

OPTIMAL STOPPING FOR CLINICAL TRIALS WITH ECONOMIC COSTS: A SIMULATION-BASED APPROACH

Amandeep Chanda¹, Michael Fu², Eric Slud^{1,3}

¹Department of Mathematics, University of Maryland, College Park, MD, USA

²Robert H. Smith School of Business and Institute for Systems Research, University of Maryland, College Park, MD, USA

³Center for Statistical Research and Methodology, US Census Bureau, Washington DC, USA

ABSTRACT

We consider the problem of designing an early stopping clinical trial investigating the efficacy of a medical intervention against the available standard of care. The standard approach is to determine a stopping rule minimizing the expected number of patients required, subject to error rate constraints, not considering costs explicitly depending on the magnitude of the treatment effect. In this paper we formulate an optimal stopping problem for clinical trials with instantaneous continuous response, the objective being minimizing an overall risk comprising loss functions accounting for costs involving the treatment effect that might model ethical and economic costs. To solve the optimization problem we propose a feasible directions simulation-based algorithm requiring new stochastic gradient estimators which we derive using Smoothed Perturbation Analysis. We conduct simulation experiments to test the effectiveness of the simulation optimization algorithm and to obtain insights on the effects of the various risk factors on the optimal solution.

1 INTRODUCTION

A randomized clinical trial serves as a litmus test to assess the safety and efficacy of medical interventions in human subjects. In its simplest form, a clinical trial investigates the efficacy of a new treatment by comparing it with the control (placebo, or the currently available standard treatment) in an experiment where participants are allocated to one of the two treatments according to a random mechanism. Owing to the nature of participant entry to a trial, the recorded observations, which are assumed to be quantitative and recorded instantaneously after administering the treatment, become available sequentially. When the accumulating data in a clinical trial provides evidence of the inferiority/superiority of the drug under investigation, should the trial be terminated, or should it continue as planned? If a trial is terminated ahead of its scheduled end, what is the ethical and financial fallout of an incorrect decision? Both of these questions are central to the design of a clinical trial. Consequently, clinical trials often include provisions to analyze the accumulating data at interim time-points - subjects entering the trial are divided into groups, and repeated significance testing of the accumulated data after processing each group of subjects is used to decide whether to stop or continue the trial. A clinical trial designed to include such repeated significance testing of the accumulating data at interim time-points is referred to as a group-sequential clinical trial, and the critical values at each analysis together constitute a group-sequential boundary.

We focus on designing a group-sequential randomized controlled clinical trial with the provision to stop early and declare efficacy (superiority of the investigational treatment). In addition to assuming the response recorded from each subject to be continuously distributed and instantaneously available on administration of a drug, we assume subjects enter the trial sequentially with no overlapping periods of followup. A group-sequential design is specified by the total number of analyses $K > 1$ (of which $K - 1$ are interim), the size of each group of subjects recruited between successive analyses (n_k , $k = 1, \dots, K$), and the stopping

boundaries (or critical values) at each analysis (b_k , $k = 1, \dots, K$). The total number of analyses K are fixed in advance, as are the group size n_k , $k = 1, \dots, K$, leaving the group-sequential boundary $\mathbf{b} = (b_1, \dots, b_K)$ to be determined. Additionally, the group-sequential boundary for repeated significance testing of the accumulating data must preserve the required false positive experiment-wise error rate.

Sequential analysis of data was first proposed by Wald (1945) during World War II. For a randomized clinical trial, however, analyzing the data in groups is more pragmatic. Elfring and Schultz (1973) and McPherson (1974) laid the foundations of group-sequential testing in clinical trials. Building on this groundwork, Pocock (1977) and O'Brien and Fleming (1979) provided the impetus required for its acceptance. Their group-sequential boundaries were easy to implement and could be extended to different types of responses, including survival-time responses. An examination of the operating characteristics of such boundaries was a natural stepping stone to the determination of optimal boundaries. A group-sequential boundary was considered to be optimal if it minimized the Expected Sample Size (*ESS*) - the average number of observations required to terminate the trial - at a specific difference in treatment effects with controlled error rates. Pocock (1982) directly compared multiple boundaries, thoroughly examining their *ESS* and statistical power at specific treatment effect differences. Jennison (1987) used numerical integration techniques to derive group-sequential boundaries which minimized the *ESS* weighted over multiple values of the treatment effect difference in the case of normally distributed responses with a known variance. Eales and Jennison (1992) improved upon this by deriving the optimal group-sequential boundary as the solution to an “unconstrained Bayes sequential decision problem”, minimizing an expected total cost involving a constant loss function for an incorrectly concluded trial using dynamic programming. However, none of these contributions consider the idea of incorporating the magnitude of the treatment effect into the loss function associated with an incorrect decision.

The first contribution of this paper is extending optimal boundaries to additionally minimize losses which depend explicitly on the magnitude of the treatment effect, thereby incorporating an interaction between said treatment effect and the terminating time of the trial. The fallout from an incorrectly concluded trial is both ethical and financial - prospective recipients are either supplied an ineffective drug, or continue on the available standard of care despite the existence of a superior drug. From the economic standpoint, the trial sponsor either fails to capture a market share and loses out on the possible revenue from a superior drug, or it must bear the burden of marketing an ineffective drug. Naturally, the resulting loss is dependent on the treatment effect difference. An optimal boundary is then defined to minimize an overall risk function - a linear combination of the risk due to subjects exposed to an “inferior” drug and the risk due to an incorrectly concluded trial, both weighted over multiple values of the treatment effect - subject to a constraint on the false positive error rate.

The second contribution of this paper is proposing an optimization algorithm to minimize a complicated overall risk function subject to a constraint on the false positive error rate. Using numerical integration to minimize an overall risk function involves the computation of high-dimensional integrals. This is cumbersome, and may require significant computing overhead. The proposed algorithm utilizes a feasible directions approach to generate a subsequence of improving group-sequential boundaries whilst maintaining the desired error rate up to a given “tolerance” limit. Additionally, we derive sample path-based estimators of the gradients of the objective and constraint functions with respect to the boundary points for use in the optimization algorithm, referring to the technique of Perturbation Analysis (PA) (Ho and Cao 1991; Glasserman 1991; Fu and Hu 1997), and in particular, Smoothed Perturbation Analysis (Gong and Ho 1987).

The remainder of this paper is organized as follows. In Section 2 we model the loss functions for an incorrectly concluded trial as a function of the difference of treatment effects, and introduce the overall risk function to be minimized subject to a false positive error rate constraint. In Section 3 we set up the SPA estimators of the gradients of the objective and constraint functions with respect to the boundary points. In Section 4 an adaptation of the feasible directions approach is used to develop an optimization

algorithm. Simulation results from implementing the optimization algorithm across different configurations are presented in Section 5 followed by concluding statements in Section 6.

2 PROBLEM DEFINITION

We consider a Phase III parallel two-arm trial to investigate the efficacy of an investigational drug (I) against the available standard of care (C) using two interim analyses and one final analysis at the conclusion of the trial. The basic sequence of events is as follows: assuming fully sequential entry of subjects into the trial without overlapping periods of followup, on entry each subject is randomly allocated to one of the two treatment arms, an immediate quantitative response is recorded for each subject upon administration of the drug, and the accumulating data is tested after each group of patients has been evaluated.

We define the following:

X_{Ij} : response of the j th subject allocated to treatment I ; $X_{Ij} \stackrel{\text{IID}}{\sim} \mathcal{N}(\theta_I, \sigma^2)$, independently of

X_{Cj} : response of the j th subject allocated to treatment C ; $X_{Cj} \stackrel{\text{IID}}{\sim} \mathcal{N}(\theta_C, \sigma^2)$,

$D_j = X_{Ij} - X_{Cj}$: difference in responses for the j th pair of subjects; $D_j \stackrel{\text{IID}}{\sim} \mathcal{N}(\theta, 2\sigma^2)$.

The normal mean $\theta = \theta_I - \theta_C$ is so parameterized that a positive value indicates superior efficacy of treatment I , and σ^2 is assumed to be known. At the k th analysis ($k = 1, 2, 3$) data on n_k pairs of subjects, one from each treatment arm, denoted by D_1, D_2, \dots, D_{n_k} , is available; under the assumption of equally sized groups: $n_k = nk$.

We wish to test,

$$H_0 : \theta \leq 0 \text{ vs. } H_1 : \theta > 0$$

Even though clinical trials are almost always done two-sided to enable early stopping for both superior and inferior efficacy of the investigational treatment, we consider a simplification where the trial stops early only for superiority. At the k th analysis, the test statistic $Z_k = \frac{\sum_{j=1}^{nk} D_j}{\sqrt{2\sigma^2 nk}}$ is compared to a stopping boundary, say b_k , and if $Z_k > b_k$, H_0 is rejected and the trial is terminated. Given the nature of the test statistic, we have the relation

$$Z_k = \sqrt{\frac{k-1}{k}} Z_{k-1} + \frac{\sum_{j=n(k-1)+1}^{nk} D_j}{\sqrt{2\sigma^2 nk}}; \quad k = 1, 2, 3$$

$Z_0 = 0$, and consequently, $\mathbf{Z} = (Z_1, Z_2, Z_3) \sim \mathcal{N}_3(\mu, \Sigma)$, where $\mu = (\mu_1, \mu_2, \mu_3)$, with $\mu_k = \sqrt{\frac{nk}{2\sigma^2}} \theta$ and $\Sigma = (\Sigma_{k_1 k_2})$ with $\Sigma_{k_1 k_2} = \sqrt{\frac{\min\{k_1, k_2\}}{\max\{k_1, k_2\}}}; 1 \leq k_1, k_2 \leq 3$.

The maximum number of observations required, and consequently the total number of subjects to be admitted to the trial, is decided according to power requirements or the availability of resources. Additionally, we desire our group-sequential stopping boundary $\mathbf{b} = (b_1, b_2, b_3)$ to have a pre-specified size (false positive error rate) at $\theta = 0$,

$$\mathbb{P}_{\theta=0}(Z_1 > b_1 \text{ or } Z_2 > b_2 \text{ or } Z_3 > b_3) = \alpha$$

Defining τ to be the random variable indicating the analysis at which a stopping rule is triggered, we have

$$\begin{aligned} \tau &= \mathbb{1}\{Z_1 > b_1\} + 2\mathbb{1}\{Z_1 \leq b_1, Z_2 > b_2\} + 3\mathbb{1}\{Z_1 \leq b_1, Z_2 \leq b_2\} \\ &= \mathbb{1}\{Z_1 > b_1\} + 2\mathbb{1}\{Z_1 \leq b_1, Z_2 > b_2\} + 3\mathbb{1}\{Z_1 \leq b_1, Z_2 \leq b_2, Z_3 > b_3\} \\ &\quad + 3\mathbb{1}\{Z_1 \leq b_1, Z_2 \leq b_2, Z_3 \leq b_3\} \\ &= \gamma_{R_1} + 2\gamma_{R_2} + 3\gamma_{R_3} + 3\gamma_{NR} \\ &= 3 - 2\gamma_{R_1} - \gamma_{R_2}, \end{aligned} \tag{1}$$

since $\gamma_{R_1} + \gamma_{R_2} + \gamma_{R_3} + \gamma_{NR} = 1$, where $\mathbb{1}\{\cdot\}$ denotes the indicator function, γ_{R_k} indicates termination of the trial with rejection of H_0 at the k th analysis, and γ_{NR} indicates termination of the trial with non-rejection of H_0 at the 3rd analysis. The Expected Sample Size $ESS_\theta = \mathbb{E}_\theta(n\tau)$ represents the number of observations required on average for a stopping rule to be triggered.

Optimal group-sequential boundaries were traditionally defined to minimize the ESS at specific θ , or weighted over multiple θ , subject to statistical size and power constraints. However, an incorrect decision, whether it is incorrectly rejecting or not rejecting H_0 , has far-reaching consequences on the prospective recipients of the drug, as well as financial repercussions for the trial sponsor. Denoting the rejection and non-rejection of H_0 , respectively, by R and NR , the loss functions associated with making an incorrect decision at the conclusion of the trial can be modeled as

$$\begin{aligned} L(\theta, R) &= \mathbb{1}\{\theta \leq 0\} \cdot g(\theta) \cdot \eta \cdot (N - n\tau) \cdot \mathbb{1}\{Z_\tau > b_\tau\} \\ &= \mathbb{1}\{\theta \leq 0\} \cdot g(\theta) \cdot \eta \cdot ((N - n)\gamma_{R_1} + (N - 2n)\gamma_{R_2} + (N - 3n)\gamma_{R_3}) \end{aligned} \quad (2)$$

$$\begin{aligned} L(\theta, NR) &= \mathbb{1}\{\theta > 0\} \cdot g(\theta) \cdot \eta \cdot (N - 3n) \cdot \mathbb{1}(Z_1 \leq b_1, Z_2 \leq b_2, Z_3 \leq b_3) \\ &= \mathbb{1}\{\theta > 0\} \cdot g(\theta) \cdot \eta \cdot (N - 3n) \cdot \gamma_{NR} \end{aligned} \quad (3)$$

where, $g(\theta) = \mathbb{1}\{\theta \leq 0\} + \mathbb{1}\{\theta > 0\}(1 + a\theta)^b$, $a, b > 0$, is the penalty for terminating the trial with an incorrect decision. Failure to detect a “superior” drug has a greater adverse effect on its probable recipients as well as the trial sponsor. This is reflected in the penalty function. N refers to the number of prospective consumers of the investigational drug, and $\eta \in (0, 1)$ a deflating factor to account for incomplete [future] uptake [of the new drug]. We note that the loss functions reflecting the economic impact associated with incorrectly concluded trials may be formulated differently depending on the therapeutic area, available drugs, and prevailing market conditions. Denoting $\pi(\cdot)$ to be a discrete prior measure on the continuous parameter space Θ , the overall risk associated with the group-sequential boundary \mathbf{b} is defined as a linear combination of the risks due to sampling duration and due to an incorrectly concluded trial:

$$\begin{aligned} R(\mathbf{b}) &= \lambda \sum_{\theta \in \Theta} \pi(\theta) \mathbb{E}_\theta(n\tau) + (1 - \lambda) \sum_{\theta \in \Theta} \pi(\theta) \mathbb{E}_\theta(L(\theta, R) + L(\theta, NR)) ; \quad \lambda \in [0, 1] \\ &= \lambda \sum_{\theta \in \Theta} \pi(\theta) n \left(3 - 2\mathbb{E}_\theta(\gamma_{R_1}) - \mathbb{E}_\theta(\gamma_{R_2}) \right) \\ &\quad + (1 - \lambda) \sum_{\theta \in \Theta: \theta \leq 0} \pi(\theta) \eta \left((N - n)\mathbb{E}_\theta(\gamma_{R_1}) + (N - 2n)\mathbb{E}_\theta(\gamma_{R_2}) + (N - 3n)\mathbb{E}_\theta(\gamma_{R_3}) \right) \\ &\quad + (1 - \lambda) \sum_{\theta \in \Theta: \theta > 0} \pi(\theta) (1 + a\theta)^b \eta (N - 3n) \left(1 - \mathbb{E}_\theta(\gamma_{R_1}) - \mathbb{E}_\theta(\gamma_{R_2}) - \mathbb{E}_\theta(\gamma_{R_3}) \right) \end{aligned} \quad (4)$$

where $\mathbb{E}_\theta(\cdot)$ indicates expectation at a specific θ . Reformulating the size of the group-sequential boundary \mathbf{b} as $S(\mathbf{b}) = \mathbb{E}_{\theta=0}(\gamma_{R_1} + \gamma_{R_2} + \gamma_{R_3})$, the optimization problem is given by:

$$\begin{aligned} \min_{\mathbf{b}} \quad & R(\mathbf{b}) \\ \text{subject to} \quad & S(\mathbf{b}) \leq \alpha \end{aligned}$$

In this particular setup, both $R(\cdot)$ and $S(\cdot)$ are continuous functions of the boundary \mathbf{b} owing to the continuous nature of the responses collected from the subjects, in addition to the opposite monotone dependence of the risk and the size as functions of the individual boundary points. Consequently, the optimal boundary is obtained at size exactly α , thereby reducing the constraint on the size to an equality.

$$\begin{aligned} \min_{\mathbf{b}} \quad & R(\mathbf{b}) \\ \text{subject to} \quad & S(\mathbf{b}) = \alpha \end{aligned}$$

Further, the analytical gradients of both the overall risk function and the size constraint are readily available due to the same reasons responsible for their continuity. As is evident from the expressions of $R(\cdot)$ and $S(\cdot)$, their gradients require computation of the derivatives $\frac{\partial \mathbb{E}_\theta(\gamma_{R_k})}{\partial b_i}; 1 \leq i \leq k \leq 3$, which in turn involve computing multivariate normal orthant probabilities. Sample path estimators of the derivatives obtained using perturbation analyses provide a work-around to such calculations which are particularly cumbersome in trials with a higher number of interim analyses. In the next section we derive the estimators of these derivatives using the technique of Smoothed Perturbation Analysis (SPA).

3 GRADIENT ESTIMATION

Considering our problem setting with normally distributed responses to the drugs with three group-sequential analyses, the true gradients of both $R(\cdot)$ and $S(\cdot)$ can be expressed in terms of the normal p.d.f $\phi(\cdot)$ and c.d.f $\Phi(\cdot)$, and hence readily computed. In this section we focus on using perturbation analysis to set up the estimators of the required derivatives, relegating the true gradients to Section 5 when we discuss the simulation results.

We illustrate deriving the estimator of $\frac{\partial \mathbb{E}_\theta(\gamma_{R_3})}{\partial b_1}$, including a brief explanation for exchanging differentiation and expectation; the estimators of the other derivatives follow accordingly. Considering the right hand limit only, by definition

$$\frac{\partial \mathbb{E}_\theta(\gamma_{R_3})}{\partial b_1} = \lim_{\Delta b_1 \rightarrow 0+} \frac{\mathbb{E}_\theta(\gamma_{R_3}(b_1 + \Delta b_1) - \gamma_{R_3}(b_1))}{\Delta b_1}$$

Deriving a sample path estimator of the above derivative using Infinitesimal Perturbation Analysis (IPA) is not feasible given the discrete nature of the indicator function.

Consequently, we refer to Smoothed Perturbation Analysis (SPA), an extension of IPA introduced by (Gong and Ho 1987; Fu and Hu 1997). Considering γ_{R_3} , the trial terminates with rejection of H_0 at the final analysis; if the test statistic Z_1 at the first analysis is at or below the nominal boundary point b_1 , there is no difference in the sequence of events between the nominal and the perturbed path. On the opposite end, if Z_1 is above the positively perturbed boundary point $b_1 + \Delta b_1$, the trial on both the nominal and the perturbed path is terminated at the first analysis with rejection of H_0 , again resulting in the same sequence of events. If, however, $b_1 < Z_1 \leq b_1 + \Delta b_1$, the trial on the nominal path deviates from the one on the perturbed path in that it rejects H_0 at the first analysis, resulting in:

$$\begin{aligned} \mathbb{E}_\theta(\gamma_{R_3}(b_1 + \Delta b_1) - \gamma_{R_3}(b_1)) &= \mathbb{E}_\theta(\mathbb{1}\{b_1 < Z_1 \leq b_1 + \Delta b_1\} \mathbb{1}\{Z_2 \leq b_2, Z_3 > b_3\}) \\ &= \mathbb{E}_\theta[\mathbb{1}\{Z_2 \leq b_2, Z_3 > b_3\} \mathbb{P}\{b_1 < Z_1 \leq b_1 + \Delta b_1 | Z_2, Z_3\}]. \end{aligned}$$

Assuming the interchange of integration and differentiation to be justified:

$$\begin{aligned} \frac{\partial \mathbb{E}_\theta(\gamma_{R_3})}{\partial b_1} &= \lim_{\Delta b_1 \rightarrow 0+} \frac{\mathbb{E}_\theta(\gamma_{R_3}(b_1 + \Delta b_1) - \gamma_{R_3}(b_1))}{\Delta b_1} \\ &= \lim_{\Delta b_1 \rightarrow 0+} \frac{\mathbb{E}_\theta[\mathbb{1}\{Z_2 \leq b_2, Z_3 > b_3\} \mathbb{P}\{b_1 < Z_1 \leq b_1 + \Delta b_1 | Z_2, Z_3\}]}{\Delta b_1} \\ &= \mathbb{E}_\theta \left[\mathbb{1}\{Z_2 \leq b_2, Z_3 > b_3\} \lim_{\Delta b_1 \rightarrow 0+} \frac{\mathbb{P}\{b_1 < Z_1 \leq b_1 + \Delta b_1 | Z_2, Z_3\}}{\Delta b_1} \right] \\ &= \mathbb{E}_\theta[\mathbb{1}\{Z_2 \leq b_2, Z_3 > b_3\} \phi_{Z_1|Z_2, Z_3}(b_1 | z_2, z_3)] \end{aligned}$$

leading to an unbiased estimator of the derivative:

$$\frac{\partial \mathbb{E}_\theta(\gamma_{R_3})}{\partial b_1} \hat{=} \mathbb{1}\{Z_2 \leq b_2, Z_3 > b_3\} \phi_{Z_1|Z_2, Z_3}(b_1 | z_2, z_3) \quad (5)$$

where, $\phi_{Z_1|Z_2,Z_3}(\cdot|z_2,z_3)$ is the conditional normal p.d.f of Z_1 given Z_2 and Z_3 . Considering the left hand limit, instead, will proceed in a similar manner resulting in the same estimator.

To justify the interchange of differentiation and expectation, it is sufficient for the conditional density $\phi_{Z_1|Z_2,Z_3}(\cdot|z_2,z_3)$ to be bounded above for all z_2 and z_3 . Noting that $\text{Var}(Z_1|Z_2,Z_3) = \sigma_{1|2,3}^2 = \frac{1}{2}$ we have for all $x \in \mathbb{R}$.

$$\phi_{Z_1|Z_2,Z_3}(x|z_2,z_3) < (2\pi\sigma_{1|2,3}^2)^{-\frac{1}{2}}$$

The sample path estimators for the remaining derivatives are obtained in a similar fashion.

$$\begin{aligned} \frac{\partial \mathbb{E}_\theta(\gamma_{R_1})}{\partial b_1} &\hat{=} -\phi_{Z_1}(b_1), \quad \frac{\partial \mathbb{E}_\theta(\gamma_{R_2})}{\partial b_1} \hat{=} \mathbb{1}\{Z_2 > b_2\}\phi_{Z_1|Z_2}(b_1|z_2); \\ \frac{\partial \mathbb{E}_\theta(\gamma_{R_2})}{\partial b_2} &\hat{=} -\mathbb{1}\{Z_1 \leq b_1\}\phi_{Z_2|Z_1}(b_2|z_1), \quad \frac{\partial \mathbb{E}_\theta(\gamma_{R_3})}{\partial b_2} \hat{=} \mathbb{1}\{Z_1 \leq b_1, Z_3 > b_3\}\phi_{Z_2|Z_1,Z_3}(b_2|z_1,z_3); \\ \frac{\partial \mathbb{E}_\theta(\gamma_{R_3})}{\partial b_3} &\hat{=} -\mathbb{1}\{Z_1 \leq b_1, Z_2 \leq b_2\}\phi_{Z_3|Z_1,Z_2}(b_3|z_1,z_2) \end{aligned}$$

The estimators of $\frac{\partial \mathbb{E}_\theta(\gamma_{NR})}{\partial b_i}; i = 1, 2, 3$; can be obtained using the relation $\gamma_{R_1} + \gamma_{R_2} + \gamma_{R_3} + \gamma_{NR} = 1$ and the above estimators. Finally, plugging in the appropriate estimators into the overall risk function and the size constraint provide the desired gradients.

4 OPTIMIZATION ALGORITHM

The proposed optimization algorithm is based on a tailored version of the classical feasible directions method (Zoutendijk 1960; Bashyam and Fu 1998) to solve a nonlinear program with an equality constraint. Given our problem setting, we require simulations to estimate only the desired gradients, whereas the objective function and the constraint can be computed directly if σ^2 , the prior p.m.f $\pi(\cdot)$ on the parameter space Θ , and the other constants are specified. In lieu of establishing its convergence properties, we exhaustively scrutinize the performance of the proposed optimization algorithm across different parameter configurations, including the number of Monte Carlo replications used to estimate the derivatives.

We first set up the required notation, along with a basic outline of the algorithm, leaving the specific details about its implementation for later in the section. All vectors, indicated by boldface lowercase letters, are column vectors with $\langle \mathbf{x} \rangle$ denoting a vector \mathbf{x} that has been normalized to have unit length. $\hat{g}_{ik}(\theta, \mathbf{b})$ denotes the SPA estimator of $\frac{\partial \mathbb{E}_\theta(\gamma_{R_k})}{\partial b_i}; 1 \leq i \leq k \leq 3$ computed using B Monte Carlo simulations. $\hat{\nabla}R(\mathbf{b})$ and $\hat{\nabla}S(\mathbf{b})$ respectively denote the estimators of the gradients of the overall risk function and the size constraint, computed at the group-sequential boundary \mathbf{b} , incorporating the derivative estimators. Let $D(\mathbf{b})$ refer to the normalized direction vector, and $\{a_n\}$ the sequence of step-sizes for the iterative updates common across all boundary points.

Each iteration of the algorithm follows

$$\mathbf{b}_{(n+1)} = \mathbf{b}_{(n)} + a_n D(\mathbf{b}_{(n)}) \quad (6)$$

where,

$$D(\mathbf{b}) = \begin{cases} -\langle \hat{\nabla}R(\mathbf{b}) \rangle & ; \text{ if } S(\mathbf{b}) \leq \alpha \\ \langle D_f(\mathbf{b}) \rangle & ; \text{ if } \alpha < S(\mathbf{b}) < \alpha_u \end{cases}$$

where $\alpha_u > \alpha$, and $D_f(\mathbf{b})$ denotes the vector of feasible directions computed at the group-sequential boundary \mathbf{b} . The optimization algorithm aims to generate a subsequence of group-sequential boundaries with decreasing overall risk and size in the interval $[\alpha, \alpha_u]$. The normalized directions vector and the common step-size are both central to achieving this. When $S(\mathbf{b}) < \alpha$, $D(\mathbf{b}) = -\langle \hat{\nabla}R(\mathbf{b}) \rangle$ represents a direction producing a concurrent reduction of $R(\cdot)$ and increase of $S(\cdot)$, the extent of the trade-off controlled

by the common step-size. The contrasting behavior of $R(\cdot)$ and $S(\cdot)$ is explained by their opposite direction of monotone dependence as functions of each individual boundary point. Within the interval $\mathcal{I} = (\alpha, \alpha_u)$ a move along $-\langle \hat{\nabla}R(\mathbf{b}) \rangle$ is not “feasible” for the same reason. Consequently, the vector of feasible directions $D(\mathbf{b}) = \langle \nabla D_f(\mathbf{b}) \rangle$ induces a shrinkage in the size, as well as a reduction in the overall risk for appropriate choice of constants in the linear program to determine it.

A classical approach (e.g., Luenberger 1973) to determine the vector $D_f = (d_1, d_2, d_3)$ of feasible directions is to solve the following linear program (LP):

$$\begin{aligned} \max_{v_0, d_1, d_2, d_3} \quad & v_0 \\ \text{subject to} \quad & \langle \hat{\nabla}R(\mathbf{b}) \rangle^T D_f \leq -k_R v_0 \\ & \langle \hat{\nabla}S(\mathbf{b}) \rangle^T D_f \leq -k_S v_0 \\ & -1 \leq d_1, d_2, d_3 \leq 1 \end{aligned} \tag{7}$$

where the constants k_R and k_S are positive. A feasible solution to the above LP is $d_1 = d_2 = d_3 = 0$ with $v_0 = 0$. Therefore, if it exists, optimization will provide a solution for which $v_0 > 0$. The constraints on the gradient vectors are so formulated that the solution represents a direction of concurrent decrease in the overall risk and the size. In addition to this, the constraints also induce a trade-off between the magnitudes of decrease in the two functions, controlled by the constants k_R and k_S - a smaller k_R to k_S ratio leads to a larger reduction in the size at the expense of a possibly smaller reduction in the overall risk. This can be used advantageously with larger reductions in the overall risk favored initially.

The common step-size a_n is determined at each iteration to produce the largest reduction in the overall risk without grossly violating the size constraint. When $S(\mathbf{b}_{(n-1)}) < \alpha$, the algorithm moves to decrease the overall risk and increase the size to the interior of \mathcal{I} . The step-size a_n is then chosen such that $S(\mathbf{b}_{(n)}) \lesssim \alpha_u$. This produces the largest possible decrease in the overall risk and ensures $S(\mathbf{b}_{(n)}) \in \mathcal{I}$. On the other hand, when $S(\mathbf{b}_{(n-1)}) \in \mathcal{I}$ moving along the feasible directions vector causes $S(\mathbf{b}_{(n)})$ to shrink towards α . The common step-size is then chosen to produce the largest possible reduction in the overall risk.

5 SIMULATION RESULTS

We consider the case of a clinical trial with two interim analyses and one final analysis. The response recorded for each subject is assumed to be normally distributed with known variance $\sigma^2 = 1$. The total number of observations $3n$ is determined from the number of observations required by a one-sided, single analysis, size (false positive error rate) $\alpha = 0.05$ Z-test to detect a minimal clinically relevant normal mean $\theta^* = 0.3$ with statistical power (1 – false negative error rate) $1 - \beta = 0.9$.

$$3n = \kappa \left(\frac{\Phi^{-1}(1 - \alpha) + \Phi^{-1}(1 - \beta)}{\frac{\theta^*}{\sqrt{2}\sigma}} \right)^2$$

where $\Phi^{-1}(\cdot)$ is the inverse of standard normal c.d.f, $\kappa = 1.3$ is an inflating factor to account for the additional observations required by a group-sequential test to detect θ^* with the same power as a test with no interim analyses. For the loss functions (2) and (3) we set the parameters in the penalty function $g(\cdot)$ at $a = 5$ and $b = 5$. We have found these values to appropriately penalize trials which fail to reject the null hypothesis even when the investigational drug is sufficiently “superior”. We choose the number of prospective recipients of the investigational drug $N = 1000$, and consider the deflating factor $\eta = 0.1$ to account for probable over-estimation of the market share of the drug. We consider two discrete uniform prior measures: $\pi_1(\cdot)$ - a non-informative uniform prior, and $\pi_2(\cdot)$ - a subjective prior.

$$\pi_1(\theta) = \begin{cases} 1/12 & ; \text{ if } \theta = -0.05, -0.001, 0, 0.001, 0.05, 0.1, 0.2, 0.25, 0.275, 0.3, 0.325, 0.35 \\ 0 & ; \text{ otherwise} \end{cases}$$

$$\pi_2(\theta) = \begin{cases} 0.4 & ; \text{ if } \theta = 0 \\ 0.2 & ; \text{ if } \theta = -0.001, 0.001 \\ 0.03 & ; \text{ if } \theta = 0.275, 0.300, 0.325 \\ 0.02 & ; \text{ if } \theta = -0.050, 0.050, 0.100, 0.200, 0.250 \\ 0.01 & ; \text{ if } \theta = 0.350 \end{cases}$$

Both prior measures place positive probability on multiple θ in the vicinity of 0. This is motivated by the fact that nine investigational drugs out of ten fail to see the light of day due to lack of sufficient evidence supporting their efficacy (Sun et al. 2022). We use two different linear combinations in (4) with $\lambda = 0.2$, and 0.7.

The estimates $\hat{g}_{ik}(\theta, \mathbf{b})$ of the derivatives $\frac{\partial \mathbb{E}_\theta(\gamma_k)}{\partial b_i}; 1 \leq i \leq k \leq 3$ are computed using $B = 10000$ simulations for each iteration of the optimization algorithm. We set $\mathcal{I} = (\alpha, \alpha_u) = (0.05, 0.0501)$ as a reasonable interval for the size of the optimal boundaries. Additionally, an iteration of the algorithm may require computation of the feasible directions vectors, and the common step-size. Solving the LP (7) for the feasible directions vector involves choosing the positive constants k_R and k_S . We set $k_R = 1$, and consider the three choices of k_S determined by $\frac{k_R}{k_S} = 0.5, 1$, and 2. Each choice of (k_R, k_S) leads to a decrease in $S(\cdot)$, and an appropriately chosen common step-size leads to a concurrent decrease in $R(\cdot)$. If such a common step-size exists, the pair which produces the largest reduction in $R(\cdot)$ is used in the LP. Essentially, at each iteration of the algorithm requiring computation of the vector of feasible directions, we choose the best out of three probable moves. If there does not exist a common step-size which results in a decrease in both the overall risk and the size, the common step-size is set to 0, and the SPA estimates are recomputed using B new Monte Carlo simulations. When the algorithm produces $m = 5$ such iterations consecutively, it is terminated, and the results are reported.

As stated previously in Section 3, given our simple problem setting the true derivatives $\frac{\partial \mathbb{E}_\theta(\gamma_k)}{\partial b_i}$ for $1 \leq i \leq k \leq 3$ can be computed explicitly, for example.

$$\frac{\partial \mathbb{E}_\theta(\gamma_{R_3})}{\partial b_1} = \phi_{Z_1}(b_1) \mathbb{P}_\theta(Z_2 \leq b_2, Z_3 > b_3 | Z_1 = b_1)$$

The results from operating the proposed optimization algorithm using the SPA estimators of the gradients are compared to those from using the true gradients. The optimization is initiated with the group-sequential boundary due to O'Brien and Fleming (1979), abbreviated OBF, henceforth. The results for the different configurations are reported in Table 1 and Figures 1 and 2: 'SPA' indicates the optimal boundary obtained by using the SPA estimators of the gradients in the optimization algorithm, while 'True' indicates the optimal boundary obtained by using the true gradients.

The results indicate the following:

- The optimal boundaries obtained by using the estimates of the gradients are very similar to those obtained by using the true gradients. The overall risk for these optimal boundaries are nearly identical.
- The OBF boundary has exactly the required size, whereas each of the optimal boundaries has size within the desired interval \mathcal{I} . In fact, each of the optimal boundaries can be further adjusted to have exact size $\alpha = 0.05$ at the expense of a minor increase in their overall risk.
- For each of the four configurations considered, the algorithm converges at nearly the same rate for either of the two gradients, although it does so more gradually when using the true gradient.

Table 1: Optimal group-sequential boundaries and corresponding overall risk $R(\cdot)$ and size $S(\cdot)$ under different linear combination of the risks and choices of the prior $\pi(\cdot)$. The boundary points and the risks have been rounded to the second decimal place, while the size has been rounded to the sixth decimal place.

		b_1	b_2	b_3	$R(\mathbf{b})$	$S(\mathbf{b})$
$\pi = \pi_1$	OBF	2.961	2.094	1.710	254.70	0.050000
	SPA	3.206	2.621	1.651	247.07	0.050096
	True	3.252	2.644	1.645	247.07	0.050092
$\pi = \pi_1$	OBF	2.961	2.094	1.710	228.93	0.050000
	SPA	2.436	2.217	1.719	226.59	0.050091
	True	2.443	2.213	1.718	226.59	0.050091
$\pi = \pi_2$	OBF	2.961	2.094	1.710	118.01	0.050000
	SPA	3.127	2.605	1.652	115.93	0.050092
	True	3.244	2.624	1.650	115.92	0.050094
$\pi = \pi_2$	OBF	2.961	2.094	1.710	192.96	0.050000
	SPA	2.316	2.155	1.753	191.99	0.050087
	True	2.317	2.159	1.752	191.99	0.050086

- For the two configurations with $\lambda = 0.7$, the OBF and optimal boundaries have similar values of the risk function. This is not unexpected because the OBF boundary implicitly minimizes the ESS . A relatively larger difference in the risk function is observed in the remaining two configurations with $\lambda = 0.2$, indicating that the OBF boundary is likely sub-optimal.

6 CONCLUSION

In this paper we have considered the problem of deriving group-sequential boundaries which are optimal in the sense of minimizing an overall risk function, the components of which depend explicitly on the magnitude of treatment effect differences. We have restricted ourselves to the setting of a clinical trial to determine the efficacy of an investigational treatment in comparison with the available standard of care, having three planned analyses of which two are interim. Further, the response of each patient to the treatment they receive is instantaneously available and assumed to be normally distributed with a known variance, and the unknown mean represents the treatment effect.

We have proposed an algorithm to solve the constrained optimization problem – minimizing an overall risk function subject to an equality constraint on the false positive error rate. In this simplified setting, the true gradients of the objective and constraint functions can be computed, and serve as the benchmark. Additionally, we have derived sample path-based estimators of the gradients using the technique of Smoothed Perturbation Analysis. In the case of trials with a higher number of analyses than we have considered here, these estimators provide an alternative to the high-dimensional integration required to directly compute the true gradients. The results of the algorithm employing the simulation-based estimated gradients compare favorably to those employing the true gradient.

The setting of this paper serves as a basis for considering more realistic clinical trials. This includes trials with more than one treatment regimen under consideration, a higher number of planned analyses, and/or time-to-event responses – the response for each patient is the time of the first observed occurrence of the event of interest measured from randomization. In these situations, the true gradients of the corresponding risks, and the risks themselves, do not have an analytical form that can be computed easily even in the low-dimensional setting. Consequently, a simulation-based approach for deriving the optimal group sequential boundary should be a promising avenue to pursue.




Figure 1: Convergence of the overall risk and the size when using the SPA gradients versus the true gradients for $\pi = \pi_1$ and the two choices of λ .




Figure 2: Convergence of the overall risk and the size when using the SPA gradients versus the true gradients for $\pi = \pi_2$ and the two choices of λ .

ACKNOWLEDGMENTS

This work was supported in part by the Air Force of Scientific Research (AFOSR) under Grant FA95502010211 and by the National Science Foundation (NSF) under Grant IIS-2123684.

REFERENCES

Bashyam, S. and M. C. Fu. 1998. "Optimization of (s, S) Inventory Systems with Random Lead Times and a Service Level Constraint". *Management Science* 44(12):S243–S256.

Eales, J. D., and C. Jennison. 1992. "An Improved Method for Deriving Optimal One-Sided Group Sequential Tests". *Biometrika* 79:13–24.

Elfring, G. L. and J. R. Schultz. 1973. "Group Sequential Designs for Clinical Trials". *Biometrics* 29:471–477.

Fu, M. C. and J. Q. Hu. 1997. *Conditional Monte Carlo: Gradient Estimation and Optimization Applications*. Boston, MA. Kluwer Academic Publishers.

Glasserman, P. 1991. *Gradient Estimation Via Perturbation Analysis*. Boston, MA. Kluwer Academic Publishers.

Gong, W. B. and Y. C. Ho. 1987. "Smoothed Perturbation Analysis of Discrete-Event Dynamic Systems". *IEEE Transactions on Automatic Control* 32(10):858–867.

Ho, Y.C. and Cao, X.R. 1991. *Perturbation Analysis of Discrete Event Dynamic Systems*. London. Kluwer Academic Publishers.

Jennison, C. 1987. "Efficient Group Sequential Tests with Unpredictable Group Sizes." *Biometrika* 74:155–165.

Luenberger, D. G. 1973. *Introduction to Linear and Nonlinear Programming*. Reading, MA. Addison-Wesley.

McPherson, K. 1974. "Statistics: The Problem of Examining Accumulating Data More Than Once". *New England Journal of Medicine* 290:501–502.

O'Brien, P. C. and Fleming, T. R. 1979. "A Multiple Testing Procedure for Clinical Trials". *Biometrics* 35:549–556.

Pocock, S. J. 1977. "Group Sequential Methods in the Design and Analysis of Clinical Trials". *Biometrika* 64:191–199.

Pocock, S. J. 1982. "Interim Analyses for Randomized Clinical Trials: The Group Sequential Approach". *Biometrics* 38:153–162.

Sun, D., W. Gao, H. Hu, and S. Zhou. 2022. "Why 90% of Clinical Drug Development Fails and How to Improve it?" *Acta Pharmaceutica Sinica. B* 12(7):3049–3062.

Wald, A. 1945. "Sequential Tests of Statistical Hypotheses". *Annals of Mathematical Statistics* 16:117–186.

Zoutendijk, G. 1960. *Methods of Feasible Directions: A Study in Linear and Non-Linear Programming*. Elsevier Publishing Company.

AUTHOR BIOGRAPHIES

AMANDEEP CHANDA is a Ph.D. Candidate in the Mathematical Statistics program at the Department of Mathematics, University of Maryland. His research interests include clinical trial design, stochastic gradient estimation, and simulation optimization. His email address is achanda@umd.edu.

MICHAEL C. FU holds the Smith Chair of Management Science in the Smith School of Business, University of Maryland, College Park. His research interests include simulation optimization and stochastic gradient estimation. He served as Program Chair for WSC 2011 and received the INFORMS Simulation Society Distinguished Service Award in 2018. He is a Fellow of INFORMS and IEEE. His email address is mfu@umd.edu, and his Web page is <https://www.rhsmith.umd.edu/directory/michael-fu>.

ERIC V. SLUD Eric Slud is a Professor at the University of Maryland, College Park, in the Statistics Program within the Mathematics Department, where he has been on the faculty since 1976. He is also the Area Chief for Mathematical Statistics in the Center for Statistical Research and Methodology at the US Census Bureau, since 2011. His research interests include Survival Analysis and Biostatistics, Sample-Survey Inference, and statistical inference for stochastic processes. He is a Fellow of the Institute for Mathematical Statistics and of the American Statistical Association. His email address is slud@umd.edu, and his Web page is <https://www.math.umd.edu/~slud/>.