

Group sequential testing of a treatment effect using a surrogate marker

Layla Parast^{id*} and Jay Bartroff

Department of Statistics and Data Sciences, The University of Texas at Austin, Austin, TX 78712, United States

*Corresponding author: Layla Parast, Department of Statistics and Data Sciences, The University of Texas at Austin, 2317 Speedway STOP D9800, Austin, TX, 78712, USA (parast@austin.utexas.edu).

ABSTRACT

The identification of surrogate markers is motivated by their potential to make decisions sooner about a treatment effect. However, few methods have been developed to actually use a surrogate marker to test for a treatment effect in a future study. Most existing methods consider combining surrogate marker and primary outcome information to test for a treatment effect, rely on fully parametric methods where strict parametric assumptions are made about the relationship between the surrogate and the outcome, and/or assume the surrogate marker is measured at only a single time point. Recent work has proposed a nonparametric test for a treatment effect using only surrogate marker information measured at a single time point by borrowing information learned from a prior study where both the surrogate and primary outcome were measured. In this paper, we utilize this nonparametric test and propose group sequential procedures that allow for early stopping of treatment effect testing in a setting where the surrogate marker is measured repeatedly over time. We derive the properties of the correlated surrogate-based nonparametric test statistics at multiple time points and compute stopping boundaries that allow for early stopping for a significant treatment effect, or for futility. We examine the performance of our proposed test using a simulation study and illustrate the method using data from two distinct AIDS clinical trials.

KEYWORDS: clinical trial; futility stopping; group sequential testing; stopping boundaries; surrogate marker.

1 INTRODUCTION

The ultimate promise of surrogate markers is that if they can be identified, then they can be used to make decisions about a treatment sooner. A surrogate marker is a measurement that replaces a primary outcome and is expected to predict the effect of a treatment (Temple, 1999; FDA, 2023). In studies where the primary outcome measurement necessitates long follow-up or is invasive or expensive, surrogate markers may lead to more timely decisions. Rigorous statistical methods have been developed to identify valid surrogate markers (Elliott, 2023). However, far fewer methods have been developed to actually use a surrogate marker to test for a treatment effect in a future study, which as mentioned above, is the ultimate goal. This is a particularly difficult problem to consider when the surrogate marker is not perfect e.g., it may not capture the entire treatment effect on the primary outcome. This concept is shown in Figure 1 such that there is a prior study, Study A, where the validity of the earlier (or less expensive) surrogate marker replacing the primary (longer term or more expensive) outcome has been examined. Interest lies in conducting a future study, Study B, and using only the surrogate marker information to test for a treatment effect.

Certainly, there is extensive previous work on utilizing a surrogate marker *in combination* with the primary outcome to test

for a treatment effect. For example, Li et al. (2022) propose an approach to use both surrogate marker and outcome information to adaptively calculate conditional power in a group sequential trial. When both the primary outcome and surrogate marker are time-to-event outcomes, Cook and Farewell (1996) and Lin (1991) propose a weighted global test statistic in a group sequential trial allowing for the evaluation of the utility of the potential surrogate endpoint. Under the assumption that the surrogate and outcome are bivariate normal, Tang et al. (1989) explicitly demonstrates the advantage of a group sequential test using both endpoints in terms of reduced needed sample size. Similarly, within a bivariate normal framework, Anderer et al. (2022) develop an approach to combine the outcome and surrogate in a Bayesian adaptive design context.

However, in many settings, the question is not how to combine the surrogate marker and the primary outcome. Instead, the aim is to understand how to test for a treatment effect with the surrogate marker measurements *only*, and thus avoid measuring the primary outcome. Methods that do address this question tend to either rely on fully parametric methods where strict parametric assumptions are made about the relationship between the surrogate and the outcome or assume the surrogate marker is measured at only a single time point. For example, Price et al.

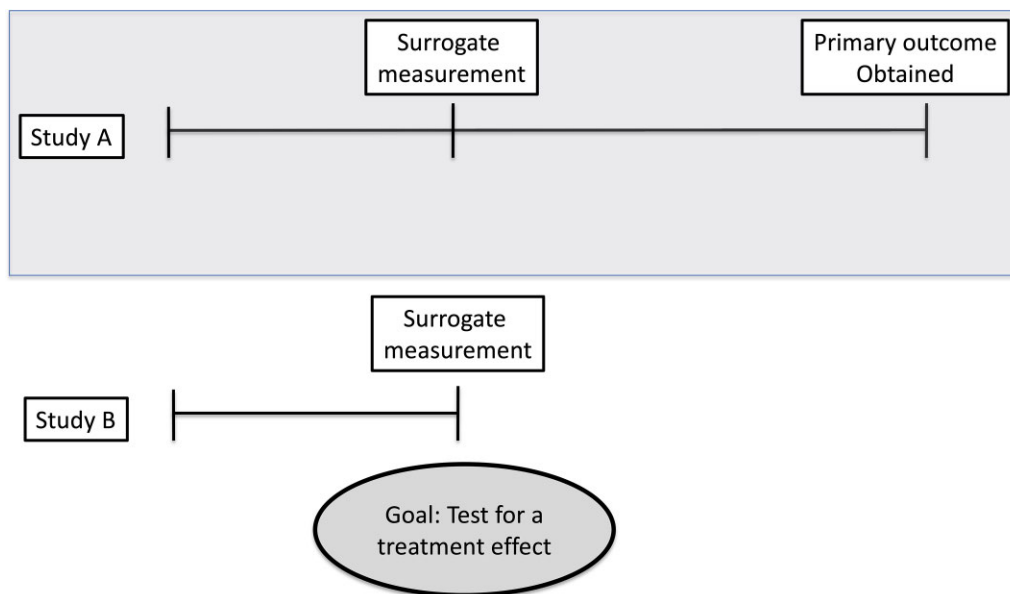


FIGURE 1 Studies A and B: Evaluating a surrogate in Study A and using the surrogate in Study B to test for a treatment effect.

(2018) propose to test for a treatment effect using a defined optimal surrogate at a single time point which aims to predict the primary outcome with estimation via the super-learner and targeted super-learner. Quan et al. (2023) and Saint-Hilary et al. (2019) proposed methods that use prior information about the treatment effect on a surrogate to plan a future study under the assumption that the true effects on the surrogate and primary outcome are bivariate normal.

Recent work (Parast et al., 2019, 2023) has proposed model-free procedures to test for a treatment effect using only surrogate marker information measured in the future study (Study B) by borrowing information learned from the prior study (Study A). Parast et al. (2023) proposed a kernel-based test statistic that is calculated using surrogate marker measurements from Study B obtained at a single, earlier time point. However, in practice, the surrogate marker is often measured repeatedly over time (eg, every 6 months) during the course of the study and thus, there is significant interest in applying sequential and group sequential testing methodology (Jennison and Turnbull, 2000; Bartroff et al., 2013) to the surrogate setting. In this paper, we build upon the kernel-based testing framework by proposing group sequential procedures that allow for early stopping to declare efficacy, and a version that also allows for the possibility of futility stopping, i.e., early stopping to declare failure to reject the null hypothesis of no effect. The use of the longitudinal surrogate in the test statistics prevents them from having the independent increment structure which simplifies the design of many group sequential procedures (see Spiessens et al., 2000; Kim and Tsiatis, 2020). Instead, we compute the correlation structure of the surrogate-based nonparametric test statistics and estimate it using the Study A data. The estimates of the correlation structure are then used to compute the group sequential procedures' stopping boundaries. We examine the performance of our procedures using a simulation study and illustrate the method using data from two distinct AIDS clinical trials.

2 SETTING, NOTATION, AND EXISTING APPROACH

2.1 Setting and notation

In Figure 2, we expand on Figure 1 to illustrate our setting of interest which has surrogate marker measurements over time. Let Y denote the continuous primary outcome measured at study completion in Study A and let S_j denote a continuous surrogate marker which is measured at multiple time points, t_j , $j = 1, \dots, J$, during the study. Without loss of generality, we assume $Y \geq 0$. Let G denote the treatment indicator where treatment is randomized and $G \in \{0, 1\}$ (ie, treatment vs. control). Our aim is to use information learned in Study A about the relationship between the primary outcome, the surrogate marker measurements, and the treatment to test for a treatment effect at an earlier time point in Study B, such that the duration of follow-up needed for Study B can be shortened and the primary outcome does not have to be measured in Study B. We use a subscript L to explicitly denote the study and use potential outcomes notation where each person in Study L has a potential $\{Y_L^{(1)}, Y_L^{(0)}, S_{jL}^{(1)}, S_{jL}^{(0)}\}$ where $Y_L^{(1)}$ is the outcome under treatment, $Y_L^{(0)}$ is the outcome under control, $S_{jL}^{(1)}$ is the surrogate at time t_j under treatment, and $S_{jL}^{(0)}$ is the surrogate at time t_j under control in Study L . For individual i in treatment group g in Study L , the *observed* surrogate at time t_j is S_{Lgji} . For individual i in treatment group g in Study A, the *observed* outcome will be denoted as Y_{Agi} ; the outcome in Study B is never measured. Let n_{Lg} denote the number of individuals in treatment group g in Study L , and $n_L = n_{L0} + n_{L1}$ where $n_{Lg}/n_L > 0$. When feasible, for ease, we will drop the subscript L from notation.

Our primary goal is to test for a treatment effect on the primary outcome in Study B quantified as $\Delta \equiv E(Y_B^{(1)} - Y_B^{(0)}) = 0$, *without* measuring the primary outcome, Y , in Study B. In the following sections, we describe an existing kernel-based test for

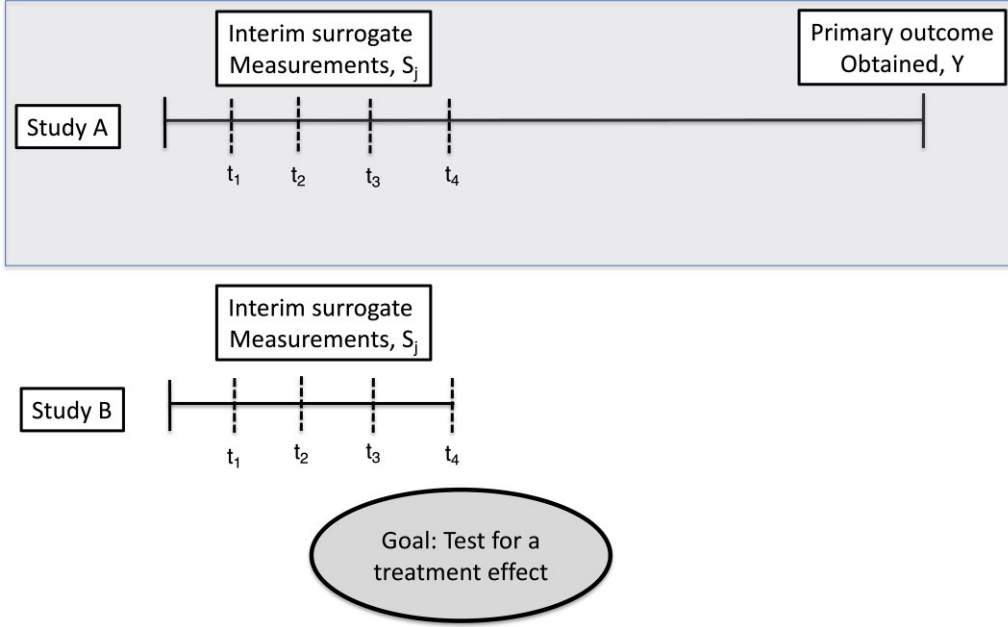


FIGURE 2 Studies A and B: Sequential surrogate setting.

a treatment effect in Study B using a surrogate measured at only a single time t_j (Parast et al., 2023) and possible naive approaches to sequential testing (Section 2.2), and then propose a novel approach to sequentially test for a treatment effect in Study B, when the surrogate is measured repeatedly over time (Section 3). In Section 4, we detail our needed assumptions as well as methods to empirically examine these assumptions.

2.2 Existing approach

Suppose there is only a single time point t_j , i.e., that the surrogate marker is measured at only one time point after randomization, but *before* the end of the study. Here, we drop the subscript j from S as it is unnecessary, and denote the potential surrogate as $S_L^{(g)}$ and the observed surrogate for person i as S_{Lgi} , $L = A, B$ and $g = 0, 1$. The goal is to take advantage of information from Study A to test $\Delta = 0$ using the surrogate marker measurement only from Study B. To achieve this goal we first consider using the testing procedure of Parast et al. (2023), described below. Note that Δ can be expressed as

$$\begin{aligned} \Delta &= E\left(Y_B^{(1)}\right) - E\left(Y_B^{(0)}\right) \\ &= \int \mu_{B1}(s) dF_{B1}(s) - \int \mu_{B0}(s) dF_{B0}(s), \end{aligned}$$

where $\mu_{Lg}(s) = E(Y_L^{(g)} | S_L^{(g)} = s)$ and $F_{Bg}(s)$ is the cumulative distribution function of $S_B^{(g)}$. Of course, this expression involves $Y_B^{(g)}$, which we do not observe. Parast et al. (2023) suggest to instead focus on

$$\Delta^* \equiv \int \mu_{A0}(s) dF_{B1}(s) - \int \mu_{A0}(s) dF_{B0}(s) \quad (1)$$

and they refer to Δ^* as an *earlier* treatment effect because it is defined before the end of Study B using only (1) the conditional mean, $\mu_{A0}(s)$, from Study A, and (2) the surrogate marker in

Study B i.e., $Y_B^{(g)}$ does not appear in Δ^* . The motivation for the exact construction of Δ^* is as follows. The troublesome components of Δ , in terms of trying to examine the treatment effect earlier, are $\mu_{B0}(s)$ and $\mu_{B1}(s)$ because these involve $Y_B^{(g)}$, which is not measure until the end of Study B. Suppose, we simply decided to replace these components with their parallel components from Study A, i.e., define Δ^* as $\int \mu_{A1}(s) dF_{B1}(s) - \int \mu_{A0}(s) dF_{B0}(s)$. It is true that this construction could still be referred to as an earlier treatment effect because it is defined without $Y_B^{(g)}$. However, there are two problems with this construction. First, we would have to assume that $\mu_{B0}(s) = \mu_{A0}(s)$ and $\mu_{B1}(s) = \mu_{A1}(s)$; i.e., that the conditional mean in each treatment group is transportable between the two studies. Second, it can be shown that unless S is a perfect surrogate marker, it is possible that $\Delta^* > \Delta$ and in an extreme case where S is a poor surrogate, Δ^* may erroneously indicate a treatment effect when in fact, $\Delta = 0$. It is this second problem that is particularly worrisome. It is never the case in practice that we have a perfect surrogate, and the risk of erroneously concluding there is a treatment effect using the surrogate when there is not a treatment effect on the outcome is not a risk that we are willing to take. Therefore, we specifically want a Δ^* that can give us some guarantee, e.g., that Δ^* is a lower bound on Δ . The definition in (1) does guarantee this property. In Web Appendix A, we show that under our assumptions (see Section 4), $\Delta^* \leq \Delta$ and that when $\Delta = 0 \Rightarrow \Delta^* = 0$. In addition, we show that a valid test of the null hypothesis that $\Delta^* = 0$ would also result in a valid test of the null hypothesis that $\Delta = 0$ in the sense that the type 1 error rate is $\leq \alpha$. Thus, to test the null hypothesis that $\Delta = 0$, one could instead consider testing the null hypothesis that $\Delta^* = 0$. Notably, Δ^* still requires a form of transportability between the two studies: $\mu_{B0}(s) = \mu_{A0}(s)$ (see Section 4), though this is less strong than requiring this equality in *both* treatment groups.

In practice, $\mu_{A0}(s)$ is unknown and furthermore, Study A data are fixed in this testing framework. By “fixed”, we mean that all of our probability calculations will condition on Study A data and n_A does not $\rightarrow \infty$. This is an important point. We specifically address the situation of designing Study B *after* Study A has concluded, and thus, an important statistical assumption made throughout this paper is that the Study A data is conditioned on, and hence treated as fixed quantities in the probabilistic calculations. Therefore, this earlier treatment effect must be defined in a way that acknowledges and makes explicit this reliance on Study A; specifically, we define

$$\Delta_E \equiv \int \widehat{\mu}_{A0}(s) dF_{B1}(s) - \int \widehat{\mu}_{A0}(s) dF_{B0}(s), \text{ where } (2)$$

$$\widehat{\mu}_{A0}(s) = \frac{\sum_{i=1}^{n_{A0}} K_h(S_{A0i} - s) Y_{A0i}}{\sum_{i=1}^{n_{A0}} K_h(S_{A0i} - s)} (3)$$

is a consistent estimate of $\mu_{A0}(s)$, and Δ_E is estimated by

$$\widehat{\Delta}_E = n_{B1}^{-1} \sum_{i=1}^{n_{B1}} \widehat{\mu}_{A0}(S_{B1i}) - n_{B0}^{-1} \sum_{i=1}^{n_{B0}} \widehat{\mu}_{A0}(S_{B0i}). (4)$$

Here, $K_h(\cdot) = K(\cdot/h)/h$ where $K(\cdot)$ is a smooth symmetric density function with finite support (eg, standard normal density) and h is a specified bandwidth, which may be data dependent. Note that this estimate only uses S data from Study B (no Y data from Study B is used or assumed to be measured) and S, Y data from Study A in group $Z = 0$ only. To test the null hypothesis, $\Delta = 0$, Parast et al. (2023) suggest testing the null hypothesis: $\Delta_E = 0$, (based on the properties above regarding Δ^*) using the Wald-type test statistic $W_E = \sqrt{n_B} \widehat{\Delta}_E / \widehat{\sigma}_E$. The test rejects the null hypothesis when $|W_E| > \Phi^{-1}(1 - \alpha/2)$. In Web Appendix B, we provide the form of $\widehat{\sigma}_E$ and discuss the asymptotic properties of $\widehat{\Delta}_E$.

To use this existing test in our setting in Figure 2, one could consider testing $\Delta_E = 0$ using W_E computed at each time t_j , $j = 1, \dots, J$, but this would have inflated Type 1 error probability. To remedy this, one could utilize a Bonferroni correction where the null hypothesis is rejected at the earliest t_j such that $|W_E| > \Phi^{-1}(1 - \alpha/(2J))$. However, one would expect this to be overly conservative, i.e., have low power. In the next section, we propose a compromise between these two approaches via group sequential procedures.

3 GROUP SEQUENTIAL PROCEDURES

Here, we propose a group sequential procedure that allows for early testing of a treatment effect, derive the properties of the correlated surrogate-based nonparametric test statistics at multiple time points, and develop stopping boundaries that allow for early stopping for a significant treatment effect or for futility. We assume that in Study A, we have measured the surrogate marker at time points t_j , $j = 1, \dots, J$, where $J \geq 2$ is the maximum number of interim analyses, and the same time points will be utilized for measuring the surrogate marker in Study B. Throughout this section, let $[J] := \{1, 2, \dots, J\}$.

3.1 Group sequential statistic

We will utilize the estimators (3)–(4) but repeatedly calculated with data from each time point t_j . Thus, we first adjust the notation to reflect this. For $j \in [J]$, recall that S_{Lgi} is the Study L ($L = A, B$) surrogate for the i th patient in group g ($g = 0, 1$) measured at time t_j and Y_{Agi} is the Study A outcome for the i th patient in group g ($g = 0, 1$). For each t_j , for any s within the support of S_{A0ij} , we define the kernel estimator using Study A data from time t_j :

$$\widehat{\mu}_{A0j}(s) = \frac{\sum_{i=1}^{n_{A0}} K_h(S_{A0ij} - s) Y_{A0i}}{\sum_{i=1}^{n_{A0}} K_h(S_{A0ij} - s)}, \text{ and let } (5)$$

$$\widehat{\Delta}_{Ej} = \frac{1}{n_{B1}} \sum_{i=1}^{n_{B1}} \widehat{\mu}_{A0j}(S_{B1ij}) - \frac{1}{n_{B0}} \sum_{i=1}^{n_{B0}} \widehat{\mu}_{A0j}(S_{B0ij}) (6)$$

be the estimator of Δ_E at time t_j . Let $\Delta_E = (\Delta_{E1}, \Delta_{E2}, \dots, \Delta_{EJ})^T$ where

$$\Delta_{Ej} = \int \widehat{\mu}_{A0j}(s) dF_{B1j}(s) - \int \widehat{\mu}_{A0j}(s) dF_{B0j}(s),$$

and $F_{Bgj}(s)$ is the cumulative distribution function of S_{Bgj} . To reflect the group sequential nature of this set up, we replace the null hypothesis, $\Delta_E = 0$ by

$$H_E : \Delta_E = \mathbf{0}, (7)$$

the latter denoting the J -dimensional zero vector. We test (7) via group sequential monitoring of the longitudinal process $\widehat{\Delta}_E = (\widehat{\Delta}_{E1}, \widehat{\Delta}_{E2}, \dots, \widehat{\Delta}_{EJ})^T$.

3.2 Covariance and limiting distribution

Conditional on $\{\widehat{\mu}_{A0j}(s)\}_{j \in [J]}$, it can be shown that as $n_{B0}, n_{B1} \rightarrow \infty$, the suitably-normalized distribution of $(\widehat{\Delta}_E - \Delta_E)$ approaches a J -dimensional multivariate normal random vector with mean zero (see Web Appendix B), and we recall that Δ_E is the zero vector under the null hypothesis (7). For $j, j' \in [J]$, the covariance matrix of $\widehat{\Delta}_E$ has elements

$$\begin{aligned} \sigma_{jj'} &:= \text{cov}(\widehat{\Delta}_{Ej}, \widehat{\Delta}_{Ej'}) \\ &= \text{cov} \left(\frac{1}{n_{B1}} \sum_{i=1}^{n_{B1}} \widehat{\mu}_{A0j}(S_{B1ij}) - \frac{1}{n_{B0}} \sum_{i=1}^{n_{B0}} \widehat{\mu}_{A0j}(S_{B0ij}), \right. \\ &\quad \left. \frac{1}{n_{B1}} \sum_{i=1}^{n_{B1}} \widehat{\mu}_{A0j'}(S_{B1ij'}) - \frac{1}{n_{B0}} \sum_{i=1}^{n_{B0}} \widehat{\mu}_{A0j'}(S_{B0ij'}) \right) (8) \end{aligned}$$

$$\begin{aligned} &= \frac{1}{n_{B1}^2} \text{cov} \left(\sum_{i=1}^{n_{B1}} \widehat{\mu}_{A0j}(S_{B1ij}), \sum_{i=1}^{n_{B1}} \widehat{\mu}_{A0j'}(S_{B1ij'}) \right) \\ &\quad + \frac{1}{n_{B0}^2} \text{cov} \left(\sum_{i=1}^{n_{B0}} \widehat{\mu}_{A0j}(S_{B0ij}), \sum_{i=1}^{n_{B0}} \widehat{\mu}_{A0j'}(S_{B0ij'}) \right) (9) \end{aligned}$$

$$\begin{aligned}
 &= \frac{1}{n_{B1}^2} \sum_{i=1}^{n_{B1}} \text{cov}(\widehat{\mu}_{A0j}(S_{B1ij}), \widehat{\mu}_{A0j'}(S_{B1ij'})) \\
 &\quad + \frac{1}{n_{B0}^2} \sum_{i=1}^{n_{B0}} \text{cov}(\widehat{\mu}_{A0j}(S_{B0ij}), \widehat{\mu}_{A0j'}(S_{B0ij'})) \\
 &= \frac{1}{n_{B1}} \text{cov}(\widehat{\mu}_{A0j}(S_{B11j}), \widehat{\mu}_{A0j'}(S_{B11j'})) \\
 &\quad + \frac{1}{n_{B0}} \text{cov}(\widehat{\mu}_{A0j}(S_{B01j}), \widehat{\mu}_{A0j'}(S_{B01j'})), \quad (10)
 \end{aligned}$$

where terms involving covariances between S_{B1ij} and $S_{B0ij'}$, which are equal to zero because $S_{B1ij} \perp S_{B0ij'}$ as they are from different treatment groups, are dropped from (8) to (9). Importantly, (10) relies on Study A being conditioned on, as mentioned in Section 2.2 and further discussed in Section 7. Taking $j = j'$ in (10) gives the variance

$$\begin{aligned}
 \sigma_{jj} &= \text{Var}(\widehat{\Delta}_{Ej}) \\
 &= \frac{1}{n_{B1}} \text{Var}(\widehat{\mu}_{A0j}(S_{B11j})) + \frac{1}{n_{B0}} \text{Var}(\widehat{\mu}_{A0j}(S_{B01j})). \quad (11)
 \end{aligned}$$

Let \mathbf{D} be the $J \times J$ diagonal matrix with entries $\sigma_{11}^{-1/2}, \sigma_{22}^{-1/2}, \dots, \sigma_{JJ}^{-1/2}$, and let $\boldsymbol{\Sigma} = \text{cov}(\mathbf{D}\widehat{\boldsymbol{\Delta}}_E)$ which is the $J \times J$ covariance matrix with (j, j') entry equal to $\sigma_{jj'}/\sqrt{\sigma_{jj}\sigma_{j'j'}}$ and, consequently, diagonal entries equal to 1. The null case of the multivariate normal limit mentioned above is that, under (7),

$$\boldsymbol{\Sigma}^{-1/2} \mathbf{D}\widehat{\boldsymbol{\Delta}}_E \rightarrow N_J(\mathbf{0}, \mathbf{I}) \quad (12)$$

as the $n_{Bg} \rightarrow \infty$, where \mathbf{I} is the $J \times J$ identity matrix. Although $\widehat{\boldsymbol{\Delta}}_E$ has the multivariate normal limit (12), it does *not* have the independent increment structure (Spiessens et al., 2000; Kim and Tsiatis, 2020) often exploited to simplify group sequential analyses since, for $j < j'$, the distribution of $S_{Bg1j'}$ differs in general from that of S_{Bg1j} , thus so do the distributions of $\widehat{\mu}_{A0j}(S_{Bg1j'})$ and $\widehat{\mu}_{A0j}(S_{Bg1j})$, thus $\text{cov}(\widehat{\mu}_{A0j}(S_{Bg1j}), \widehat{\mu}_{A0j'}(S_{Bg1j'})) \neq \text{Var}(\widehat{\mu}_{A0j}(S_{Bg1j}))$ in general, and so $\text{cov}(\widehat{\Delta}_{Ej}, \widehat{\Delta}_{Ej'}) \neq \text{Var}(\widehat{\Delta}_{Ej})$. In spite of this, in Section 3.3, we provide a way to compute stopping boundaries for $\widehat{\boldsymbol{\Delta}}_E$ utilizing the limiting normal distribution and the covariance structure (10), which can be implemented by Monte Carlo simulation.

3.3 Group sequential tests without futility stopping

Replacing the variances in (11) by the corresponding sample variances, for $j \in [J]$ define the variance estimator $\widehat{\sigma}_{jj} = \sum_{g=0}^1 \frac{1}{n_{Bg}} \left\{ \frac{1}{n_{Bg}} \sum_{i=1}^{n_{Bg}} (\widehat{\mu}_{A0j}(S_{Bgi j}))^2 - \left(\frac{1}{n_{Bg}} \sum_{i=1}^{n_{Bg}} \widehat{\mu}_{A0j}(S_{Bgi j}) \right)^2 \right\}$ of σ_{jj} . We consider the group sequential procedure that stops and rejects the null hypothesis (7) at the first analysis $j \in [J]$ such that

$$|W_{Ej}| \geq b_j, \quad \text{where } W_{Ej} := \widehat{\Delta}_{Ej}/\sqrt{\widehat{\sigma}_{jj}}, \quad (13)$$

and otherwise terminates at the J th analysis and declares failure to reject the null hypothesis. The statistics $W_{E1}, W_{E2}, \dots, W_{EJ}$ extend the univariate statistic, W_E , to the group sequential

setting. In (13), the b_j are pre-determined stopping boundaries; in general the b_j will also depend on J but we suppress this to simplify the notation. This test structure does not permit stopping before the J th analysis to declare failure to reject the null hypothesis, known as futility stopping, and in Section 3.4, we discuss a modification of this procedure that allows futility stopping. To compute the stopping boundaries b_j to achieve approximately a prescribed type I error probability α , we consider a set up that includes the popular Pocock (1977), O'Brien-Fleming (1979), and Wang-Tsiatis (1987) "power family" type boundaries, although a similar approach can be used for other families of group sequential boundaries (see Jennison and Turnbull, 2000). For these families of stopping boundaries, the b_j in (13) are of the form

$$b_j = b \cdot \widetilde{b}_j, \quad (14)$$

where \widetilde{b}_j is a known function of j (and J) and b is a constant determined to satisfy the type I error probability constraint. For example, for $j \in [J]$, the Pocock boundaries use $\widetilde{b}_j = 1$, the O'Brien-Fleming boundaries use $\widetilde{b}_j = \sqrt{J/j}$, and the Wang-Tsiatis boundaries use $\widetilde{b}_j = (j/J)^{\delta-1/2}$ for a chosen value of the power parameter δ . The latter includes the Pocock ($\delta = 1/2$) and O'Brien-Fleming ($\delta = 0$) boundaries as special cases. These \widetilde{b}_j are intended for equally space analyses, i.e., analysis time points satisfy $t_j/t_j = j/J$. For unequally space analyses, these functions of j/J should be replaced by t_j/t_j .

The normal limit (12) still holds when the covariances in (10) are replaced by any consistent estimators $\widetilde{\sigma}_{jj'}$ of $\sigma_{jj'}$, for $j, j' \in [J]$. Let $\widetilde{\boldsymbol{\Sigma}}$ and $\widetilde{\mathbf{D}}$ denote the corresponding matrices with the $\sigma_{jj'}$ replaced by $\widetilde{\sigma}_{jj'}$, i.e., $\widetilde{\mathbf{D}}$ is the $J \times J$ diagonal matrix with entries $\{\widetilde{\sigma}_{jj}^{-1/2}\}$ and $\widetilde{\boldsymbol{\Sigma}}$ is the $J \times J$ matrix with (j, j') entry equal to $\widetilde{\sigma}_{jj'}/\sqrt{\widetilde{\sigma}_{jj}\widetilde{\sigma}_{j'j'}}$ and diagonal entries equal to 1. Then,

$$\widetilde{\boldsymbol{\Sigma}}^{-1/2} \widetilde{\mathbf{D}}\widehat{\boldsymbol{\Delta}}_E \rightarrow N_J(\mathbf{0}, \mathbf{I}) \quad (15)$$

under the null. The W_{Ej} in (13) are the entries of the vector $\widehat{\mathbf{D}}\widehat{\boldsymbol{\Delta}}_E$, where $\widehat{\mathbf{D}}$ is the diagonal matrix with entries $\{\widehat{\sigma}_{jj}^{-1/2}\}$. A normal limit analogous to (12) and (15) holds for $\widehat{\mathbf{D}}\widehat{\boldsymbol{\Delta}}_E$:

$$\widehat{\boldsymbol{\Sigma}}^{-1/2} \widehat{\mathbf{D}}\widehat{\boldsymbol{\Delta}}_E \rightarrow N_J(\mathbf{0}, \mathbf{I}) \quad (16)$$

under the null, where $\widehat{\boldsymbol{\Sigma}}$ is the variance-covariance matrix of $\widehat{\mathbf{D}}\widehat{\boldsymbol{\Delta}}_E$, of which $\widehat{\boldsymbol{\Sigma}}$ is a consistent estimator. Since the vector $\widehat{\mathbf{D}}\widehat{\boldsymbol{\Delta}}_E$ of the W_{Ej} has the same asymptotic limit under the null as $\widehat{\mathbf{D}}\widehat{\boldsymbol{\Delta}}_E$ when both are suitably standardized, the event (13) has the same asymptotic probability under the null as the event

$$\begin{aligned}
 |X_j| \geq b_j, \quad \text{where } \mathbf{X} &= (X_1, \dots, X_J)^T := \widehat{\boldsymbol{\Sigma}}^{1/2} \mathbf{Z} \\
 \text{and } \mathbf{Z} &\sim N_J(\mathbf{0}, \mathbf{I}). \quad (17)
 \end{aligned}$$

For any stopping boundaries b_j of the form (14), moving \widetilde{b}_j to the other side of the inequality, (17) becomes $|X_j|/\widetilde{b}_j \geq b$. From this, we see that choosing b to be the upper- α quantile of the distribution of

$$\max_{j \in [J]} |X_j|/\widetilde{b}_j \quad (18)$$

makes the test (13) approximately level- α .

Since, in general, the elements X_j of $\mathbf{X} = \tilde{\Sigma}^{1/2} \mathbf{Z}$ in (18) are not independent, standard tools from order statistics or extreme value theory cannot be used to calculate (18). However, it can be quickly and accurately computed using Monte Carlo by simulating a large number B of independent multivariate standard normals $\mathbf{Z}^{(1)}, \dots, \mathbf{Z}^{(B)} \sim N_j(\mathbf{0}, \mathbf{I})$, setting $\mathbf{X}^{(k)} = (X_1^{(k)}, X_2^{(k)}, \dots, X_j^{(k)})^T := \tilde{\Sigma}^{1/2} \mathbf{Z}^{(k)}$ for all $k \in [B]$, and taking b to be the upper- α sample quantile of $\{\max_{j \in [J]} |X_j^{(1)}|/\tilde{b}_j, \max_{j \in [J]} |X_j^{(2)}|/\tilde{b}_j, \dots, \max_{j \in [J]} |X_j^{(B)}|/\tilde{b}_j\}$

3.4 Group sequential tests with futility stopping

The distribution theory above makes possible a variety of modifications to the group sequential stopping rule defined above to allow futility stopping, i.e., early stopping before the J th stage to declare failure to reject the null. In this section, we describe one way of doing this with a flexible family of tests known as the power family of two-sided inner wedge tests; these are generalizations of stopping rules of Pampallona and Tsiatis (1994) due to Jennison and Turnbull (2001). Let W_{Ej} , $j \in [J]$, be as in (13). Given stopping boundaries $\{(a_j, b_j)\}_{j \in [J]}$ with

$$0 \leq a_j < b_j \text{ for all } j \in [J-1] \text{ and } a_j = b_j, \quad (19)$$

a general form of a two-sided test with an inner wedge for futility stopping stops at the first analysis $j \in [J]$ such that

$$|W_{Ej}| \notin [a_j, b_j], \quad (20)$$

rejecting H_E if the outer boundary is crossed $|W_{Ej}| \geq b_j$, or declaring failure to reject the null if the statistic enters the “inner wedge” $|W_{Ej}| < a_j$. By virtue of the restriction $a_j = b_j$, the test is guaranteed to stop at the J th stage if it has not stopped prior to that. The power family inner wedge tests (see Jennison and Turnbull, 2000, p. 118) utilize boundaries of the form

$$b_j = b(j/J)^{\delta-1/2} \text{ for } j \in [J], \quad (21)$$

$$a_j = \begin{cases} (a+b)(j/J)^{1/2} - a(j/J)^{\delta-1/2} & \text{for } j_0 \leq j \leq J, \\ 0 & \text{for } 1 \leq j < j_0, \end{cases} \quad (22)$$

with constants a, b, δ , and $j_0 \in [J]$, the latter denoting the earliest analysis at which it is desired to allow futility stopping. Choosing

$$b > 0, \quad \delta \leq 1, \quad \text{and} \quad a \in \left(-b, \frac{b}{(J/j_0)^{1-\delta} - 1}\right) \quad (23)$$

guarantees that the resulting boundaries (21)–(22) will satisfy (19). In Web Appendix C, we provide a method for calculating the constants a, b in the boundaries (21)–(22) to achieve an approximately level- α test.

4 ASSUMPTIONS AND EMPIRICAL ASSESSMENT

We assume the following for the surrogate marker at each time t_j :

$$(C1) \ E(Y_L^{(0)} | S_{Lj}^{(0)} = s) \text{ is a monotone increasing function of } s \text{ for } L = A, B;$$

$$(C2) \ E(Y_L^{(1)} | S_{Lj}^{(1)} = s) \geq E(Y_L^{(0)} | S_{Lj}^{(0)} = s) \text{ for all } s \text{ for } L = A, B;$$

$$(C3) \ P(S_{Lj}^{(1)} > s) \geq P(S_{Lj}^{(0)} > s) \text{ for all } s \text{ for } L = A, B.$$

Assumptions (C1)–(C3) are not unique to our approach and parallel those required in general in surrogate marker research (Wang and Taylor, 2002; Chen et al., 2007; Wu et al., 2011; VanderWeele, 2013; Parast et al., 2017). Assumption (C1) implies that the surrogate marker is positively related to the time of the primary outcome, (C2) implies that there is a non-negative residual treatment effect beyond that on the surrogate marker, and (C3) implies that there is a non-negative treatment effect on the surrogate marker. Our testing approach requires the following additional set of assumptions for the surrogate marker at each t_j :

$$(C4) \ \text{The surrogate captures a reasonably large proportion of the treatment effect on the primary outcome, described below;}$$

$$(C5) \ \text{The support of } S_{Bj}^{(0)} \text{ and } S_{Bj}^{(1)} \text{ must be contained within the support of } S_{Aj}^{(0)};$$

$$(C6) \ E(Y_A^{(0)} | S_{Aj}^{(0)} = s) = E(Y_B^{(0)} | S_{Bj}^{(0)} = s) \text{ for all } s;$$

$$(C7) \ Y_L^{(1)} \perp S_{Lj}^{(0)} | S_{Lj}^{(1)} \text{ and } Y_L^{(0)} \perp S_{Lj}^{(1)} | S_{Lj}^{(0)} \text{ for } L = A, B.$$

While we don't require a perfect surrogate, we do need the surrogate to be useful in some sense of surrogacy where in Assumption (C4), we measure surrogacy using the proportion of the treatment effect on the primary outcome that is explained by the treatment effect on the surrogate marker (Freedman et al., 1992; Wang and Taylor, 2002). If the surrogate is weak in the sense of explaining a low proportion of the treatment effect, we would expect to have lower power to test for a treatment effect on the primary outcome. In this paper, we assume that surrogate marker validation has been previously completed and there is some level of agreement clinically and statistically that the S being considered is a reasonable surrogate. Assumption (C5) is needed for kernel estimation. Assumption (C6) means that in the control groups, the two studies share the same conditional expectations for $Y^{(0)}$ given $S_j^{(0)}$. Though this is certainly a strong assumption, some strong conditions about the transportability between Study A and Study B are necessary in order to borrow information from Study A. If no assumptions about transportability can be made, it is likely not feasible to borrow information from Study A. This assumption is at least only specific to transportability of the control group conditional means and is likely reasonable when the control groups are, e.g., usual care or placebo. Assumption (C7) is needed for identifiability.

It is reasonable to ask whether these assumptions hold and what happens if they do not hold. These questions have largely not been addressed in existing work with the exception of recent work by Elliott et al. (2015) and Shafie Khorassani et al. (2023) in a meta-analytic framework. Assumptions (C6) and (C7) generally cannot be empirically explored with observed data because we do not have the needed quantities i.e., we only have $Y^{(g)}$ and $S^{(g)}$ for an individual if they were in group g , and we do

TABLE 1 Expected stopping time $E(T)$ and probability $P(\text{Rej. } H_E)$ of rejecting the null hypothesis H_E for fixed sample and group sequential procedures without futility stopping, at most $J = 8$ analyses, and nominal level $\alpha = .05$ based on 1000 replications of each procedure in each setting.

Test	Setting 1: no treatment effect		Setting 2: treatment effect	
	$E(T)$ (SE)	$P(\text{Rej. } H_E)$ (SE)	$E(T)$ (SE)	$P(\text{Rej. } H_E)$ (SE)
Fixed sample test	8.000 (.000)	.053 (.007)	8.000 (.000)	.822 (.012)
GS unadjusted	7.314 (.060)	.151 (.011)	5.103 (.073)	.874 (.010)
GS Bonferroni	7.921 (.020)	.021 (.005)	6.694 (.056)	.586 (.016)
Pocock	7.792 (.034)	.045 (.007)	6.173 (.065)	.723 (.014)
O'Brien-Fleming	7.938 (.012)	.054 (.007)	6.609 (.043)	.808 (.012)
Wang-Tsiatis ($\delta = .4$)	7.866 (.024)	.046 (.007)	6.210 (.058)	.762 (.013)

Standard errors SE are in parentheses.

not have any primary outcome measurements, Y , from Study B. However, Assumptions (C1)–(C5) can be empirically explored with available data to some extent. Assumptions (C1)–(C2) can be explored for Study A, but not for Study B (because we do not have primary outcome information); (C3) can be examined in both studies. Assumption (C4) can be explored using only Study A data, while Assumption (C5) can be explored in both studies because it involves only the surrogate marker in both studies. We illustrate assessment of Assumptions (C1)–(C5) in our data application in Section 6 and specifically point out the result of misuse of our method if it is used when an assumption is violated.

5 SIMULATION STUDIES

Here, we present simulation studies of the performance of the proposed group sequential procedures and compare them with a fixed sample procedure and two naive group sequential procedures in two settings: Setting 1 in which there is no treatment effect and Setting 2 in which there is a treatment effect. The data-generating mechanisms for these settings as well as additional simulation results and discussion are provided in [Web Appendix D](#). In both settings, the surrogate marker is possibly measured at $J = 8$ times points $t_j = 1, 2, \dots, 8$ and the primary outcome is measured at $t_8 = 8$. Data generation (for Study B) and testing procedures are replicated 1000 times in each setting to estimate the expected stopping time $E(T)$, where T is the time point at which the procedure terminates and rejects or declares failure to reject H_E , and the probability $P(\text{Rej. } H_E)$ of the procedure rejecting the null. The results are summarized in Tables 1 and 2 which contain the operating characteristics of the proposed group sequential procedures without and with futility stopping, respectively. Both tables include the fixed sample size test, which rejects H_E if and only if (13) holds for $j = J = 8$. The tables also include two naive group sequential procedures, denoted by “GS unadjusted” and “GS Bonferonni,” which stop and reject H_E at the earliest analysis j such that (13) occurs, with $b_1 = \dots = b_j = z_{\alpha/2}$, the upper $\alpha/2$ standard normal quantile, in the unadjusted case, and with $b_1 = \dots = b_j = z_{\alpha/2}/8$ in the Bonferroni-adjusted case. Both of these procedures declare failure to reject the null at the $J = 8$ th analysis if (13) does not occur for any $j \in [J] = [8]$.

Table 1 shows the performance of the group sequential procedures described in Section 3.3 with Pocock, O'Brien-Fleming,

and Wang-Tsiatis power family boundaries (hereafter abbreviated P, OF, and WT). These are given by stopping rule (13) with boundaries of the form (14) with $\tilde{b}_j = 1, \sqrt{j}/j$, and $(j/J)^{\delta-1/2}$ with $\delta = .4$ for the P, OF, and WT boundaries, respectively. The value b in (14) was computed by Monte Carlo as described in Section 3.3. In Setting 1, there is very little savings in expected stopping time $E(T)$ below the maximum value $T = J = 8$, reflecting that the null is true and none of these procedures in this setting allow early stopping to declare failure to reject the null. The P, OF, and WT procedures maintain a type I error probability close to the nominal level of $\alpha = .05$, while the GS unadjusted procedure has a type I error probability that is more than three times $\alpha = .05$. In contrast, the GS Bonferroni stopping value of $z_{\alpha/2}/8 = 2.734$, which is quite conservative, results in a type I error probability of .021. In Setting 2, the P, OF, and WT procedures provide roughly 1–2 analyses of savings on average over the fixed sample test, which always takes $J = 8$ analyses. The P procedure is the most aggressive in early stopping, with the OF procedure the least, and the WT procedure in the middle. The GS unadjusted procedure has the smallest expected stopping time at 5.103, but this is due to its extreme type I error probability exceedance in Setting 1. In Setting 2, the powers of the P, OF, and WT procedures are substantially higher than the GS Bonferroni procedure, and comparable to but slightly less than the fixed sample procedure despite their savings in expected stopping time. Table 2 shows the performance of the group sequential procedures that incorporate futility stopping described in Section 3.4 with $\delta = 1/2, 0$, and .4, respectively. For these procedures, futility stopping begins at the $j_0 = 4$ th analysis. The values (a, b) in (21)–(22) were computed by Monte Carlo as described in Section 3.4. The fixed sample and unadjusted and Bonferroni GS procedures are the same as in Table 1. The P, OF, and WT procedures with futility stopping show more savings in $E(T)$ in both Settings 1 and 2, with the P procedure being the most aggressive in early stopping, followed by the WT and OF procedures. These show substantial savings over the fixed sample and naive GS procedures, which do not allow futility stopping, in Setting 1, and even savings over these procedures in Setting 2, with the exception of the GS unadjusted procedure whose type I error probability is more than 3 times $\alpha = .05$. The power of the P, OF, and WT procedures is higher than the GS Bonferroni procedure, which is conservative in this sense of being underpowered in Setting 2 as well as having a type I error probability of .021 in Setting 1.

TABLE 2 Expected stopping time $E(T)$ and probability $P(\text{Rej. } H_E)$ of rejecting the null hypothesis H_E for fixed sample and group sequential procedures with futility stopping beginning at the $j_0 = 4$ th analysis, at most $J = 8$ analyses, and nominal level $\alpha = .05$ based on 1,000 replications of each procedure in each setting.

Test	Setting 1: no treatment effect		Setting 2: treatment effect	
	$E(T)$ (SE)	$P(\text{Rej. } H_E)$ (SE)	$E(T)$ (SE)	$P(\text{Rej. } H_E)$ (SE)
Fixed sample test	8.000 (.000)	.053 (.007)	8.000 (.000)	.822 (.012)
GS unadjusted	7.314 (.060)	.151 (.011)	5.103 (.073)	.874 (.010)
GS Bonferroni	7.921 (.020)	.021 (.005)	6.694 (.056)	.586 (.016)
Pocock	4.756 (.033)	.051 (.007)	5.572 (.056)	.662 (.015)
O'Brien-Fleming	5.340 (.032)	.053 (.007)	5.619 (.040)	.706 (.014)
Wang-Tsiatis ($\delta = .4$)	4.908 (.032)	.051 (.007)	5.333 (.049)	.674 (.015)

Standard errors SE are in parentheses.

6 APPLICATION TO AIDS CLINICAL TRIALS

We illustrate our proposed testing procedure using data from two randomized clinical trials from the AIDS Clinical Trial Group (ACTG) Network: ACTG 320 (Study A) and ACTG 193A (Study B). ACTG 320 was a randomized, double-blind trial among 1156 HIV-infected patients that compared a two-drug regimen (two nucleoside reverse transcriptase inhibitors [NRTI]; $n = 579$) with a three-drug regimen (two NRTIs plus indinavir; $n = 577$) (Hammer et al., 1997). ACTG 193A was a randomized, double-blind trial among 1,313 HIV-infected patients comparing four daily treatment regimes (Henry et al., 1998). We focus on comparing a two-drug regimen (zidovudine plus zalcitabine, 2 NRTIs; $n = 326$, group 0) with a three-drug regimen (two NRTIs plus nevirapine; $n = 330$, group 1) in this study. For our analysis, the primary outcome is change in RNA from baseline to 40 weeks post-baseline which was measured in Study A but was *not* measured for all participants in Study B. The surrogate marker of interest is change in CD4 count from baseline to t_j , $0 < t_j \leq 40$. That is, we aim to test for a treatment effect on change in RNA at 40 weeks in Study B (ACTG 320 Study) using the surrogate marker of change in CD4 count via our proposed group sequential testing approach. Importantly, the exact times of the CD4 measurements do not perfectly align between studies. In Study A, CD4 was measured at baseline, 4, 8, 24, and 40 weeks; in Study B, CD4 was measured at baseline, 8, 16, 24, and 40 weeks. Before we discuss how to handle this, we first assessed Assumptions (C1)–(C3). Detailed results, provided in Web Appendix E, show that we should not borrow information about the surrogate marker at weeks 4 or 8 from Study A. Furthermore, there is no measurement available at 4 weeks in Study B. In Study B, we only consider testing at 16, 24, and 40 weeks. However, because the surrogate was not measured at 16 weeks in Study A, we construct the test statistic at 16 weeks by borrowing the conditional mean function $\hat{\mu}_{A0}(s)$ at 24 weeks.

Figure 3(a) shows the resulting test statistics with dashed lines drawn at the naive, Bonferroni, Pocock, O'Brien-Fleming, and Wang-Tsiatis (with $\delta = .4$) boundaries as described in Section 3.3 with $J = 3$ analyses. The constant b in (14) for each of these procedures was calculated by Monte Carlo with the quantity j/J appearing in the \tilde{b}_j expressions replaced by t_j/t_j with $t_j = 16, 24$, and 40 because of the unequally spaced analyses. For all boundaries, we would conclude a significant treatment effect at 40 weeks using the surrogate marker, CD4 count. To illustrate the importance of empirically examining the needed

assumptions (which we do in Web Appendix E), we show in Figure 3(b), the results we *would have* obtained had we included testing at 8 weeks in Study B, using the surrogate marker information at 8 weeks from Study A, even though we have empirical evidence of possible assumption violation. Here, we would have incorrectly concluded a significant negative treatment effect using the naive, Bonferroni, and Pocock boundary, but not with the O'Brien-Fleming or and Wang-Tsiatis boundaries. In practice, it is important that these assumptions be examined empirically as described in Section 4, to the extent possible with available data, to reduce the risk of an inappropriate conclusion (such as this one) based on the surrogate marker.

7 DISCUSSION

There is a tremendous amount of hope in the promise of surrogate markers that can identify effective (and ineffective) treatments sooner. In this paper, we offer group sequential procedures that allow for early stopping to declare efficacy of the treatment effect or early futility stopping. If the surrogate measurements over time were independent, this would be a quite straightforward problem. However, it would be unreasonable to assume such independence, at least in any clinical setting, and thus, our statistical development of appropriate procedures was complicated by and accounted for the correlation between the surrogate marker measurements over time. An R package that implements our procedures, named `SurrogateSeq`, is available on GitHub (Parast and Bartroff, 2024).

The simulation studies in Section 5 show the proposed group sequential procedures can be made to control the type I error close to a nominal level α . The versions of these procedures without futility stopping provide modest savings in sample size while maintaining comparable power, and the versions incorporating futility stopping provide more savings in sample size while having higher power than the Bonferroni-adjusted, naive group sequential procedure. Our results show that the naive, unadjusted group sequential procedure is not a viable option because its type I error probability grossly exceeds the nominal level, a well-understood phenomenon in many sequential and group sequential settings. In practice, when deciding between the Pocock, O'Brien-Fleming, and Wang-Tsiatis versions of the proposed procedures, we recommend the Pocock boundaries if aggressive early stopping or simplicity of stop-

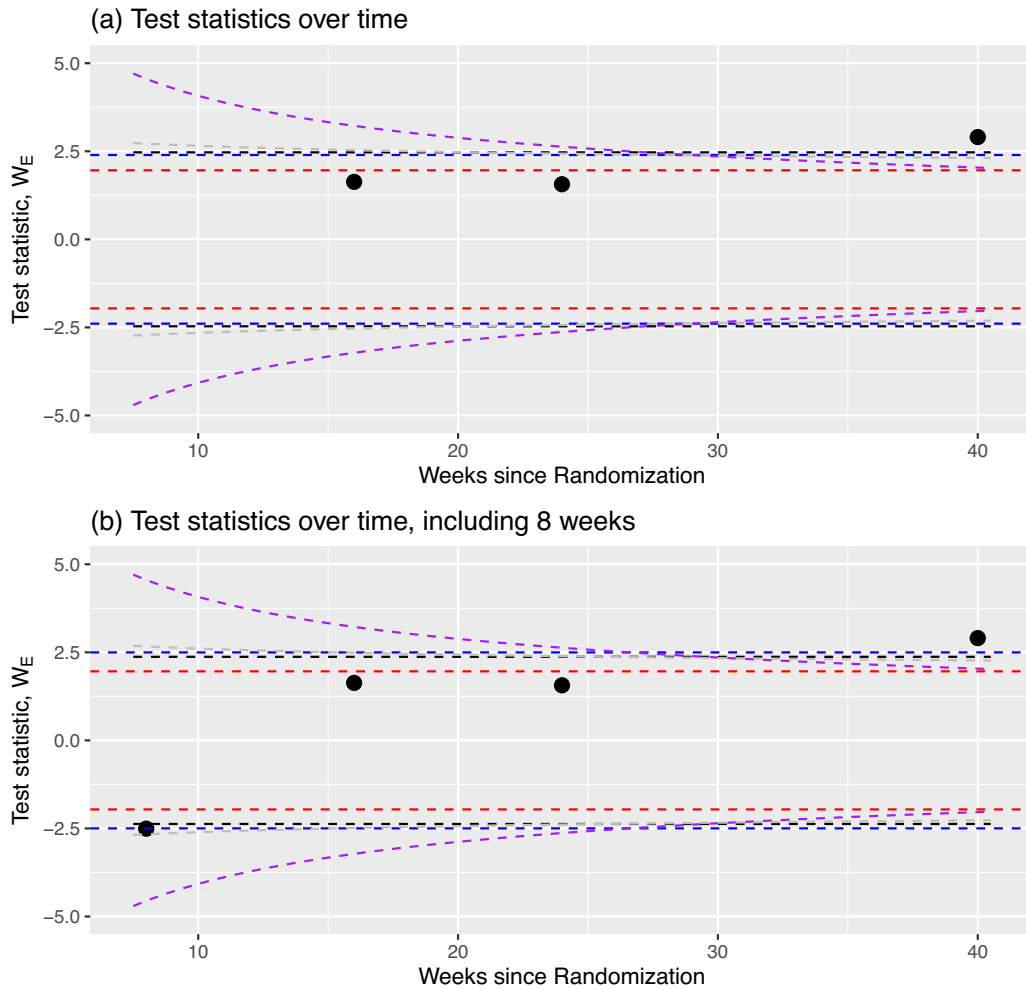


FIGURE 3 AIDS clinical trial testing results where solid dots indicate the estimated test statistic, W_E , at each time point t_j for weeks 16, 24, and 40 (a), and weeks 8, 16, 24, and 40 (b); a red dashed line (closest to center) is drawn at $\pm\Phi^{-1}(1 - 0.05/2) = 1.96$, a blue dashed line (at about 2.5) is drawn at the Bonferroni boundary, a purple dashed line (outermost line) is drawn at the O’Brien-Fleming boundary, a black dashed line (between 1.96 and Bonferroni) is drawn at the Pocock boundary, and a gray dashed line (between Bonferroni and outermost line) is drawn at the Wang-Tsiatis boundary.

ping boundaries (ie, constant) is desired, the O’Brien-Fleming boundaries if maintaining high power is a primary objective; otherwise, some instance of the Wang-Tsiatis boundaries can be chosen as a “middle ground.”

While a major advantage of our proposed approach is the lack of reliance on parametric assumptions, we still require a number of assumptions described in Section 4. These assumptions are not unique to our approach; some version of these assumptions is generally required in surrogate marker research (VanderWeele, 2013). Importantly, we do not claim that these are replacing parametric assumptions and are less strict. They are strict assumptions, and while we can empirically examine some of these with available data, we cannot examine or test Assumptions (C6) and (C7). Assumption (C6) has some parallels to assumptions used in domain adaptation (Kouw and Loog, 2019) and Assumption (C7) has some parallels to (though is not equivalent to) the widely used assumption of no unmeasured confounding in causal inference (VanderWeele and Arah, 2011). Future work to develop methods to investigate these assump-

tions via simulation or the construction of bounds would be valuable.

The proposed method for calculating the stopping boundaries of these procedures in Section 3.3 relies on Assumption (C6) which permits estimating the covariance structure of the Study B test statistics using that from Study A. If Study A does not provide an accurate estimate of the Study B covariance, an alternative approach may be to compute the boundaries adaptively using a type I error-spending function $\alpha(t)$, increasing from $\alpha(0) = 0$ to $\alpha(1) = \alpha$. In the context of the group sequential tests in Section 3.3 with stopping rule (13), after the $(j - 1)$ st analysis has been completed and data from the j th analysis is available, the stopping boundary b_j would be computed to satisfy $P(|W_{E1}| < b_1, \dots, |W_{Ej-1}| < b_{j-1}, |W_{Ej}| \geq b_j) = \alpha(t_j/t_j) - \alpha(t_{j-1}/t_j)$ where this probability is under the null hypothesis. This event can be simulated via Monte Carlo using a similar approach to (17) but where \mathbf{Z} is length j and $\tilde{\Sigma}$ is replaced by the $j \times j$ leading principal submatrix of $\tilde{\Sigma}$, a function of the available data from the first j analyses, producing

length j vectors \mathbf{X} whose entries have the same distribution as $W_{E1}, W_{E2}, \dots, W_{Ej}$ under the null.

While this work is a unique and novel contribution to the field, the practical utility of this approach depends immensely on the quality, size, and representativeness of the Study A data. Notably, our testing framework makes explicit the reliance on Study A data via the definition of Δ_E which is defined with the $\hat{\mu}_{A0j}(s)$ term. (However, for completeness, we discuss our approach with Study A as random in [Web Appendix F](#).) While it is less common to see an estimator within a defined null hypothesis, this reflects the true use of surrogates in practice. It would be unrealistic to develop a testing procedure and asymptotic results under the assumption that $n_A \rightarrow \infty$ because, in practice, Study A is generally not large and is not set by the designers of Study B. When a study team is considering using a surrogate marker to test for a treatment effect, they are using information about surrogacy from some previously completed study. To the designers of Study B, the design is such that this prior Study A is indeed fixed. This is *not* the same thing as assuming that the results of Study A were “exact” or not subject to its own variation, but just that they they will not change before Study B begins. If Study A was very small or not representative or in some way not providing a reasonable estimate of $\mu_{A0j}(s)$, there is no “magic bullet” (that we know of) that would produce a legitimate testing procedure that uses the surrogate as a replacement of the primary outcome, borrowing from Study A. Ultimately, we admit that this is a difficult problem. Surrogate markers are imperfect and it is important for us to be explicit about the reliance of this testing procedure (and more generally, any procedure that uses a surrogate marker to replace the primary outcome) on prior data that is being used to assess surrogacy.

ACKNOWLEDGEMENTS

We are grateful to the AIDS Clinical Trial Group (ACTG) Network for supplying the ACTG data.

SUPPLEMENTARY MATERIAL

Supplementary material is available at [Biometrics](#) online.

Referenced Web Appendices and Tables and a zip file with R code to reproduce the simulations and example referenced in Sections 2, 3, 5, 6, and 7 are available with this paper at the Biometrics website on Oxford Academic. The R package implementing these procedures, named `SurrogateSeq`, is available on GitHub at <https://github.com/layl aparast/SurrogateSeq>.

FUNDING

This work was supported by NIDDK grant R01DK118354.

CONFLICT OF INTEREST

None declared.

DATA AVAILABILITY

The AIDS clinical trial data used in this paper are publicly available in our R package `SurrogateSeq` available at <https://github.com/layl aparast/SurrogateSeq>.

REFERENCES

- Anderer, A., Bastani, H. and Silberholz, J. (2022). Adaptive clinical trial designs with surrogates: When should we bother? *Management Science*, 68, 1982–2002.
- Bartroff, J., Lai, T. L. and Shih, M. (2013). *Sequential Experimentation in Clinical Trials: Design and Analysis*, New York, NY: Springer.
- Chen, H., Geng, Z. and Jia, J. (2007). Criteria for surrogate end points. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 69, 919–932.
- Cook, R. J. and Farewell, V. T. (1996). Incorporating surrogate endpoints into group sequential trials. *Biometrical journal*, 38, 119–130.
- Elliott, M. R. (2023). Surrogate endpoints in clinical trials. *Annual Review of Statistics and its Application*, 10, 75–96.
- Elliott, M. R., Conlon, A. S., Li, Y., Kaciroti, N. and Taylor, J. M. (2015). Surrogacy marker paradox measures in meta-analytic settings. *Biostatistics*, 16, 400–412.
- FDA (2023). *Surrogate endpoint resources for drug and biologic development*, <https://www.fda.gov/drugs/development-resources/surrogate-endpoint-resources-drug-and-biologic-development> [Accessed 1 August 2024].
- Freedman, L. S., Graubard, B. I. and Schatzkin, A. (1992). Statistical validation of intermediate endpoints for chronic diseases. *Statistics in Medicine*, 11, 167–178.
- Hammer, S. M., Squires, K. E., Hughes, M. D., Grimes, J. M., Demeter, L. M., Currier, J. S. et al. (1997). A controlled trial of two nucleoside analogues plus indinavir in persons with human immunodeficiency virus infection and cd4 cell counts of 200 per cubic millimeter or less. *New England Journal of Medicine*, 337, 725–733.
- Henry, K., Erice, A., Tierney, C., Balfour, H. H., Fischl, M. A., Kmark, A. et al. (1998). A randomized, controlled, double-blind study comparing the survival benefit of four different reverse transcriptase inhibitor therapies (three-drug, two-drug, and alternating drug) for the treatment of advanced aids. *Journal of Acquired Immune Deficiency Syndromes and Human Retrovirology*, 19, 339–349.
- Jennison, C. and Turnbull, B. W. (2000). *Group Sequential Methods with Applications to Clinical Trials*. New York, NY: Chapman & Hall/CRC.
- Jennison, C. and Turnbull, B. W. (2001). On group sequential tests for data in unequally sized groups and with unknown variance. *Journal of Statistical Planning and Inference*, 96, 263–288.
- Kim, K. and Tsiatis, A. (2020). Independent increments in group sequential tests: a review. *SORT: Statistics and Operations Research Transactions*, 44, 0223–264.
- Kouw, W. M. and Loog, M. (2019). A Review of Domain Adaptation Without Target Labels. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43, 766–785.
- Li, Q., Lin, J., Liu, M., Wu, L. and Liu, Y. (2022). Using surrogate endpoints in adaptive designs with delayed treatment effect. *Statistics in Biopharmaceutical Research*, 14, 661–670.
- Lin, D. (1991). Nonparametric sequential testing in clinical trials with incomplete multivariate observations. *Biometrika*, 78, 123–131.
- O’Brien, P. C. and Fleming, T. R. (1979). A multiple testing procedure for clinical trials. *Biometrics*, 35, 549–556.
- Pampallona, S. and Tsiatis, A. A. (1994). Group sequential designs for one-sided and two-sided hypothesis testing with provision for early stopping in favor of the null hypothesis. *Journal of Statistical Planning and Inference*, 42, 19–35.
- Parast, L. and Bartroff, J. (2024). *SurrogateSeq* r package, version 1.0, <https://github.com/layl aparast/SurrogateSeq> [Accessed 1 August 2024].
- Parast, L., Cai, T. and Tian, L. (2017). Evaluating surrogate marker information using censored data. *Statistics in Medicine*, 36, 1767–1782.

- Parast, L., Cai, T. and Tian, L. (2019). Using a surrogate marker for early testing of a treatment effect. *Biometrics*, 75, 1253–1263.
- Parast, L., Cai, T. and Tian, L. (2023). Using a surrogate with heterogeneous utility to test for a treatment effect. *Statistics in Medicine*, 42, 68–88.
- Pocock, S. J. (1977). Group sequential methods in the design and analysis of clinical trials. *Biometrika*, 64, 191–199.
- Price, B. L., Gilbert, P. B. and van der Laan, M. J. (2018). Estimation of the optimal surrogate based on a randomized trial. *Biometrics*, 74, 1271–1281.
- Quan, H., Xu, Z., Luo, J., Paux, G., Cho, M. and Chen, X. (2023). Utilization of treatment effect on a surrogate endpoint for planning a study to evaluate treatment effect on a final endpoint. *Pharmaceutical Statistics*, 22, 633–649.
- Saint-Hilary, G., Barboux, V., Pannaux, M., Gasparini, M., Robert, V. and Mastrantonio, G. (2019). Predictive probability of success using surrogate endpoints. *Statistics in Medicine*, 38, 1753–1774.
- Shafie Khorassani, F., Taylor, J. M., Kaciroti, N. and Elliott, M. R. (2023). Incorporating covariates into measures of surrogate paradox risk. *Stats*, 6, 322–344.
- Spiessens, B., Lesaffre, E., Verbeke, G., Kim, K. and DeMets, D. L. (2000). An overview of group sequential methods in longitudinal clinical trials. *Statistical Methods in Medical Research*, 9, 497–515.
- Tang, D.-I., Gnecco, C. and Geller, N. L. (1989). Design of group sequential clinical trials with multiple endpoints. *Journal of the American Statistical Association*, 84, 775–779.
- Temple, R. (1999). Are surrogate markers adequate to assess cardiovascular disease drugs? *JAMA*, 282, 790–795.
- VanderWeele, T. J. (2013). Surrogate measures and consistent surrogates. *Biometrics*, 69, 561–565.
- VanderWeele, T. J. and Arah, O. A. (2011). Bias formulas for sensitivity analysis of unmeasured confounding for general outcomes, treatments, and confounders. *Epidemiology*, 22, 42–52.
- Wang, S. K. and Tsiatis, A. A. (1987). Approximately optimal one-parameter boundaries for group sequential trials. *Biometrics*, 43, 193–200.
- Wang, Y. and Taylor, J. M. (2002). A measure of the proportion of treatment effect explained by a surrogate marker. *Biometrics*, 58, 803–812.
- Wu, Z., He, P. and Geng, Z. (2011). Sufficient conditions for concluding surrogacy based on observed data. *Statistics in Medicine*, 30, 2422–2434.