6 CONCLUSIONS

This paper demonstrates an alternative usage of the framework of Chick et al. (2019) for correlated mean responses in highly-sequential multi-arm clinical trials to optimize drug dosing. We apply those results to the context of factorial clinical trials, to test combinations of multiple treatments in a multi-arm trial for a given clinical condition. The approach also allows for the use of a value-based trial framework (Chick et al. 2017), which uses monetary and health benefit metrics to determine trial lengths as compared to the use of arbitrary statistical metrics.

We applied the framework to data from a factorial trial involving the presence and absence of each of two treatments for allergic rhinitis. The benefit of modeling such correlations in this combinatorial treatment setting was less significant than in the dose optimization experiment of Chick et al. (2019), as our linear model for treatment effects allowed for each independent interaction to be modeled individually. We hypothesize that benefits would be greater in the presence of additional structure regarding correlations (as in the case of dose-optimization trials) or when the largest mean effect of a therapy in the combination for two arms is greater than the mean size of the second order-interactions of that therapy with other, less-significant arms. However, our approach performed better in this example than the cKG approach, in that the new approach uses allocation indices for assigning treatments to patients which have a more complete assessment of the expected value of information of sampling. In simulations, this observation was relatively robust to mis-estimation of the correlation in the mean response across arms.

ACKNOWLEDGMENTS

This research was supported by the Wharton Risk Management Center's Russell Ackoff Doctoral Student Fellowship and the Wharton-INSEAD Alliance. Chick acknowledges the support of the European Union through the MSCA-ESA-ITN project (676129).

A PROOF OF PROPOSITION 1

In Section 4, we consider a problem that starts at time $t = 0$ with prior parameters $\left(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right)$ with two arms, $M' = 2$. We solve the problem here with population size $P = 1$, then adjust the solution for general P. We are sampling only from arm i, and we denote the arm we are not sampling from as $-i$.

First, we write the optimal stopping boundary for *independent* arms following Chick and Frazier (2012).

$$\text{EVI}_i^*(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0) \;=\; \sup_{\pi_i} \mathbb{E}_{\pi_i}\left[-c_i T_i \;+\; \max\left\{\mu_i^{T_i}, \mu_{-i}^{T_i}\right\} \;\middle|\; \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right] - \max\left\{\mu_i^0, \mu_{-i}^0\right\}. \tag{14}$$

Because we are not sampling from arm $-i$, $\mu_{-i}^t = \mu_{-i}^0$ for any $t > 0$. Thus, the reward from arm $-i$ can be thought of as a known standard. Then, the optimal stopping boundary in (μ_i^t, n_i^t) coordinates is

$$\mu_{-i}^0 \pm c_i^{1/3} \lambda_i^{1/3} \partial_{CF}(\lambda_i^{1/3} / (c_i^{2/3} n_i^t)). \tag{15}$$

Next, we allow arms to be correlated. We use Eq. (8) and Eq. (9) to write the problem in Eq. (7) as

$$\text{EVI}_i^*(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0) \;=\; \sup_{\pi_i} \mathbb{E}_{\pi_i}\left[-c_i T_i \;+\; a_{(1)} + b_{(1)} Z_i^{T_i} + (b_{(2)} - b_{(1)}) \left(-d_{(1)} + Z_i^{T_i}\right)^+ \;\middle|\; \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right] - a_{g(0)}. \tag{16}$$

Here, we start with the value of the arm with the lowest slope, $a_{(1)} + b_{(1)} Z_i^{T_i}$, and add the incremental value from the other alternative. Then, we rearrange terms to obtain

$$\begin{aligned} \text{EVI}_i^*(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0) \;&=\; \sup_{\pi_i} \mathbb{E}_{\pi_i}\left[-c_i T_i \;+\; (b_{(2)} - b_{(1)}) \left(-d_{(1)} + Z_i^{T_i}\right)^+ \;\middle|\; \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right] + a_{(1)} - a_{g(0)} && (17) \\ &=\; (b_{(2)} - b_{(1)}) \sup_{\pi_i} \mathbb{E}_{\pi_i}\left[-\frac{c_i}{b_{(2)} - b_{(1)}} T_i \;+\; \left(-d_{(1)} + Z_i^{T_i}\right)^+ \;\middle|\; \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right] + a_{(1)} - a_{g(0)}. && (18) \end{aligned}$$

We then define following variables: $\tilde{c}_i = \frac{c_i}{b_{(2)} - b_{(1)}}$, $m = 0$, $\tilde{\mu}_i^t = -d_{(1)} + Z_i^t$. Notice that, for any given t, $\tilde{\mu}_i^t \sim N(-d_{(1)}, \sigma_{Z_i^t}^2)$. Recall from Eq. (9) that $d_{(1)} = \frac{a_{(1)} - a_{(2)}}{b_{(2)} - b_{(1)}}$. The problem in these new variables is

$$\text{EVI}_i^*(\boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0) \;=\; (b_{(2)} - b_{(1)}) \sup_{\pi_i} \mathbb{E}_{\pi_i}\left[-\tilde{c}_i T_i \;+\; \max\left\{\tilde{\mu}_i^{T_i}, m\right\} \;\middle|\; \boldsymbol{\mu}^0, \boldsymbol{\Sigma}^0\right] + a_{(1)} - a_{g(0)}. \tag{19}$$

Using results from Chick and Frazier (2012), the stopping boundary in $(\tilde{\mu}_i^t, n_i^t)$ coordinates is then

$$0 \pm \left(c_i / (b_{(2)} - b_{(1)})\right)^{1/3} \lambda_i^{1/3} \partial_{CF}\left(\lambda_i^{1/3} [c_i / (b_{(2)} - b_{(1)})]^{-2/3} / n_i^t\right). \tag{20}$$

Notice that Eq. (15) is in (μ_i^t, n_i^t) coordinates, while Eq. (20) is in $(\tilde{\mu}_i^t, n_i^t)$ coordinates. To be able to compare the two boundaries, we need to convert them to the same coordinate system. By definition, $\tilde{\mu}_i^t = -d_{(1)} + Z_i^t = -d_{(1)} + \mu_i^t - \mu_i^0$. If we write the continuation region in Eq. (20) explicitly, we have

$$|-d_{(1)} + \mu_i^t - \mu_i^0| \le \left(c_i / (b_{(2)} - b_{(1)})\right)^{1/3} \lambda_i^{1/3} \partial_{CF}\left(\lambda_i^{1/3} [c_i / (b_{(2)} - b_{(1)})]^{-2/3} / n_i^t\right). \tag{21}$$

Add $d_{(1)} + \mu_i^0$ to both sides of (21) to rewrite the optimal stopping boundary in (μ_i^t, n_i^t) coordinates as

$$d_{(1)} + \mu_i^0 \pm \left(c_i / (b_{(2)} - b_{(1)})\right)^{1/3} \lambda_i^{1/3} \partial_{CF}\left(\lambda_i^{1/3} [c_i / (b_{(2)} - b_{(1)})]^{-2/3} / n_i^t\right). \tag{22}$$

Thus, the optimal stopping boundary is again symmetric, but the center is $d_{(1)} + \mu_i^0$.

Recall Equations (1) and (2). An optimal solution with sampling costs c_i/P and population size 1 is also optimal with sampling costs c_i and population size P. The result directly follows. □

REFERENCES

Berry, D. A. 2012. "Adaptive Clinical Trials in Oncology". *Nature Reviews Clinical Oncology* 9(4):199–207.

Boeree, M. J., N. Heinrich, R. Aarnoutse, A. H. Diacon, R. Dawson, S. Rehal, G. S. Kibiki, G. Churchyard, I. Sanne, N. E. Ntinginya et al. 2017. "High-dose Rifampicin, Moxifloxacin, and SQ109 for Treating Tuberculosis: A Multi-arm, Multi-stage Randomised Controlled Trial". *The Lancet Infectious Diseases* 17(1):39–49.

CADTH. 2016. "Common Drug Review Clinical Review Report for Dymista". Canadian Agency for Drugs and Technologies in Health, https://www.cadth.ca/sites/default/files/cdr/clinical/SR0408_Dymista_CL_Report.pdf, Accessed 16 Aug 2019.

Chick, S. E., M. Forster, and P. Pertile. 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, 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., and N. Gans. 2009. "Economic Analysis of Simulation Selection Problems". *Management Science* 55(3):421–437.

Chick, S. E., N. Gans, and O. Yapar. 2019. "Bayesian Sequential Learning for Clinical Trials of Multiple Correlated Medical Interventions". Working paper, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3184758, Accessed 16 Aug 2019.

Draper, D. 2013. "Discussion on 'Group sequential tests for delayed responses' (by L. Hampson and C. Jennison).". *Journal of the Royal Statistical Society: Series B* 75(1):48.

EMA 2007. "Reflection Paper on Methodological Issues in Confirmatory Clinical Trials Planned with an Adaptive Design". European Medicines Agency, http://www.ema.europa.eu/docs/en_GB/document_library/Scientific_guideline/2009/09/WC500003616.pdf, Accessed 16 Aug 2019.

EMA 2017. "Paediatric Gaucher disease: A Strategic Collaborative Approach from EMA and FDA". European Medicines Agency, http://www.ema.europa.eu/docs/en_GB/document_library/Scientific_guideline/2017/06/WC500230342.pdf, Accessed 16 Aug 2019.

FDA 2010. "Adaptive Design Clinical Trials for Drugs and Biologics". Rockwell MD: Food and Drug Administration, https://www.fda.gov/downloads/drugs/guidances/ucm201790.pdf, Accessed 7 May 2018.

FDA 2016. "Adaptive Designs for Medical Device Clinical Studies". Rockwell MD: Food and Drug Administration, https://www.fda.gov/downloads/medicaldevices/deviceregulationandguidance/guidancedocuments/ucm446729.pdf, Accessed May 7, 2018.

Frazier, P. I., W. Powell, and S. Dayanik. 2009. "The Knowledge-Gradient Policy for Correlated Normal Beliefs". *INFORMS Journal on Computing* 21(4):599–613.

Gold, M. R., J. E. Siegel, L. B. Russell, and M. C. Weinstein. 1996. *Cost-effectiveness in Health and Medicine*. Oxford: Oxford University Press.

HAS 2012. "Choices in Methods for Economic Evaluation". Haute Autorité de Santé, https://www.has-sante.fr/portail/jcms/r_1499251/en/choices-in-methods-for-economic-evaluation, Accessed 7 May 2018.

Jennison, C., and B. W. Turnbull. 1999. *Group Sequential Methods with Applications to Clinical Trials*. CRC Press.

Pallmann, P., A. W. Bedding, B. Choodari-Oskooei, M. Dimairo, L. Flight, L. V. Hampson, J. Holmes, A. P. Mander, L. Odondi, M. R. Sydes, S. S. Villar, J. M. Wason, C. J. Weir, G. M. Wheeler, C. Yap, and T. Jaki. 2018. "Adaptive Designs in Clinical Trials: Why Use Them, and How to Run and Report Them". *BMC Medicine* 16(29).

Petersen, K. D., C. Kronborg, J. N. Larsen, R. Dahl, and D. Gyrd-Hansen. 2013. "Patient Related Outcomes in a Real Life Prospective Follow up Study: Allergen Immunotherapy Increase Quality of Life and Reduce Sick Days". *World Allergy Organization Journal* 6(1):15.

Settipane, R. A. 2001. "Demographics and Epidemiology of Allergic and Nonallergic Rhinitis". *Allergy and Asthma Proceedings* 22(4):185–189.

UK 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, Accessed 7 May 2018.

AUTHOR BIOGRAPHIES

STEPHEN E. CHICK is Novartis Chair of Healthcare Management at INSEAD. He enjoys Bayesian statistics, math modeling for health applications, and simulation. His email address is stephen.chick@insead.edu.

NOAH GANS is a Professor at the Wharton School. He is interested in call center operations and enjoys stochastic models and applied probability. His email address is gans@wharton.upenn.edu.

OZGE YAPAR is a PhD Candidate at the Wharton School. She will join the ODT Department of the Indiana University in July 2019. She uses tools from applied probability, stochastic processes, and health economics. Her email address is oyapar@iu.edu.