Received: 13 September 2022 | Accepted: 2 March 2023

DOI: 10.1111/biom.13854

BIOMETRIC METHODOLOGY

Biometrics A JOURNAL OF THE INTERNATIONAL BIOMETRIC SOCIETY WILEY

Interim monitoring of sequential multiple assignment randomized trials using partial information

Cole Manschot[1] | **Eric Laber**[2] | **Marie Davidian**[1]

[1]Department of Statistics, North Carolina State University, Raleigh, North Carolina, USA

[2]Department of Statistical Science and Department of Biostatistics & Bioinformatics, Duke University, Durham, North Carolina, USA

Correspondence
Cole Manschot, Department of Statistics, North Carolina State University, Raleigh, NC 27695, USA.
Email: cmansch@ncsu.edu

Abstract

The sequential multiple assignment randomized trial (SMART) is the gold standard trial design to generate data for the evaluation of multistage treatment regimes. As with conventional (single-stage) randomized clinical trials, interim monitoring allows early stopping; however, there are few methods for principled interim analysis in SMARTs. Because SMARTs involve multiple stages of treatment, a key challenge is that not all enrolled participants will have progressed through all treatment stages at the time of an interim analysis. Wu et al. (2021) propose basing interim analyses on an estimator for the mean outcome under a given regime that uses data only from participants who have completed all treatment stages. We propose an estimator for the mean outcome under a given regime that gains efficiency by using partial information from enrolled participants regardless of their progression through treatment stages. Using the asymptotic distribution of this estimator, we derive associated Pocock and O'Brien-Fleming testing procedures for early stopping. In simulation experiments, the estimator controls type I error and achieves nominal power while reducing expected sample size relative to the method of Wu et al. (2021). We present an illustrative application of the proposed estimator based on a recent SMART evaluating behavioral pain interventions for breast cancer patients.

KEYWORDS

augmented inverse probability weighting, clinical trials, double robustness, dynamic treatment regimes, early stopping, group sequential analysis

1 | INTRODUCTION

Treatment of chronic diseases and disorders involves a series of treatment decisions made at critical points in the progression of a patient's health status. To optimize long-term health outcomes, these decisions must adapt to evolving patient information, including response to previous treatments. Strategies for adapting treatment decisions over time are formalized as treatment regimes, which comprise a sequence of decision rules, one per stage of intervention, that map accrued patient information to a recommended treatment (Chakraborty & Moodie, 2013; Tsiatis et al., 2020). The value of a regime is the expected utility if the regime is used to select treatments in the population of interest. A regime is optimal if it has maximal value. Much of the statistical literature on treatment regimes has focused on estimation and inference for optimal regimes (Kosorok & Laber, 2019). However, scientific

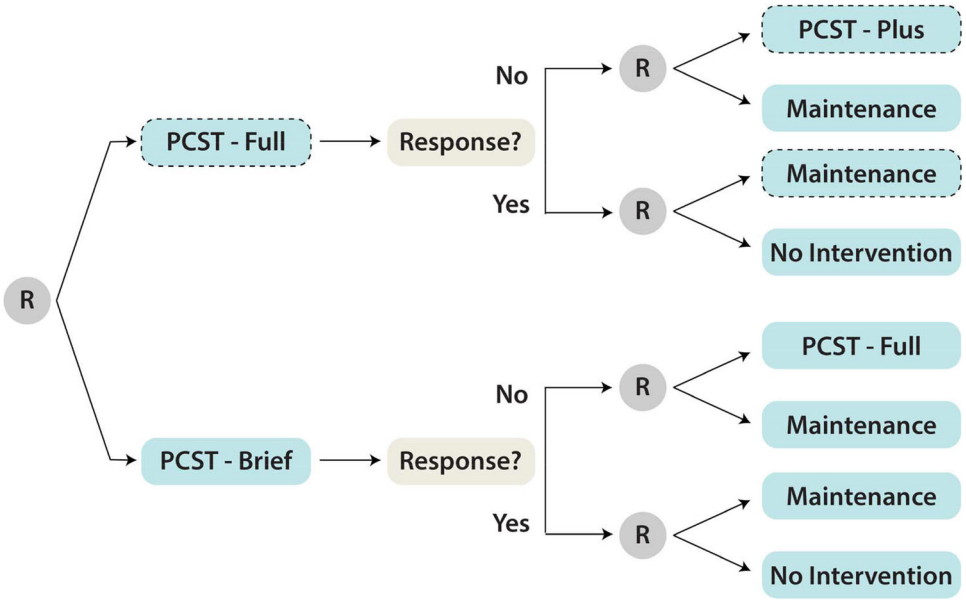

FIGURE 1 Schema for the sequential multiple assignment randomized trial (SMART) evaluating regimes involving behavioral interventions for pain management in breast cancer patients embedding eight regimes of the form "Give intervention a; if nonresponse, give b; otherwise, if response give c." The embedded regime determined by a = PCST (pain coping skills training)-Full, b = PCST-Plus, c = Maintenance is shown with dashed lines around the treatments. Regimes $l = 1, \ldots, 8$ take (a, b, c) to be (Full, Plus, Maintenance), (Full, Plus, No Intervention), (Full, Maintenance, Maintenance), (Full, Maintenance, No Intervention), (Brief, Full, Maintenance), (Brief, Full, No Intervention), (Brief, Maintenance, Maintenance), and (Brief, Maintenance, No Intervention), respectively. This figure appears in color in the electronic version of this article.

interest often focuses on comparison of a small number of prespecified treatment regimes, either with each other or against a control, on the basis of mean outcome.

The gold standard for data collection for the evaluation of treatment regimes is the sequential multiple assignment randomized trial (SMART; Lavori & Dawson, 2004; Murphy, 2005). A SMART contains multiple stages of randomization, with each stage corresponding to a key decision point. In a SMART, if, when, and to whom a treatment might be randomly assigned is allowed to depend on a patient's treatment and outcome history, leading to a rich and flexible class of designs. In the past decade, the use of SMARTs has increased dramatically; SMARTs have been conducted in a range of disease and disorder areas, including cancer (Kelleher et al., 2017; Thall, 2015; Wang et al., 2012), behavioral sciences (Almirall et al., 2014; Kidwell & Hyde, 2016), and mental health (Manschreck & Boshes, 2007; Sinyor et al., 2010). For a comprehensive list of SMARTs, see Bigirumurame et al. (2022).

Every SMART can be equivalently represented as randomizing subjects at baseline among a set of fixed regimes known as the trial's "embedded regimes." Primary analyses in a SMART often focus on comparisons of the embedded regimes against each other or a control (Lavori & Dawson, 2004; Murphy, 2005). These comparisons are often used for sizing a SMART (Artman et al., 2020; Seewald et al., 2020). For example, Figure 1 shows a two-stage SMART schema for behavioral interventions for pain management in cancer patients with eight embedded regimes (Kelleher et al., 2017; ClinicalTrials.gov, 2021). Each embedded regime takes the form "give intervention a; if response, give b; if nonresponse, give c;" for example, give pain coping skills training (PCST) Full initially; if response, give maintenance; otherwise, give PCST-Plus. We discuss this study further in Section 7.

Interim monitoring allows early stopping for efficacy or futility, which can reduce cost and accelerate evaluation of candidate treatments. Group sequential methods allowing early stopping are well established for conventional clinical trials (Jennison & Turnbull, 2000). However, analogous methodology for SMARTs is limited. Wu et al. (2021) propose an interim test for a difference in mean outcome among embedded regimes in two-stage SMARTs. However, their approach is based on the inverse probability weighted estimator (IPWE), which does not incorporate baseline and accruing patient information that could be used to enhance efficiency (Zhang et al., 2013). Chao et al. (2020) consider interim analysis for a small-n, two-stage SMART restricted to the specific situation in which the same treatments are available at each stage and the goal is to remove futile treatments.

We develop a class of interim analysis methods for SMARTs based on an augmented inverse probability weighted estimator (AIPWE) for the value of a regime that increases statistical efficiency by using partial information from individuals with incomplete regime trajectories.

Our method applies to SMARTs with an arbitrary number of stages and treatments, as well as those in which the set of allowable treatments depends on a patient's history. We present the statistical framework in Section 2. In Section 3, we review the AIPWE for the value of a regime when all participants have progressed through all stages. We introduce the proposed Interim AIPWE in Section 4. In Section 5, we discuss testing procedures, stopping boundaries, and sample size formulas for interim analysis. In Section 6, we evaluate the empirical performance of the proposed procedure in a series of simulation experiments, and we present a case study based on the cancer pain management SMART in Section 7.

2 | STATISTICAL FRAMEWORK

Consider a SMART with K stages and a planned total sample size of N. Each subject completing the trial generates a trajectory of the form $(\mathbf{X}_1, A_1, \mathbf{X}_2, A_2, \dots, \mathbf{X}_K, A_K, Y)$, where $A_k \in \mathcal{A}_k$, $k = 1, \dots, K$, is the treatment assigned at stage k; $\mathcal{A}_k$ is a finite set of treatment options at decision k; $\mathbf{X}_1 \in \mathbb{R}^{p_1}$ comprises baseline subject variables; $\mathbf{X}_k \in \mathbb{R}^{p_k}$, $k = 2, \dots, K$, comprises subject variables collected between stages $k-1$ and k; and $Y \in \mathbb{R}$ is an outcome measured at the end of follow-up, coded so that higher values are better. Let $\overline{\mathbf{X}}_k = (\mathbf{X}_1, \dots, \mathbf{X}_k)$ and $\overline{\mathbf{A}}_k = (A_1, \dots, A_k)$, and define $\mathbf{H}_1 = \mathbf{X}_1$ and $\mathbf{H}_k = (\overline{\mathbf{X}}_k, \overline{\mathbf{A}}_{k-1})$, $k \geq 2$, so that $\mathbf{H}_k$ is the information available at the time A_k is assigned. Let $\mathcal{H}_k = \text{dom}\, \mathbf{H}_k$ and let $2^{\mathcal{A}_k}$ denote the power set of $\mathcal{A}_k$. We assume there exists a set-valued function $\Psi_k : \mathcal{H}_k \to 2^{\mathcal{A}_k}$ so that the set of allowable treatments for a subject with $\mathbf{H}_k = \mathbf{h}_k$ at stage k is $\Psi_k(\mathbf{h}_k) \subseteq \mathcal{A}_k$ (Tsiatis et al., 2020; van der Laan & Petersen, 2007).

In this setting, a treatment regime is a sequence of decision rules, $\mathbf{d} = (d_1, \dots, d_K)$, where $d_k : \mathcal{H}_k \to \mathcal{A}_k$ and $d_k(\mathbf{h}_k) \in \Psi_k(\mathbf{h}_k)$ for all $\mathbf{h}_k \in \mathcal{H}_k$, $k = 1, \dots, K$. Let $Y^*(\overline{\mathbf{a}}_K)$ denote the potential outcome under treatment sequence $\overline{\mathbf{a}}_K = (a_1, \dots, a_K)$, and let $\mathbf{X}_k^*(\overline{\mathbf{a}}_{k-1})$ denote the potential intermediate variables under sequence $\overline{\mathbf{a}}_{k-1}$ at stage $k \geq 2$. Define $\overline{\mathbf{X}}_k^*(\overline{\mathbf{a}}_{k-1}) = \{\mathbf{X}_1, \mathbf{X}_2^*(a_1), \dots, \mathbf{X}_k^*(\overline{\mathbf{a}}_{k-1})\}$, $\mathbf{H}_k^*(\overline{\mathbf{a}}_{k-1}) = \{\overline{\mathbf{X}}_k^*(\overline{\mathbf{a}}_{k-1}), \overline{\mathbf{a}}_{k-1}\}$, and $\mathbf{H}_1^*(a_0) = \mathbf{H}_1$. The potential covariates and outcome for an individual receiving treatment according to regime $\mathbf{d}$ are

$$\mathbf{X}_k^*(\mathbf{d}) = \sum_{\overline{\mathbf{a}}_{k-1} \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_{k-1}} \mathbf{X}^*(\overline{\mathbf{a}}_{k-1}) \prod_{k=1}^{k-1} I\left[a_k = d_k\left\{\mathbf{H}_k^*(\overline{\mathbf{a}}_{k-1})\right\}\right], k = 2, \dots, K,$$

$$Y^*(\mathbf{d}) = \sum_{\overline{\mathbf{a}}_K \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_K} Y^*(\overline{\mathbf{a}}_K) \prod_{k=1}^{K} I\left[a_k = d_k\left\{\mathbf{H}_k^*(\overline{\mathbf{a}}_{k-1})\right\}\right].$$

The mean outcome, or value, for a regime $\mathbf{d}$ is $V(\mathbf{d}) = \mathbb{E}\{Y^*(\mathbf{d})\}$.

In a SMART, primary analyses often focus on inference on $V(\mathbf{d})$ for regimes $\mathbf{d}$ that are embedded in the trial. Let $\pi_k(a_k, \mathbf{h}_k) = P(A_k = a_k | \mathbf{H}_k = \mathbf{h}_k)$ be the probability (propensity) of being randomized to treatment $a_k \in \Psi_k(\mathbf{h}_k)$ at stage k for a subject with history $\mathbf{h}_k$. It is well known that $V(\mathbf{d})$ is identifiable under the following conditions: positivity ($\pi_k(a_k, \mathbf{h}_k) > 0$ for all $\mathbf{h}_k \in \mathcal{H}_k$ and $a_k \in \Psi_k(\mathbf{h}_k)$); sequential randomization ($\{\mathbf{X}_1, \mathbf{X}_2^*(a_1), \dots, \mathbf{X}_k^*(\overline{\mathbf{a}}_{k-1}), Y^*(\overline{\mathbf{a}}_K)\}_{\overline{\mathbf{a}}_K \in \mathcal{A}_1 \times \cdots \times \mathcal{A}_K} \perp\!\!\!\perp A_k | \mathbf{H}_k$ at each stage k for all $\overline{\mathbf{a}}_K$, where $\perp\!\!\!\perp$ denotes independence), which holds by design in a SMART; consistency, $Y = Y^*(\overline{\mathbf{A}}_K)$ and $\mathbf{H}_k = \mathbf{H}_k^*(\overline{\mathbf{A}}_{k-1})$; and no interference among subjects (Tsiatis et al., 2020). Hereafter, we assume that these conditions hold.

Take $\mathbf{d}^\ell$, $\ell = 1, \dots, L$, to be the regimes embedded in the SMART and $\mathbf{d}^0$ a possible control, for example, a treatment or regime representing the standard of care. For definiteness, we consider two null hypotheses that address the efficacy of the embedded regimes:

$$\text{Homogeneity} \qquad H_{0H} : V(\mathbf{d}^1) = \cdots = V(\mathbf{d}^L), \tag{1}$$

$$\text{Superiority} \quad H_{0D} : V(\mathbf{d}^\ell) - V(\mathbf{d}^0) \leq \delta \text{ for all } \ell = 1, \dots, L. \tag{2}$$

These hypotheses are analogous to those used in multi-arm, multi-stage and platform trials (Jennison and Turnbull, 2000, Chap. 16; Wason, 2019). The control value $V(\mathbf{d}^0)$ may be fixed or estimated from an additional control arm. The methods presented here apply to futility testing with minor modification. Hypotheses that stop the trial for either a single regime or all regimes falling below an efficacy boundary are possible in this construction.

3 | AIPWE FOR COMPLETE DATA

We briefly review the AIPWE of the value in the setting where one observes N complete independent and identically distributed trajectories $\{\mathbf{X}_{1,i}, A_{1,i}, \dots, \mathbf{X}_{K,i}, A_{K,i}, Y_i\}_{i=1}^N$. For any regime $\mathbf{d}$, define $C_{\mathbf{d},k} = \prod_{j=1}^k I\{A_j = d_j(\mathbf{H}_j)\}$ to be an indicator that treatment is consistent with $\mathbf{d}$ through the first k decisions, and let $C_{\mathbf{d},0} = 1$. For each $k = 1, \dots, K$, let $\pi_k(a_k, \mathbf{h}_k; \theta_k)$ be a posited model for $\pi_k(a_k, \mathbf{h}_k)$ indexed by $\theta_k \in \Theta_k$. Although the propensities are known in a SMART, estimating them based on correctly specified models can increase efficiency (Tsiatis, 2006b). Let $\widehat{\theta}_{k,N}$ be an estimator of θ_k. The form of the AIPWE for $V(\mathbf{d})$ is (Zhang et al., 2013; Tsiatis et al., 2020, Section 6.4.4)

$$\widehat{V}_{\mathrm{A}}(\mathbf{d}) = N^{-1}\sum_{i=1}^{N}\left[\frac{Y_i\, C_{\mathbf{d},K,i}}{\prod_{k=1}^{K}\pi_k(A_{k,i},\mathbf{H}_{k,i};\widehat{\theta}_{k,N})} + \sum_{k=1}^{K}\left\{\frac{C_{\mathbf{d},k-1,i}}{\prod_{v=1}^{k-1}\pi_v(A_{v,i},\mathbf{H}_{v,i};\widehat{\theta}_{v,N})} - \frac{C_{\mathbf{d},k,i}}{\prod_{v=1}^{k}\pi_v(A_{v,i},\mathbf{H}_{v,i};\widehat{\theta}_{v,N})}\right\}L_k(\overline{\mathbf{X}}_{ki})\right], \quad (3)$$

where $L_k(\overline{\mathbf{x}}_k)$ is an arbitrary function of $\overline{\mathbf{x}}_k$ and we define $\prod_{v=1}^{0}\pi_v(A_{v,i},\mathbf{H}_{v,i};\widehat{\theta}_{v,N}) = 1$. Setting $L_k(\overline{\mathbf{x}}_k) \equiv 0$ yields the IPWE that forms the basis for the approach of Wu et al. (2021); the IPWE uses only the observed outcomes and no covariate information. It can be an inefficient estimator for $V(\mathbf{d})$ when there are covariates that are correlated with the outcome. The efficient choice for L_k is $L_k^{\mathbf{d}}(\overline{\mathbf{x}}_k) = E\{Y^*(\mathbf{d})|\overline{\mathbf{X}}_k = \overline{\mathbf{x}}_k, \overline{\mathbf{A}}_{k-1} = \overline{\mathbf{d}}_{k-1}(\overline{\mathbf{x}}_{k-1})\}$, where $\overline{\mathbf{d}}_k(\overline{\mathbf{x}}_k)$ is defined as follows: $\overline{d}_1(\mathbf{x}_1) = d_1(\mathbf{x}_1)$, $\overline{\mathbf{d}}_2(\overline{\mathbf{x}}_2) = [d_1(\mathbf{x}_1), d_2\{\overline{\mathbf{x}}_2, d_1(\mathbf{x}_1)\}], \ldots, \overline{\mathbf{d}}_k(\overline{\mathbf{x}}_k) = [d_1(\mathbf{x}_1), d_2\{\overline{\mathbf{x}}_2, d_1(\mathbf{x}_1)\}, \ldots, d_k\{\overline{\mathbf{x}}_k, \overline{\mathbf{d}}_{k-1}(\overline{\mathbf{x}}_{k-1})\}]$, $k = 1, \ldots, K$; and $\overline{\mathbf{A}}_k = \overline{\mathbf{d}}_k(\overline{\mathbf{X}}_k)$ is the event that all treatments received are consistent with $\mathbf{d}$ through decision k.

In practice, the functions $L_k^{\mathbf{d}}(\overline{\mathbf{x}}_k)$, $k = 1, \ldots, K$ are unknown, but they can be estimated using Q-learning as follows (Tsiatis et al., 2020, Section 6.4.2). Posit a model $Q_K(\overline{\mathbf{x}}_K, \overline{\mathbf{a}}_K; \boldsymbol{\beta}_K)$ for $Q_K(\overline{\mathbf{x}}_K, \overline{\mathbf{a}}_K) = \mathbb{E}(Y|\overline{\mathbf{X}}_K = \overline{\mathbf{x}}, \overline{\mathbf{A}}_K = \overline{\mathbf{a}}_K)$ indexed by $\boldsymbol{\beta}_K \in \mathcal{B}_K \subseteq \mathbb{R}^{P_{Q_K}}$. Obtain an estimator $\widehat{\boldsymbol{\beta}}_{K,N}$ for $\boldsymbol{\beta}_K$ by an appropriate regression method, for example, least squares, and take $\widehat{L}_K^{\mathbf{d}}(\overline{\mathbf{x}}_K) = Q_K\{\overline{\mathbf{x}}_K, \overline{\mathbf{d}}_K(\overline{\mathbf{x}}_K); \widehat{\boldsymbol{\beta}}_{K,N}\}$. Define the pseudo-outcomes at stage $k = K-1, \ldots, 1$ as $Q_{k+1}^{\mathbf{d}}[\overline{\mathbf{X}}_{k+1,i}, \{\overline{\mathbf{A}}_{k,i}, d_{k+1}(\overline{\mathbf{X}}_{k+1,i}, \overline{\mathbf{A}}_{k,i})\}; \widehat{\boldsymbol{\beta}}_{k+1,N}]$, the predicted outcome using the fitted model when individuals receive consistent treatments at stage $k+1$. Then, obtain $\widehat{\boldsymbol{\beta}}_{k,N}^{\mathbf{d}}$ by a suitable regression method using the pseudo-outcomes as the response, for example, for least squares,

$$\widehat{\boldsymbol{\beta}}_{k,N}^{\mathbf{d}} = \arg\min_{\boldsymbol{\beta}_k}\sum_{i=1}^{n}\left\{Q_{k+1}^{\mathbf{d}}\left[\overline{\mathbf{X}}_{k+1,i}, \left\{\overline{\mathbf{A}}_{k,i}, d_{k+1}(\overline{\mathbf{X}}_{k+1,i}, \overline{\mathbf{A}}_{k,i})\right\}; \widehat{\boldsymbol{\beta}}_{k+1,N}\right] - Q_k^{\mathbf{d}}\left(\overline{\mathbf{X}}_k, \overline{\mathbf{A}}_k; \boldsymbol{\beta}_k\right)\right\}^2,$$

and $\widehat{L}_k^{\mathbf{d}}(\overline{\mathbf{x}}_k) = Q_k^{\mathbf{d}}\{\overline{\mathbf{x}}_k, \overline{\mathbf{d}}_k(\overline{\mathbf{x}}_k); \widehat{\boldsymbol{\beta}}_{k,N}\}$. For individuals with only one treatment available at stages k to k', we use pseudo-outcome $Q_{k'+1}^{\mathbf{d}}[\overline{\mathbf{X}}_{k'+1,i}, \{\overline{\mathbf{A}}_{k',i}, d_{k'+1}(\overline{\mathbf{X}}_{k'+1,i}, \overline{\mathbf{A}}_{k',i})\}; \widehat{\boldsymbol{\beta}}_{k'+1,N}]$ for $k' < K-1$ and Y for $k' = K-1$ (Tsiatis et al., 2020, Section 6.4.2).

The estimator (3) is doubly robust, that is, it consistently estimates $V(\mathbf{d})$ if either of the sets of models $\pi_k(a_k, \mathbf{h}_k; \theta_k)$, $k = 1, \ldots, K$, or $Q_K(\overline{\mathbf{x}}_K, \overline{\mathbf{a}}_K; \boldsymbol{\beta}_K), Q_k^{\mathbf{d}}(\overline{\mathbf{x}}_k, \overline{\mathbf{a}}_k; \boldsymbol{\beta}_k), k = 1, \ldots, K$, is correctly specified (Han, 2014; Luedtke et al., 2018; Tsiatis et al., 2020; Vermeulen & Vansteelandt, 2015). Consistency is guaranteed in SMARTs because propensities are known.

4 | INTERIM AIPW ESTIMATOR

The interim AIPW estimator (IAIPWE) uses partial information from individuals who have yet to complete follow-up at the times interim analyses are conducted; the IAIPWE includes the IPWE and AIPWE for complete data as special cases. Assume that the enrollment process is independent of all subject information and that the time between stages is fixed, as is the case for many SMARTs. Let S be the number of planned analyses. Let $\Gamma(t) \in \{0, 1\}$ be an indicator that a participant has enrolled in the SMART at study time t, where $t = 0$ denotes the start of the study (in calendar time). In addition, let $\kappa(t) \in \{0, 1, \ldots, K\}$ be the furthest stage reached by a participant at time t with $\Gamma(t) = 0 \Rightarrow \kappa(t) = 0$; and let $\Delta(t) \in \{0, 1\}$ be an indicator that a participant has completed follow-up, that is, they have completed all K stages and have had their outcome ascertained. Thus, the number of participants enrolled at time t is $n(t) = \sum_{i=1}^{N}\Gamma_i(t)$. We evaluate either the fixed set of L embedded regimes $\{\mathbf{d}^\ell\}_{\ell=1}^{L}$ for null hypothesis (1), or the embedded regimes along with a control regime, $\mathbf{d}^0$, for null hypothesis (2). The control regime may be estimated from a separate trial arm or may have a predetermined fixed value. We use superscript ℓ to indicate that a quantity is being computed for regime $\mathbf{d}^\ell$, for example, $\widehat{\boldsymbol{\beta}}_{k,N}^{\ell}$ is shorthand for $\widehat{\boldsymbol{\beta}}_{k,N}^{\mathbf{d}^\ell}$.

We define the "full data" under regime $\mathbf{d}^\ell$ as $W_{\mathbf{d}^\ell}^* = \{Y^*(\mathbf{d}^\ell), \overline{\mathbf{X}}_K^*(\mathbf{d}^\ell)\}$, which comprises the potential outcome $Y^*(\mathbf{d}^\ell)$ and associated potential covariates $\overline{\mathbf{X}}_K^*(\mathbf{d}^\ell) = \{\mathbf{X}_1, \mathbf{X}_2^*(\mathbf{d}^\ell), \ldots, \mathbf{X}_K^*(\mathbf{d}^\ell)\}$. The observed data for an individual at time t are therefore $W(t) = \Gamma(t)[1, \kappa(t), \mathbf{X}_1, A_1, I\{\kappa(t) > 1\}\mathbf{X}_2, I\{\kappa(t) > 1\}A_2, \ldots, I\{\kappa(t) > K-1\}\mathbf{X}_K, I\{\kappa(t) > K-1\}A_K, \Delta(t), \Delta(t)Y]$. For a given time t and regime $\mathbf{d}^\ell$, let $\mathcal{R}^\ell(t) \in \{1, \ldots, 2K, \infty\}$ be a discrete coarsening variable, which is defined as follows:

$$\begin{aligned}
\mathcal{R}^\ell(t) &= 1 && \text{if} \quad A_1 \neq d_1^\ell(\mathbf{H}_1)\\
\mathcal{R}^\ell(t) &= 2 && \text{if} \quad C_{\mathbf{d}^\ell,1} = 1, \kappa(t) = 1\\
\mathcal{R}^\ell(t) &= 3 && \text{if} \quad C_{\mathbf{d}^\ell,1} = 1, \kappa(t) = 2, A_2 \neq d_2^\ell(\mathbf{H}_2)\\
\mathcal{R}^\ell(t) &= 4 && \text{if} \quad C_{\mathbf{d}^\ell,2} = 1, \kappa(t) = 2\\
&\vdots\\
\mathcal{R}^\ell(t) &= 2k-1 && \text{if} \quad C_{\mathbf{d}^\ell,k-1} = 1, \kappa(t) = k, A_k \neq d_k^\ell(\mathbf{H}_k)
\end{aligned}$$

$$\mathcal{R}^{\ell}(t) = 2k \quad \text{if} \quad C_{\mathbf{d}^{\ell},k} = 1, \kappa(t) = k$$
$$\vdots$$
$$\mathcal{R}^{\ell}(t) = 2K \quad \text{if} \quad C_{\mathbf{d}^{\ell},K} = 1, \kappa(t) = K, \Delta(t) \neq 1$$
$$\mathcal{R}^{\ell}(t) = \infty \quad \text{if} \quad C_{\mathbf{d}^{\ell},K} = 1, \kappa(t) = K, \Delta(t) = 1.$$

Thus, $\mathcal{R}^{\ell}(t) = \infty$ corresponds to a participant having completed follow-up and being consistent with $\mathbf{d}^{\ell}$ for all treatment decisions at time t. For $\mathcal{R}^{\ell}(t) < \infty$, $\lfloor \mathcal{R}^{\ell}(t)/2 \rfloor$ is the number of stages at which a participant is consistent with $\mathbf{d}^{\ell}$ at time t, and $\mathcal{R}^{\ell}(t) \mod 2$ encodes whether the number of consistent stages is due to time-related censoring, that is, not having yet completed the current stage, or having been assigned a treatment that is inconsistent with $\mathbf{d}^{\ell}$. See Web Appendix A for an example of how $\mathcal{R}^{\ell}(t)$ is determined.

The observed data $W(t)$ are a coarsened version of the full data $W^{*}_{\mathbf{d}^{\ell}}$. The coarsening is monotone in that the full data coarsened to level $\mathcal{R}^{\ell}(t) = r$ are a many-to-one function of the full data coarsened to level $\mathcal{R}^{\ell}(t) = r + 1$ at time t. Moreover, the data are coarsened at random, as $P\{\mathcal{R}^{\ell}(t) = r | W^{*}_{\mathbf{d}^{\ell}}\} = P\{\mathcal{R}^{\ell}(t) = r | W(t)\}$ (Tsiatis, 2006b, Chap. 7; Zhang et al., 2013), which follows from the consistency and sequential randomization assumptions in Section 2. Define the coarsening hazard function $\lambda^{\ell}_{r}(t) = P\{\mathcal{R}^{\ell}(t) = r | \mathcal{R}^{\ell}(t) \geq r, W(t)\}$ to be the conditional probability that an individual is coarsened to level r given that they are at risk of being coarsened. Because the data are coarsened at random, $\lambda^{\ell}_{r}(t)$ is a function of the observed data. Let the probability that an individual is coarsened after r be $K^{\ell}_{r}(t) = P\{\mathcal{R}^{\ell}(t) > r | W(t)\}$, which is also a function of the observed data. Let $\widehat{K}^{\ell}_{r,n(t)}(t)$ be an estimator of $K^{\ell}_{r}(t)$. Let $\nu_k(t) = P\{\kappa(t) \geq k | \Gamma(t) = 1\}$, $k = 1, \ldots, K$; $\nu_{K+1}(t) = P\{\Delta(t) = 1 | \Gamma(t) = 1\}$; and $k(r)$ map coarsening level r to corresponding decision k. We can express both $\lambda^{\ell}_{r}(t)$ and $K^{\ell}_{r}(t)$ in terms of propensities $\pi_k(A_k, \mathbf{H}_k)$ and $\nu_k(t)$ for $k = 1, \ldots, K+1$. For $\mathcal{R}^{\ell}(t) = r$, r odd, $\lambda^{\ell}_{r}(t) = \pi_{k(r)}(A_{k(r)}, \mathbf{H}_{k(r)})^{1 - C_{d^{\ell},k(r)}} \{1 - \pi_k(A_{k(r)}, \mathbf{H}_{k(r)})\}^{C_{d^{\ell},k(r)}}$ and $K^{\ell}_{r}(t) = \nu_{k(r)}(t)\pi_{k(r)}(A_{k(r)}, \mathbf{H}_{k(r)})^{C_{d^{\ell},k(r)}} \{1 - \pi_k(A_{k(r)}, \mathbf{H}_{k(r)})\}^{1 - C_{d^{\ell},k(r)}} \prod_{v=1}^{k(r)-1} \pi_v(A_v, \mathbf{H}_v)$. For $\mathcal{R}^{\ell}(t) = r$, r even, $\lambda^{\ell}_{r}(t) = \{\nu_{k(r)}(t) - \nu_{k(r)+1}(t)\}/\nu_{k(r)}(t)$ and $K^{\ell}_{r}(t) = \nu_{k(r)+1}(t) \prod_{v=1}^{k(r)} \pi_v(A_v, \mathbf{H}_v)$. It is straightforward to posit models for $\pi_k(A_k, \mathbf{H}_k)$ or $\nu_k(t)$ using logistic regression or simple averages and estimate $\lambda^{\ell}_{r}(t)$ and $K^{\ell}_{r}(t)$. The form of the IAIPWE for regime $\mathbf{d}^{\ell}$ at time t is

$$\widehat{V}^{\ell}_{\text{IA}}(t) = n(t)^{-1} \sum_{i=1}^{N} \Gamma_i(t) \left[\frac{I\{\mathcal{R}^{\ell}_{i}(t) = \infty\}}{\widehat{K}^{\ell}_{2K,i,n(t)}(t)} Y_i \right.$$
$$\left. + \sum_{r=1}^{2K} \frac{I\{\mathcal{R}^{\ell}_{i}(t) = r\} - \widehat{\lambda}^{\ell}_{r,i}(t) I\{\mathcal{R}^{\ell}_{i}(t) \geq r\}}{\widehat{K}^{\ell}_{r,i,n(t)}(t)} L^{\ell}_{k(r)}\left(\overline{\mathbf{X}}_{k(r),i}\right) \right], \quad (4)$$

where $L^{\ell}_{k(r)}(\overline{\mathbf{x}}_{k(r)})$ is an arbitrary function of $\overline{\mathbf{x}}_{k(r)}$. The estimator is doubly robust and thus guaranteed to be consistent in a SMART with a specified enrollment process. We include a proof in Web Appendix B. Similar to the AIPWE, we estimate the efficient choice of unknown functions $L^{\ell}_{k(r)}(\overline{\mathbf{x}}_{k}) = \mathbb{E}\{Y^{*}(\mathbf{d}) | \overline{\mathbf{X}}_{k} = \overline{\mathbf{x}}_{k}, \overline{\mathbf{A}}_{k-1} = \overline{\mathbf{d}}_{k-1}(\overline{\mathbf{x}}_{k-1})\}$ using Q-learning; however, because the IAIPWE uses individuals with incomplete treatment trajectories, the Q-learning procedure for $L_{k(r)}(\overline{\mathbf{x}}_{k(r)})$ is more complicated. Posit a model $Q^{\ell}_{K}\{\overline{\mathbf{x}}_{K}, \overline{\mathbf{a}}_{K}; \boldsymbol{\beta}_{K}(t)\}$ for $Q_K(\overline{\mathbf{x}}_{K}, \overline{\mathbf{a}}_{K}) = \mathbb{E}(Y | \overline{\mathbf{X}}_{K} = \overline{\mathbf{x}}_{K}, \overline{\mathbf{A}}_{K} = \overline{\mathbf{a}}_{K})$ indexed by $\boldsymbol{\beta}_{K}(t) \in \mathcal{B}_{K} \subseteq \mathbb{R}^{P_{Q^{\ell}_{K}}}$. Construct an estimator $\widehat{\boldsymbol{\beta}}_{K}$ for $\boldsymbol{\beta}_{K}$ by an appropriate regression method, for example, least squares, using only individuals who have completed all treatment stages, that is, $\Delta(t) = 1$, and subsequently take $\widehat{L}^{\ell}_{K}(\overline{\mathbf{x}}_{K}) = Q^{\ell}_{K}\{\overline{\mathbf{x}}_{K}, \overline{\mathbf{d}}_{K}(\overline{\mathbf{x}}_{K}); \widehat{\boldsymbol{\beta}}_{K}\}$. Posit models $Q^{\ell}_{k(r)}(\overline{\mathbf{x}}_{k(r)}, \overline{\mathbf{a}}_{k(r)}; \boldsymbol{\beta}_{k(r)})$ for $Q^{\ell}_{k(r)}(\overline{\mathbf{x}}_{k(r)}, \overline{\mathbf{a}}_{k(r)}) = \mathbb{E}\left(Q^{\ell}_{k(r)+1}\left[\overline{\mathbf{X}}_{k(r)+1}, \left\{\overline{\mathbf{A}}_{k(r)}, d_{k(r)+1}(\overline{\mathbf{X}}_{k(r)+1}, \overline{\mathbf{A}}_{k(r)})\right\}\right] | \overline{\mathbf{X}}_{k(r)} = \overline{\mathbf{x}}_{k(r)}, \overline{\mathbf{A}}_{k(r)} = \overline{\mathbf{a}}_{k(r)}\right)$ for $k(r) = K-1, \ldots, 1$. Estimating $\widehat{\boldsymbol{\beta}}_{k(r)}$ requires pseudo-outcomes, which may be missing when using individuals who have been observed through stage $k(r)+1$, that is, $\kappa(t) > k(r)$, but have no observed outcome Y or estimable pseudo-outcome from stages $k(r)+2$ or later. In such cases, we define the pseudo-outcomes for estimating $\widehat{\boldsymbol{\beta}}_{k(r)}$ as

$$\begin{aligned}
&\tilde{Q}^{\ell}_{k(r)+1}(\overline{\mathbf{x}}_{k(r)+1}, \overline{\mathbf{a}}_{k(r)+1}; \widehat{\boldsymbol{\beta}}_{k(r)+1}, \ldots, \widehat{\boldsymbol{\beta}}_{K}) \\
&= \left[I\{|\Psi_{k(r)+1}(\mathbf{h}_{k(r)+1})| \neq 1\} \right. \\
&\left. + I\{\kappa(t) = k(r)+1, \Delta(t) = 0, |\Psi_{k(r)+1}(\mathbf{h}_{k(r)+1})| = 1\} \right] \\
&\quad Q^{\ell}_{k(r)+1}(\overline{\mathbf{x}}_{k(r)+1}, \overline{\mathbf{a}}_{k(r)+1}; \widehat{\boldsymbol{\beta}}_{k(r)+1}) \\
&+ I\{\kappa(t) = k(r)+2, \Delta(t) = 0, |\Psi_{k(r)+1}(\mathbf{h}_{k(r)+1})| = 1\} \\
&\quad Q^{\ell}_{k(r)+2}(\overline{\mathbf{x}}_{k(r)+2}, \overline{\mathbf{a}}_{k(r)+2}; \widehat{\boldsymbol{\beta}}_{k(r)+2}) \\
&+ \cdots + I\{\Delta(t) = 1, |\Psi_{k(r)+1}(\mathbf{h}_{k(r)+1})| = 1\} Y.
\end{aligned}$$

This approach uses individuals with incomplete information to fit the Q-functions for greater efficiency. When all observed individuals have completed their regimes, this strategy is equivalent to the pseudo-outcome method outlined in Section 3. We obtain $\widehat{\boldsymbol{\beta}}^{\ell}_{k(r)}$ by a suitable regression method, using $Q^{\ell}_{k(r)+1}(\overline{\mathbf{x}}_{k(r)+1}, \overline{\mathbf{a}}_{k(r)+1}; \widehat{\boldsymbol{\beta}}_{k(r)+1})$ with $\tilde{Q}^{\ell}_{k(r)+1}(\overline{\mathbf{x}}_{k(r)+1}, \overline{\mathbf{a}}_{k(r)+1}; \widehat{\boldsymbol{\beta}}_{k(r)+1}, \ldots, \widehat{\boldsymbol{\beta}}_{K})$ when necessary, and $L^{\ell}_{k(r)}(\overline{\mathbf{x}}_{k(r)}) = Q^{\ell}_{k(r)}(\overline{\mathbf{x}}_{k(r)}, \overline{\mathbf{a}}_{k(r)}; \widehat{\boldsymbol{\beta}}_{k(r)})$.

To make clear the connection between the IAIPWE and the (A)IPWE, we express $\widehat{V}^{\ell}_{\mathrm{IA}}(t)$ in (4) in an alternate form. For definiteness, consider $K = 2$ decisions at fixed times, and let $\nu_2(t)$ and $\nu_3(t)$ be estimated by $\widehat{\nu}_{2,n(t)}(t) = \sum_{i=1}^{N} I\{\kappa_i(t) = 2\}/\sum_{i=1}^{N} \Gamma_i(t)$, and $\widehat{\nu}_{3,n(t)}(t) = \sum_{i=1}^{N} \Delta_i(t)/\sum_{i=1}^{N} \Gamma_i(t)$. It is shown in Web Appendix A that in this case (4) is equivalent to

$$\begin{aligned}\widehat{V}^{\ell}_{\mathrm{IA}}(t) = n(t)^{-1} \sum_{i=1}^{N} \Gamma_i(t) \\ \Bigg(\frac{\Delta_i(t) C^{\ell}_{2,i} Y_i}{\pi_1(A_{1,i}, \mathbf{H}_{1,i}; \widehat{\theta}_{1,n(t)}) \pi_2(A_{2,i}, \mathbf{H}_{2,i}; \widehat{\theta}_{2,n(t)}) \widehat{\nu}_{3,n(t)}(t)} \\ - \left[\frac{I\{A_{1i} = d^{\ell}_1(\mathbf{H}_{1i})\} I\{\kappa_i(t) = 2\}}{\pi_1(A_{1,i}, \mathbf{H}_{1,i}; \widehat{\theta}_{1,n(t)}) \widehat{\nu}_{2,n(t)}(t)} - 1 \right] L^{\ell}_1(\overline{\mathbf{X}}_{1i}) \\ - \frac{I\{A_{1i} = d^{\ell}_1(\mathbf{H}_{1i})\} I\{\kappa_i(t) = 2\}}{\pi_1(A_{1,i}, \mathbf{H}_{1,i}; \widehat{\theta}_{1,n(t)}) \widehat{\nu}_{2,n(t)}(t)} \\ \left[\frac{I\{A_{2i} = d^{\ell}_2(\mathbf{H}_{2i})\} \Delta_i(t) \widehat{\nu}_{2,n(t)}(t)}{\pi_2(A_{2,i}, \mathbf{H}_{2,i}; \widehat{\theta}_{2,n(t)}) \widehat{\nu}_{3,n(t)}(t)} - 1 \right] L^{\ell}_2(\overline{\mathbf{X}}_{2i}) \Bigg). \end{aligned} \tag{5}$$

If $\Gamma_i(t) = \Delta_i(t) = 1$ for all i, so that $n(t) = N$, as at the time of the final analysis, (5) reduces to the AIPWE (3) with $K = 2$. The augmentation terms in (4) use partial information from participants who are enrolled at the time of an interim analysis but who do not yet have complete follow-up. In contrast, the IPWE (obtained by setting $L^{\ell}_k(\overline{\mathbf{X}}_{ki}) \equiv 0$, $k = 1, 2, \ldots, K$) uses data only from those subjects who are consistent with the regime under consideration at all stages of the study and who have completed the trial. The AIPWE (3) also uses information only from subjects who have completed the trial, but it additionally uses a series of regression models, one for each stage, to impute information for subjects who are not consistent with the regime under consideration starting from the stage at which their treatment first deviates from the regime. The IAIPWE (4) furthermore uses data from all subjects in fitting the regression models in the AIPWE and thereby uses more information and further improves efficiency.

As our goal is to use the IAIPWE for interim monitoring and analyses, we need to characterize its sampling distribution. The following result shows that the IAIPWE for the embedded regimes is asymptotically normal; we use this result to construct tests and decision boundaries in subsequent sections. A proof is given in the Web Appendix C.

Theorem 1. *Let* $\widehat{\mathcal{V}}(t) = \{\widehat{V}^{0}_{\mathrm{IA}}(t), \widehat{V}^{1}_{\mathrm{IA}}(t), \ldots, \widehat{V}^{L}_{\mathrm{IA}}(t)\}^{\top}$ *be the stacked value estimators at time t across all regimes, and let* $n(t)/N \xrightarrow{p} c$, *a constant. Under standard regularity conditions stated in the Web Appendix C,* $N^{1/2}\{\widehat{\mathcal{V}}(t) - \mathcal{V}(t)\} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$ *as* $N \to \infty$.

A consistent estimator $\widehat{\boldsymbol{\Sigma}}$ of $\boldsymbol{\Sigma}$ can be obtained using the sandwich estimator or the bootstrap. Comparisons among the $L + 1$ regimes can be constructed using a contrast vector and are asymptotically normal via a simple Taylor series argument (see Web Appendix C). When there is no control regime, $\mathcal{V}(t)$ is indexed only by $\ell = 1, \ldots, L$.

5 | INTERIM ANALYSIS FOR SMARTS

5.1 | Hypothesis testing

For simplicity, consider $S = 2$ planned analyses at study times t_1 (interim analysis) and t_2 (final analysis). We present the extension to an arbitrary S in Web Appendix D. We discuss the interim analysis procedure in the context of superiority; the procedure for homogeneity follows under minor modifications.

Define the test statistics at analysis time t_s,

$$Z^{\ell}(t_s) = \{\widehat{V}^{\ell}_{\mathrm{IA}}(t_s) - \widehat{V}^{0}_{\mathrm{IA}}(t_s)\}/\mathrm{SE}\{\widehat{V}^{\ell}_{\mathrm{IA}}(t_s) - \widehat{V}^{0}_{\mathrm{IA}}(t_s)\}, \quad \ell = 1, \ldots, L,$$

where $\widehat{V}^{0}_{\mathrm{IA}}(t_s)$ can be estimated as the sample average of response Y_i for individuals receiving $\mathbf{d}^0$ and the denominator is obtained from the approximate normal sampling distribution for $\widehat{\mathcal{V}}(t_s)$ in Theorem 1. If regime means are compared to a fixed control value V^0, replace $\widehat{V}^{0}_{\mathrm{IA}}(t_s)$ by V^0. At each analysis s, we propose to stop the trial if any test statistic exceeds a stopping boundary $c_s(\alpha)$, which will be discussed in the next section. Heuristically, the testing procedure at significance level α across all t_s is as follows:

(1) At time t_1, compute $Z^{\ell}(t_1)$, $\ell = 1, \ldots, L$. If $Z^{\ell}(t_1) > c_{\alpha}(1)$, for any ℓ, reject H_0 and terminate the trial; else, continue the trial.
(2) At time t_2, compute $Z^{\ell}(t_2)$, $\ell = 1, \ldots, L$. If $Z^{\ell}(t_2) > c_{\alpha}(2)$ for any ℓ. reject H_0; otherwise, fail to reject H_0. Terminate the trial.

A trial with more than two planned analysis repeats step (1) for all interim analyses, terminating when a test statistic is greater than the corresponding stopping boundary.

This formulation can be adapted to any set of hypotheses involving functions of the values of regimes of interest. For example, testing the homogeneity hypothesis (1) would involve calculation of chi-square test statistics based on the distributions of $\widehat{\mathcal{V}}(t_s)$, $s = 1, 2$, analogous to Wu

et al. (2021), which would be compared to corresponding stopping boundaries.

5.2 | Stopping boundaries

We discuss boundary selection and sample size calculations for superiority null hypothesis (2), which involves multiple comparisons of embedded regimes against a control regime. We seek to determine stopping boundaries $c_\alpha(s)$, $s = 1, 2$, that control the familywise error rate across all planned analyses at level α; that is,

$$P\{\text{Reject } H_{0D} | H_{0D} \text{ is true}\} = P\left[\bigcup_{\ell=1}^{L}\bigcup_{s=1}^{2}\{Z^\ell(t_s) \geq c_\alpha(s)\}\middle| H_{0D}\right] \leq \alpha. \quad (6)$$

Common approaches to calculating boundaries that satisfy (6) include the Pocock boundary, which takes $c_\alpha(s) = c_\alpha$ for some c_α for $s = 1, 2$ (Pocock, 1977); the O'Brien-Fleming (OBF) boundary $\{c_\alpha(1), c_\alpha(2)\} = \{\iota c_\alpha, c_\alpha\}$ (O'Brien & Fleming, 1979), where ι is the reciprocal of the square root of the statistical information (e.g., inverse of the variance of the numerator of the associated Z-score) available at analysis s divided by the statistical information available at final analysis S; or the broader α-spending approach (DeMets & Lan, 1994). If the information proportion between the interim and final analysis varies by regimes, practitioners may elect to use a regime-dependent ι^ℓ in the spirit of OBF. For a detailed discussion about if and when each boundary type might be preferable, see Jennison and Turnbull (2000).

Define the stacked vector of sequential test statistics

$$\mathbf{Z} = \{Z^1(t_1), \ldots, Z^L(t_1), Z^1(t_2), \ldots, Z^L(t_2)\}^\top. \quad (7)$$

Boundaries that satisfy (6) can be obtained via the joint cumulative distribution function of $\mathbf{Z}$ under null hypothesis (2).

Theorem 2. *Under the null hypothesis (2) and* $n(t)/N \xrightarrow{p} c$, *a constant, the test statistics* $\mathbf{Z}$ *satisfy* $\mathbf{Z} \xrightarrow{d} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{H_0})$ *where* $\boldsymbol{\Sigma}_{H_0}$ *is a block diagonal matrix with diagonal entries* $\text{corr}\{Z^1(t_s), \ldots, Z^L(t_s)\}$ *and off-diagonal entries* $\iota^{-1}\text{corr}\{Z^1(t_s), \ldots, Z^L(t_s)\}$, ι *is the reciprocal of the information proportion between interim analysis* s *and final analysis* S.

A proof of Theorem 2 and discussion on calculating ι and the correlation between the Z-statistics are provided in the Web Appendix E. In practice, computation of c_α can be done numerically. Either the correlation of the test statistics or the variance of all components of the estimator must be specified to compute the stopping boundaries. We approximate c_α through integration of the corresponding multivariate normal distribution of $\mathbf{Z}$. Under the information monitoring approach (Tsiatis, 2006a), the correlation between sequential test statistics for the same regime simplifies to the square root of the ratio of the information available between the two time points. Because of incomplete information for participants enrolled but who have not yet completed the trial, this quantity does not simplify to the square root of the ratio of the interim sample size to the final planned sample size. The off-diagonal elements of the covariance matrix, $\boldsymbol{\Sigma}_{H_0}$, may be non-zero for overlapping embedded regimes. For these reasons, it may be difficult to specify $\boldsymbol{\Sigma}_{H_0}$. An alternative is to specify generative models for the observed data, that is, a mean model and distributions for associated covariates, propensities, and enrollment at time of interim analyses, and estimate the correlation structure empirically via simulation.

The choice of the models and estimators for $\lambda_r^\ell(t)$, $K_r^\ell(t)$, and $L_{k(r)}^\ell(\overline{\mathbf{x}}_{k(r)})$ impact the correlation structure of $\boldsymbol{\Sigma}_{H_0}$ and can result in correlated value estimators across nonoverlapping embedded regimes; that is, regimes that involve different stage 1 treatment options. If cohorts enroll sequentially and interim analyses are planned such that all enrollment occurs within each cohort (i.e., $\Delta_i(t_s) = \Gamma_i(t_s)$ for all i for $s = 1, 2$), then the test statistics at each analysis use the standard AIPWE (3) computed using data from all participants who have entered the trial. Therefore, stopping boundaries for trials with such enrollment procedures are subsumed by this method.

5.3 | Power and sample size

With stopping boundaries $\mathbf{c}_\alpha = \{c_\alpha(1)\mathbf{1}_L, c_\alpha(2)\mathbf{1}_L\} \in \mathbb{R}^{2L}$ for $\mathbf{1}_L$ an L-vector of ones, and specified alternative H_A, the power of the testing procedure is

$$P\{\text{Reject } H_{0D} | H_A \text{ is true}\} = P\left[\bigcup_{\ell=1}^{L}\bigcup_{s=1}^{2}\{Z^\ell(t_s) \geq c_\alpha^\ell(s)\}\middle| H_A\right] = 1 - \beta.$$

The power of the test under H_A, where $\mathbf{Z}$ has expected value $\boldsymbol{\mu}_A = \boldsymbol{\mu}_A\{n(t_1), n(t_2)\}$ and covariance $\boldsymbol{\Sigma}_{H_A}$, is approximately

$$1 - \int \cdots \int_D \frac{1}{(2\pi)^L \det(\boldsymbol{\Sigma}_{H_A})^{-1/2}} \exp\{-\frac{1}{2}(\mathbf{Z} - \boldsymbol{\mu}_A)^\top \boldsymbol{\Sigma}_{H_A}^{-1}(\mathbf{Z} - \boldsymbol{\mu}_A)\} dz^1(1) dz^2(1) \cdots dz^L(2) \quad (8)$$

for domain $D = (-\infty, c_\alpha(1)] \times (-\infty, c_\alpha(1)] \times \cdots \times (-\infty, c_\alpha(2)]$. As the mean under the alternative, $\boldsymbol{\mu}_A$, is a function of the sample size, so too is (8). Thus, to

achieve nominal power $100(1-\beta)\%$, one can set (8) equal to $1-\beta$ and solve for the sample size. Although our results hold for a general alternative hypothesis H_A, we proceed under the simplifying assumption that $\boldsymbol{\Sigma}_{H_0} = \boldsymbol{\Sigma}_{H_A}$, that is, that the covariance is the same under H_{0D} and H_A. In our implementation, we use a grid search for a fixed enrollment process and ratio between interim sample sizes to find the total planned sample size N that attains the desired power. When the augmentation terms are zero, the analyst must specify the correlation among estimators of the regimes, the information proportion for analyses, the alternative mean outcomes, and the variance of the mean outcomes. When augmentation terms are nonzero, all generative models must be specified to determine the sample size and corresponding power.

Specification of all generative models required for the IAIPWE at the design stage may be challenging. Accordingly, a practical strategy would be to power the trial and thus determine N conservatively based on the IPWE but base interim analyses on the more efficient IAIPWE, which can lead to increased power and smaller expected sample size.

As previously stated, the covariance structure, $\boldsymbol{\Sigma}_{H_0}$, depends on the enrollment process through the information proportion at the time of analysis. Thus, one can compute the maximum power for a fixed sample size under differing enrollment processes using (8) adjusted for the differences in the information proportion at the time of the analysis. One can also consider other objectives such as minimizing the time to decision or the cost of the trial using these same procedures.

5.4 | Test for homogeneity

Exploiting the previous developments, we formulate a sequential testing procedure using $\mathbf{Z}(t_s) = \{Z^1(t_s), \dots, Z^L(t_s)\}$ for the global null hypothesis (1), that is, that all regimes are equal. We derive a χ^2-statistic using Theorem 2. Let $\mathbf{C} = [\mathbf{I}_{L-1} | -\mathbf{1}_{L-1}]$ where $\mathbf{I}_q$ is the $(q \times q)$ identity matrix and $\mathbf{1}_q$ a q-vector of ones. Let $\boldsymbol{\Sigma}_{H_0}(t_s)$ be the $(L \times L)$ submatrix of $\boldsymbol{\Sigma}_{H_0}$ corresponding to the covariance of $\mathbf{Z}(t_s)$, and let $\boldsymbol{\mu}_A(t_s)$ be the $(L \times 1)$ vector corresponding to the alternative mean at time t_s. The sequential Wald-type test statistic at time t_s is

$$T_{\chi^2,\upsilon}(t_s) = \mathbf{Z}^\top(t_s)\mathbf{C}^\top\{\mathbf{C}\boldsymbol{\Sigma}_{H_0}(t_s)\mathbf{C}^\top\}^{-}\mathbf{C}\mathbf{Z}(t_s), \tag{9}$$

which follows a χ^2 distribution with degrees of freedom $\upsilon = \text{rank}\{\mathbf{C}\boldsymbol{\Sigma}_{H_0}(t_s)\mathbf{C}^\top\}$ and noncentrality parameter $\phi_A = \boldsymbol{\mu}_A(t_s)^\top\{\mathbf{C}\boldsymbol{\Sigma}_{H_0}(t_s)\mathbf{C}^\top\}^{-}\boldsymbol{\mu}_A(t_s)$. Following the methods in previous sections, the stopping boundaries now come from a χ^2 distribution. Using simulation, we estimate the stopping boundaries using the correlation structure of $\mathbf{Z}$ such that $\{c_\alpha(1), c_\alpha(2)\}$ satisfy the type I error rate. The Pocock boundaries still satisfy $c_\alpha(s) = c_\alpha$; however, the OBF type boundaries satisfy $\{c_\alpha(1), c_\alpha(2)\} = \{\iota^2 c_\alpha, c_\alpha\}$ with ι as defined in Section 5.2.

After calculating the stopping boundaries, we use the distribution of $\mathbf{Z}$ for relevant power and sample size calculations. We estimate the total planned sample size required to attain power $1-\beta$ numerically; see Web Appendix F for details on implementation.

6 | SIMULATION EXPERIMENTS

We report on extensive simulations to evaluate the performance of the IAIPWE. In our simulation settings, IPWE corresponds to the proposed method of Wu et al. (2021). We present results here based on 1000 Monte Carlo replications for the schema shown in Figure 1. We evaluate the type I error rates, power, and expected sample sizes for fixed interim analysis times for the null hypothesis H_{0D} in (2) and alternative hypothesis $H_{AD} : V(\mathbf{d}^\ell) - V(\mathbf{d}^0) > \delta$ for at least one ℓ. We also investigate the benefit of leveraging partial information through the IAIPWE over an IPWE in trials with sample size determined by the IPWE. Finally, we consider how the proportion of enrolled individuals having reached different stages of the trial at an interim analysis affects performance. We consider both Pocock and OBF boundaries. We use correctly specified Q-functions for augmented estimators. Web Appendix G includes results for additional schema and settings; the results are qualitatively similar.

We generate data with a dependence between history and outcomes and explore the impact of the enrollment process on interim analyses. We generate two baseline covariates $X_{1,1} \sim \text{Normal}(47.5, 64)$ and $X_{1,2} \sim \text{Bernoulli}(0.5)$ as well as an interim outcome $X_{2,1} \sim \text{Normal}(1.25X_{1,1}, 9)$. We simulate the response status $R_2 \sim \text{Bernoulli}\{\text{expit}(0.01X_{1,1} + 0.02X_{1,2} - 0.008X_{2,1})\}$ where $\text{expit}(u) = e^u/(1+e^u)$. Individuals at stage one and responders at stage two are randomized with equal probability to feasible treatments. The outcome is normally distributed with variance 100 and mean

$$\begin{aligned}
\mu_{S2}(\overline{\mathbf{X}}_2, \overline{\mathbf{A}}_2) = {} & I\{A_1 = 0\}\big\{\beta_0 + \beta_1 X_{1,1} + \beta_2 X_{1,2} + \beta_3 R_2 X_{2,1} \\
& + \beta_4(1-R_2)X_{2,1} + R_2 A_2\big(\beta_5 + \beta_6 X_{1,1} + \beta_7 X_{1,2} + \beta_8 X_{2,1}\big) \\
& + (1-R_2)A_2\big(\beta_9 + \beta_{10} X_{1,1} + \beta_{11} X_{1,2} + \beta_{12} X_{2,1}\big)\big\} \\
& + I\{A_1 = 0\}\big\{\beta_{13} + \beta_{14} X_{1,1} + \beta_{15} X_{1,2} + \beta_{16} R_2 X_{2,1} \\
& + \beta_{17}(1-R_2)X_{2,1} + R_2 A_2\big(\beta_{18} + \beta_{19} X_{1,1} + \beta_{20} X_{1,2} + \beta_{21} X_{2,1}\big) \\
& + (1-R_2)A_2\big(\beta_{22} + \beta_{23} X_{1,1} + \beta_{24} X_{1,2} + \beta_{25} X_{2,1}\big)\big\}.
\end{aligned}$$

In the first scenario, we perform an interim analysis at day 500, and enrollment times are drawn uniformly between 0 and 1000 days with follow-up times every 100 days. We define three value patterns (VPs): (VP1), all embedded regimes have value 47.5; (VP2), regimes $\ell = 1, \dots, 8$ have values $(49.5, 49.5, 49.5, 49.5, 47.5, 47.5, 47.5, 47.5)$; and (VP3), regimes $\ell = 1, \dots, 8$ have values $(50.5, 49.0, 49.0, 47.5, 47.5, 47.5, 47.5, 47.5)$. In each case, $\boldsymbol{\beta}$ is chosen to achieve these VPs. We use the sample size determined to achieve power 80% under a specified VP and estimator. This allows us to investigate the performance of the estimators for different alternatives. In H_{0D} and H_{AD}, $V(\mathbf{d}^0)$ is a fixed control value equal to 47.5 and $\delta = 0$.

Table 1 summarizes the total planned sample size to achieve power 80% under a specified alternative (VP2 for VP1 and VP2, and VP3 for VP3), the proportion of early rejections of null (2), the proportion of total rejections of null (2), the expected sample size, and the expected stopping time. Results are given for both the total sample size to achieve the desired power for each individual estimator (a) and for the total sample size for the IPWE to achieve the desired power (b). The slight differences among the total planned sample size in (a) and (b) are due to Monte Carlo error. All estimators achieve nominal power and type I error rate. The IAIPWE requires a smaller total planned sample size to achieve nominal power. The IAIPWE also exhibits the highest early rejection rate under true alternatives demonstrating the efficiency gain from the augmentation terms and therefore lower expected sample sizes and earlier expected stopping times. The AIPWE slightly underperforms relative to the IPWE due to the overestimation of variance using the sandwich matrix for small $n(t_1)$. It is well known that the performance of the sandwich matrix can deteriorate for small samples. As such, alternative estimation of the covariance matrix, such as using the empirical bootstrap, can be used. The IAIPWE is less affected by overestimation of the variance than the AIPWE. When the total sample size is selected based on the IPWE and an augmented estimator is used, the type I error rate is controlled and the study achieves a higher power.

Table 2 summarizes estimation performance at the interim and final analyses, where a mean square error (MSE) greater than one implies that the indicated estimator is more efficient than the IPWE. The estimators are all consistent as expected. Both the AIPWE and IAIPWE are more efficient than the IPWE at both analyses, and the IAIPWE is more efficient than the AIPWE. At the interim analysis, the standard errors for the IPWE underestimate the sampling variation in most cases, whereas the standard errors for the AIPWE overestimate the sampling variation. The IAIPWE consistently estimates the sampling variation with the exception of regime 6 at the interim analysis.

In the second scenario, we investigate how different enrollment processes affect the proportion of early rejections for hypothesis (2) with $S = 2$ analyses. To vary the rate of enrollment, we select in which of four time periods ([0,500], [501,600], [601,700], and [701,1000]) an individual enrolls using a multinomial distribution. Within each, individuals enroll uniformly. Results for the Pocock stopping boundaries under (VP2) are given in Table 3. The sample sizes are determined to achieve 80% power under (VP2), and the interim analysis is conducted on day 700. Both the total planned and expected sample sizes are lower for the IAIPWE than the IPWE or AIPWE. The proportion of early rejections is higher when more individuals have progressed further through the study due to the increased information available at the time of analysis. All methods attain the desired power, and the IAIPWE achieves earlier expected stopping times and lower expected sample sizes than the IPWE and AIPWE.

In Web Appendix G, we present results for two additional, common designs: the schema in Figure 1 with a control arm and a schema in which responders are not rerandomized. The additional simulations demonstrate that the IAIPWE performs well even under misspecification of the Q-functions. In small samples, the IAIPWE variance may be overestimated, resulting in the estimated proportion of information at interim analyses being inflated. The OBF boundaries may be conservative in these cases. The IAIPWE performs well with multiple interim analyses and for the χ^2 testing procedure for H_{0H}.

7 | CASE STUDY: CANCER PAIN MANAGEMENT SMART

We present a case study based on a recently completed trial evaluating behavioral interventions for pain management in breast cancer patients (Kelleher et al., 2017; ClinicalTrials.gov, 2021). A schematic for the trial is shown in Figure 1. Initially, patients are randomized with equal probability to one of two PCST interventions: five sessions with a licensed therapist (PCST-Full) or one 60-min session (PCST-Brief) with a licensed therapist. After 8 weeks (end of stage one), participants who achieve a 30% reduction in pain from baseline are deemed responders and randomized with equal probability to maintenance therapy or no further intervention. Nonresponders who received PCST-Full are randomized with equal probability to either two full sessions (PCST-Plus) or maintenance. Nonresponders who received PCST-Brief are randomized with equal probability to PCST-Full or maintenance. The eight embedded regimes are given in Figure 1. Follow-up occurs 8 weeks after administration of stage two intervention and again 6 months later. Here, we take the outcome of interest to

TABLE 1 For the schema in Figure 1, interim analysis performance results for testing hypothesis (2) against H_{AD} with a fixed control value using Pocock and O'Brien-Fleming (OBF) boundaries. VP indicates the true value pattern. Method indicates the estimator used. The total planned sample size N is determined by either each method (a) or by the inverse probability weighted estimator (IPWE) (b). Total planned sample sizes are determined to maintain a nominal type I error rate of $\alpha = 0.05$ and achieve a power of 80% under the respective value patterns, using alternative (VP2) to determine the sample size for the null (VP1). Early reject and Total reject are the rejection rates at the first analysis and for the overall procedure, respectively. $\mathbb{E}$(SS) is the expected sample size, that is, the average number of individuals enrolled in the trial when the trial is completed. $\mathbb{E}$(Stop) is the expected stopping time, that is, the average number of days that the trial ran. Monte Carlo standard deviations are given in parentheses.

		(a) N based on Method					(b) N based on IPWE				
VP	**Method**	N	**Early reject**	**Total reject**	$\mathbb{E}$**(SS)**	$\mathbb{E}$**(Stop)**	N	**Early reject**	**Total reject**	$\mathbb{E}$**(SS)**	$\mathbb{E}$**(Stop)**
1	IPWE	1049		0.076	1049 (0)	1199 (1)	1051		0.059	1051 (0)	1199 (1)
1	AIPWE	758		0.042	758 (0)	1199 (1)	1051		0.049	1051 (0)	1199 (1)
2	IPWE	1049		0.795	1049 (0)	1199 (1)	1051		0.781	1051 (0)	1199 (1)
2	AIPWE	758		0.814	758 (0)	1199 (1)	1051		0.908	1051 (0)	1199 (1)
3	IPWE	873		0.833	873 (0)	1199 (1)	873		0.833	873 (0)	1199 (1)
3	AIPWE	586		0.801	586 (0)	1198 (2)	873		0.953	873 (0)	1199 (1)
Pocock											
1	IPWE	1212	0.040	0.072	1188 (119)	1171 (137)	1213	0.041	0.071	1188 (120)	1171 (138)
1	AIPWE	872	0.024	0.042	861 (67)	1182 (107)	1213	0.030	0.043	1195 (103)	1178 (119)
1	IAIPWE	869	0.032	0.049	855 (77)	1177 (123)	1213	0.037	0.059	1191 (112)	1174 (130)
2	IPWE	1212	0.321	0.799	1017 (283)	975 (327)	1213	0.299	0.797	1032 (277)	990 (320)
2	AIPWE	872	0.256	0.800	760 (191)	1020 (305)	1213	0.339	0.912	1008 (286)	962 (331)
2	IAIPWE	869	0.322	0.801	729 (203)	974 (327)	1213	0.399	0.915	972 (296)	920 (343)
3	IPWE	987	0.297	0.835	841 (225)	984 (323)	987	0.297	0.835	841 (225)	984 (323)
3	AIPWE	663	0.236	0.825	585 (140)	1034 (297)	987	0.364	0.950	808 (237)	945 (337)
3	IAIPWE	660	0.290	0.822	565 (149)	996 (317)	987	0.434	0.953	774 (244)	896 (347)
O'Brien-Fleming											
1	IPWE	1052	0.000	0.071	1052 (0)	1199 (1)	1051	0.000	0.059	1051 (0)	1199 (1)
1	AIPWE	758	0.000	0.042	758 (0)	1199 (1)	1051	0.000	0.050	1051 (0)	1199 (1)
1	IAIPWE	756	0.000	0.043	756 (0)	1199 (1)	1051	0.000	0.050	1051 (0)	1199 (1)
2	IPWE	1052	0.008	0.795	1048 (47)	1194 (62)	1051	0.006	0.781	1048 (40)	1195 (54)
2	AIPWE	758	0.002	0.813	757 (17)	1197 (31)	1051	0.005	0.908	1048 (37)	1196 (49)
2	IAIPWE	756	0.014	0.811	751 (44)	1189 (82)	1051	0.013	0.908	1044 (59)	1190 (79)
3	IPWE	873	0.012	0.833	868 (48)	1191 (76)	873	0.012	0.833	868 (48)	1191 (76)
3	AIPWE	585	0.003	0.801	584 (16)	1196 (38)	873	0.004	0.954	871 (28)	1196 (44)
3	IAIPWE	586	0.013	0.802	582 (34)	1189 (79)	873	0.021	0.954	864 (28)	1184 (100)

Abbreviations: IPWE inverse probability weighted estimator; AIPWE, augmented inverse probability weighted estimator; IAIPWE, interim augmented inverse probability weighted estimator.

be percent reduction in pain from baseline at the final 6-month assessment and the primary analysis to be the evaluation of the eight embedded regimes via the null hypothesis H_{0D} in (2) as described below.

Because the data from the trial are not yet published, we simulate the trial based on the protocol. We consider five baseline covariates: height $X_{1,1}$, weight $X_{1,2}$, presence/absence of comorbidities $X_{1,3}$, use of pain medication $X_{1,4}$, and whether or not the participant is receiving chemotherapy $X_{1,5}$. We observe the response status R_2, percent reduction in pain $X_{2,0}$, and degree of adherence $X_{2,1}$ at the first follow-up at the end of stage one. Participants enroll uniformly over 1000 days, the end of stage one occurs 8 weeks after enrollment, and the outcome Y is ascertained 18 weeks after the end of stage one and thus 6 months after enrollment. The distributions of covariates and outcomes are given in Web Appendix H. We take $N = 284$ to match the sample size of Kelleher et al. (2017).

An interim analysis is planned for day 500 and a final analysis at the trial conclusion, a maximum of 1182 days. We test the null hypothesis (2) against the alternative that any regime achieves greater than a 22.5% reduction in pain (fixed control value); see Web Appendix H. We consider both Pocock and OBF boundaries, for which, to achieve

TABLE 2 For the schema in Figure 1, interim analysis performance results for testing hypothesis (2) against H_{AD} with a fixed control value under Pocock Boundaries under (VP2) and sample size N based on the method. MC mean is the Monte Carlo mean of the estimates, MC SD is the Monte Carlo standard deviation of estimates, ASE is the Monte Carlo mean of the standard errors, and MSE ratio is the ratio of the Monte Carlo mean square error for the inverse probability weighted estimator (IPWE) divided by that of the indicated estimator for the tree estimates at the interim analysis (a) and final analysis (b) for $B = 1000$ simulations. The true values under (VP2) for regimes $(1, \dots, 8)$ are $(49.5, 49.5, 49.5, 49.5, 47.5, 47.5, 47.5, 47.5)$.

		(a) Interim analysis				(b) Final analysis			
Method	**Regime**	**MC mean**	**MC SD**	**ASE**	**MSE ratio**	**MC mean**	**MC SD**	**ASE**	**MSE ratio**
IPWE	1	49.47	1.52	1.44	1.00	49.50	0.80	0.79	1.00
IPWE	2	49.52	1.51	1.44	1.00	49.52	0.84	0.79	1.00
IPWE	3	49.48	1.47	1.44	1.00	49.48	0.78	0.79	1.00
IPWE	4	49.53	1.48	1.44	1.00	49.51	0.82	0.79	1.00
IPWE	5	47.51	1.47	1.43	1.00	47.51	0.82	0.79	1.00
IPWE	6	47.55	1.43	1.44	1.00	47.55	0.78	0.79	1.00
IPWE	7	47.48	1.48	1.44	1.00	47.49	0.80	0.79	1.00
IPWE	8	47.52	1.47	1.44	1.00	47.53	0.76	0.79	1.00
AIPWE	1	49.47	1.43	1.48	1.14	49.48	0.76	0.77	1.11
AIPWE	2	49.50	1.42	1.47	1.13	49.51	0.76	0.76	1.22
AIPWE	3	49.49	1.39	1.46	1.13	49.47	0.75	0.77	1.06
AIPWE	4	49.52	1.40	1.48	1.12	49.50	0.75	0.77	1.21
AIPWE	5	47.50	1.39	1.46	1.12	47.49	0.77	0.76	1.12
AIPWE	6	47.53	1.33	1.45	1.16	47.53	0.75	0.76	1.08
AIPWE	7	47.42	1.45	1.45	1.04	47.46	0.78	0.76	1.05
AIPWE	8	47.44	1.41	1.46	1.07	47.51	0.72	0.77	1.13
IAIPWE	1	49.47	1.37	1.38	1.24	49.48	0.76	0.77	1.10
IAIPWE	2	49.50	1.37	1.37	1.21	49.51	0.76	0.77	1.21
IAIPWE	3	49.50	1.33	1.37	1.23	49.47	0.76	0.77	1.05
IAIPWE	4	49.53	1.35	1.38	1.20	49.50	0.75	0.77	1.21
IAIPWE	5	47.51	1.34	1.37	1.20	47.49	0.77	0.77	1.13
IAIPWE	6	47.54	1.27	1.35	1.26	47.53	0.75	0.76	1.08
IAIPWE	7	47.43	1.40	1.36	1.12	47.46	0.79	0.76	1.05
IAIPWE	8	47.46	1.36	1.37	1.16	47.51	0.72	0.77	1.12

Abbreviations: IPWE inverse probability weighted estimator; AIPWE, augmented inverse probability weighted estimator; IAIPWE, interim augmented inverse probability weighted estimator.

a type I error of $\alpha = 0.05$ using our IAIPWE procedure, $c_{\alpha=0.05} = (2.66, 2.66)$ and $(4.20, 2.43)$, respectively. For the AIPWE and IPWE, the Pocock and OBF boundaries are $= (2.66, 2.66)$ and $(4.30, 2.44)$, respectively. In this setting, the correlation structure for $\mathbf{Z}$ is similar for all estimators. Therefore, the Pocock boundaries are the same even with the difference of available information at the interim analysis. As a result, the Pocock boundaries illustrate in part why we expect more early rejections under a true alternative for the IAIPWE than the other estimators. By construction, the different OBF boundaries demonstrate the impact of the increased information available using the IAIPWE at the interim analysis.

The interim analysis occurs at 500 days after the trial enrollment begins, at which point 51.4% of the total planned sample size N has been enrolled, 46.8% of the N planned participants have progressed to the second decision, and 34.9% have completed the trial. Figure 2 summarizes the estimated values for each regime at the time of analysis, corresponding Z-statistic. Exact numbers are recorded in a tabular format in Web Appendix I. Regime 1, which starts with PCST-Full, triggers early stopping based on the test statistic exceeding the OBF boundary. Regimes 1 and 3 trigger early stopping based on test statistics exceeding the Pocock boundary. The standard errors are smaller than those obtained using the IPWE or AIPWE, which are included in the Web Appendix I. The IPWE and AIPWE trigger early stopping with regimes exceeding the Pocock boundary, but fail to trigger early stopping under the OBF boundary. The decision to stop the trial early reduces the sample size from the total possible 284 subjects to 146 and the length of the study by 96 weeks.

TABLE 3 For the schema in Figure 1, interim analysis performance results for testing hypothesis (2) against H_{AD} with a fixed control value using Pocock boundaries under varying enrollments. The interim analysis is conducted on day 700. The percentages p_1, p_2, and p_3 are the expected percentage of individuals to have completed the trial, made it to only stage two, and to have made it to only stage one, respectively. Method indicates the estimator used. The total planned sample size N is determined by either each method. Total planned sample sizes are determined to maintain a nominal type I error rate of $\alpha = 0.05$ and achieve a power of 80% under (VP2). Early reject and Total reject are the rejection rates at the first analysis and for the overall procedure, respectively. $\mathbb{E}$(SS) is the expected sample size, that is, the average number of individuals enrolled in the trial when the trial is completed. $\mathbb{E}$(Stop) is the expected stopping time, that is, the average number of days that the trial ran. Monte Carlo standard deviations are given in parentheses.

p_1	p_2	p_3	Method	N	Early reject	Final reject	$\mathbb{E}$(SS)	$\mathbb{E}$(Stop)
50	10	10	IPWE	1179	0.472	0.802	1012 (177)	964 (249)
50	10	10	AIPWE	844	0.423	0.787	737 (125)	988 (247)
50	10	10	IAIPWE	839	0.472	0.792	720 (126)	963 (249)
40	20	10	IPWE	1190	0.392	0.826	1050 (174)	1003 (243)
40	20	10	AIPWE	856	0.367	0.811	762 (124)	1016 (241)
40	20	10	IAIPWE	851	0.436	0.814	740 (127)	981 (247)
30	30	10	IPWE	1216	0.312	0.811	1102 (170)	1043 (231)
30	30	10	AIPWE	874	0.247	0.799	809 (113)	1076 (215)
30	30	10	IAIPWE	868	0.353	0.802	776 (125)	1023 (239)
40	10	20	IPWE	1194	0.382	0.800	1057 (174)	1009 (243)
40	10	20	AIPWE	859	0.340	0.782	772 (122)	1029 (236)
40	10	20	IAIPWE	855	0.408	0.793	751 (126)	995 (245)
30	20	20	IPWE	1215	0.329	0.812	1095 (171)	1035 (235)
30	20	20	AIPWE	874	0.257	0.806	807 (115)	1071 (218)
30	20	20	IAIPWE	867	0.340	0.811	779 (123)	1029 (236)
30	10	30	IPWE	1213	0.302	0.798	1103 (167)	1048 (229)
30	10	30	AIPWE	873	0.271	0.800	802 (116)	1064 (222)
30	10	30	IAIPWE	870	0.332	0.812	784 (123)	1033 (235)

Abbreviations: IPWE inverse probability weighted estimator; AIPWE, augmented inverse probability weighted estimator; IAIPWE, interim augmented inverse probability weighted estimator.

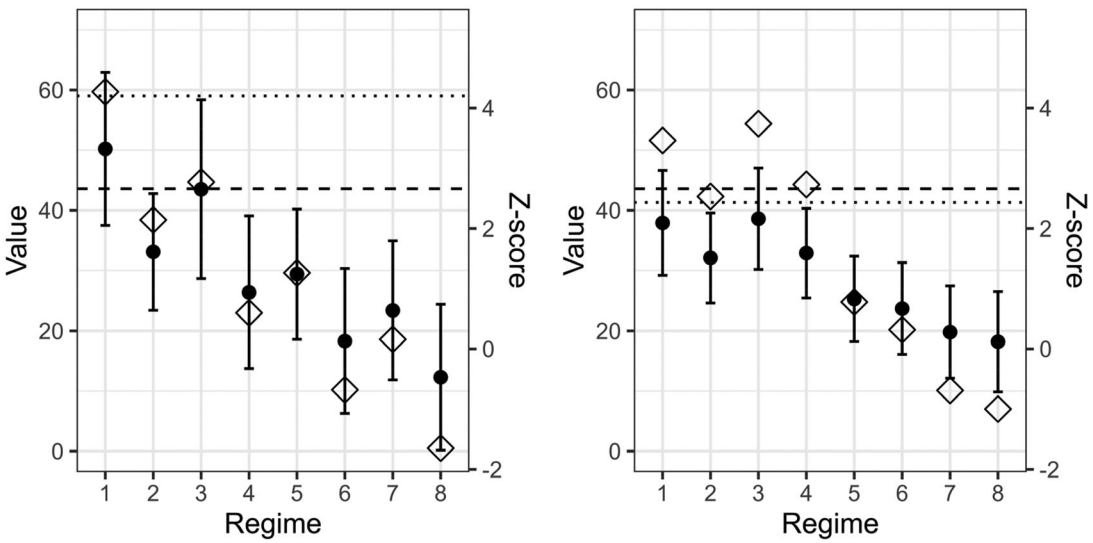

FIGURE 2 For the schema in Figure 1, interim analysis performance results for testing null hypothesis (2) against H_{AD} with a fixed control. Results include the Pocock boundaries (dashed), O'Brien-Fleming (OBF) boundaries (dotted), value estimates (circles) and 95% confidence bounds, and test statistics (rhombus) at the interim (left) and final (right) analysis time for the behavioral pain management case study data set using the interim augmented inverse probability weighted estimator (IAIPWE).

Early stopping means implementation of behavioral interventions for pain management in breast cancer patients, potentially helping more individuals and avoiding less efficacious regimes for those who otherwise would have enrolled in the trial.

8 | DISCUSSION

We proposed interim analysis methods for SMARTs that gain efficiency by using partial information from participants who have not yet completed all stages of the study. The approach yields a smaller expected sample size than competing methods while preserving type I error and power. Simulations demonstrate a potential for substantial resource savings.

We have demonstrated the methodology in the case of two-stage SMARTs with an interim analysis focused on evaluation of efficacy. However, the methods extend readily to studies with $K \geq 2$ decision points, multiple interim looks, and general hypotheses including futility. We have consider Pocock and OBF boundaries, though the approach can be adapted to any monitoring method, such as information-based monitoring (Tsiatis, 2006a) and the use of α spending functions (DeMets & Lan, 1994).

We have made the simplifying assumptions throughout that: (i) the time between stages is fixed, which is the case for many SMARTs; and (ii) the final outcome is observed on all individuals by the end of the trial (so excluding the possibility of drop out). The extension to random times between stages is nontrivial. Simulations included in Web Appendix G suggest that the IAIPWE (incorrectly assuming fixed transition times) performs well when time per stage varies with subject outcomes. Due to variability in enrollment, an analysis at a predetermined time may have a realized power slightly different from the nominal power based on the number of individuals enrolled and their realized trajectories at the time of analysis. In such cases, planning the interim analysis based on available sample size rather than a predetermined time may be preferred. Extensions for additional levels of coarsening, such as those due to drop out, attrition, or time-to-event outcomes requires additional augmentation terms or changes to the functions $\lambda_r^\ell(t), K_r^\ell(t)$, and $L_{k(r)}^\ell(\overline{\mathbf{x}}_{k(r)})$. For a comprehensive review of the considerations involved, see Chap. 8 of Tsiatis et al. (2020). A modified multiple imputation strategy may also be used for missing data following that of Shortreed et al. (2014).

As demonstrated in our simulation experiments, the sandwich covariance estimator can overestimate the variance of the values and lead to conservative stopping boundaries when the number of parameters is close to the sample size. Interim analyses typically have larger sample sizes, so this issue is unlikely to occur in practice. The information proportion can be checked at each interim analysis to verify the planned proportions against the realized values. The IAIPWE stopping boundary and sample size calculations also require the challenge of positing models. Although we have studied the performance of the IAIPWE under these conditions to evaluate fully its properties, we anticipate the trialists will prefer to power a SMART based on the IPWE to avoid making the additional model assumptions. We advocate this approach in practice as it can assuage concerns about misspecified models while still benefiting from the efficiency gains of the IAIPWE. If a trial does reach the final analysis, using the AIPWE offers efficiency gains by effectively performing covariate adjustment. Here, the covariates to be used in the Q-functions should be specified before the trial begins.

The framework presented here forms the basis for additional methodology for interim monitoring for SMARTs with random times between stages and specialized endpoints. The IAIPWE has potential use in adaptive trials in which randomization probabilities, or even the set of treatments, varies with accumulating information (Jennison and Turnbull, 2000, Chap. 17; Wang & Yee, 2019). We will report on these developments in future work.

ACKNOWLEDGMENTS

The authors thank Dr. Anastasios Tsiatis for helpful remarks and insights into the demonstration of the independent increments property and the Associate Editor and three referees for thoughtful and important comments that led to improvements in the presentation.

DATA AVAILABILITY STATEMENT

Data sharing is not applicable to this article, as no actual data sets are generated or analyzed. The methods developed are proposed to enable analyses of data from future clinical trials. The simulated data set for the case study in Section 7 is available with the paper at the Biometrics website on Wiley Online Library.

ORCID

Cole Manschot https://orcid.org/0000-0002-5054-2663
Marie Davidian https://orcid.org/0000-0002-9202-6341

REFERENCES

Almirall, D., Nahum-Shani, I., Sherwood, N.E. & Murphy, S.A. (2014) Introduction to SMART designs for the development of adaptive interventions: with application to weight loss research. *Translational Behavioral Medicine*, 4(3), 260–274.

Artman, W.J., Nahum-Shani, I., Wu, T., Mckay, J.R. & Ertefaie, A. (2020) Power analysis in a smart design: sample size estimation for determining the best embedded dynamic treatment regime. *Biostatistics*, 21(3), 432–448.

Bigirumurame, T., Uwimphuwe, G. & Wason, J. (2022) Sequential multiple assignment randomized trial studies should report all key components: a systematic review. *Journal of Clinical Epidemiology*, 142, 152–160.

Chakraborty, B. & Moodie, E. (2013) *Statistical methods for dynamic treatment regimes*. New York, NY: Springer.

Chao, Y., Braun, T., Tamura, R. & Kidwell, K. (2020) A Bayesian group sequential small *n* sequential multiple-assignment randomized trial. *Applied Statistics*, 69, 663–680.

ClinicalTrials.gov. (2021) Optimizing delivery of a behavioral cancer pain intervention using a SMART. ClinicalTrials.gov NCT02791646.

DeMets, D.L. & Lan, K.K. (1994) Interim analysis: the alpha spending function approach. *Statistics in Medicine*, 13, 1341–1352.

Han, P. (2014) Multiply robust estimation in regression analysis with missing data. *Journal of the American Statistical Association*, 109(505), 1159–1173.

Jennison, C. & Turnbull, B. (2000) *Group sequential methods with applications to clinical trials*. Boca Raton, FL: Chapman & Hall/CRC Press.

Kelleher, S.A., Dorfman, C.S., Vilardaga, J.C.P., Majestic, C., Winger, J., Gandhi, V. et al. (2017) Optimizing delivery of a behavioral pain intervention in cancer patients using a sequential multiple assignment randomized trial (SMART). *Contemporary Clinical Trials*, 57, 51–57.

Kidwell, K.M. & Hyde, L.W. (2016) Adaptive interventions and SMART designs: application to child behavior research in a community setting. *American Journal of Evaluation*, 37(3), 344–363.

Kosorok, M.R. & Laber, E.B. (2019) Precision medicine. *Annual Review of Statistics and Its Application*, 6, 263–286.

Lavori, P. & Dawson, R. (2004) Dynamic treatment regimes: practical design considerations. *Clinical Trials*, 1, 9–20.

Luedtke, A.R., Sofrygin, O., van der Laan, M.J. & Carone, M. (2018) Sequential double robustness in right-censored longitudinal models. *arXiv*. [Preprint] Available from https://arxiv.org/pdf/1705.02459

Manschreck, T.C. & Boshes, R.A. (2007) The CATIE schizophrenia trial: results, impact, controversy. *Harvard Review of Psychiatry*, 15(5), 245–258.

Murphy, S.A. (2005) An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine*, 24, 1455–1481.

O'Brien, P.C. & Fleming, T.R. (1979) A multiple testing procedure for clinical trials. *Biometrics*, 35(3), 549–556.

Pocock, S.J. (1977) Group sequential methods in the design and analysis of clinical trials. *Biometrika*, 64(2), 191–199.

Seewald, N.J., Kidwell, K.M., Nahum-Shani, I., Wu, T., McKay, J. & Almirall, D. (2020) Sample size considerations for comparing dynamic treatment regimens in a sequential multiple-assignment randomized trial with a continuous longitudinal outcome. *Statistical Methods in Medical Research*, 29(7), 1891–1912.

Shortreed, S.M., Laber, E., Scott Stroup, T., Pineau, J. & Murphy, S.A. (2014) A multiple imputation strategy for sequential multiple assignment randomized trials. *Statistics in Medicine*, 33(24), 4202–4214.

Sinyor, M., Schaffer, A. & Levitt, A. (2010) The sequenced treatment alternatives to relieve depression (STAR*D) trial: a review. *Canadian Journal of Psychiatry*, 55(3), 126–135.

Thall, P.F. (2015) SMART design, conduct, and analysis in oncology. In: Kosorok, M.R. & Moodie, E.E.M. (Eds.) *Adaptive treatment strategies in practice: Planning trials and analyzing data for personalized medicine*. Philadelphia, PA: ASA-SIAM, pp. 41–54.

Tsiatis, A. (2006a) Information-based monitoring of clinical trials. *Statistics in Medicine*, 25, 3236–3244.

Tsiatis, A. (2006b) *Semiparametric theory and missing data*. New York: Springer.

Tsiatis, A.A., Davidian, M., Holloway, S.T. & Laber, E.B. (2020) *Dynamic treatment regimes: Statistical methods for precision medicine*. Boca Raton, FL: Chapman & Hall/CRC Press.

van der Laan, M.J. & Petersen, M.L. (2007) Causal effect models for realistic individualized treatment and intention to treat rules. *International Journal of Biostatistics*, 3(1), 1–52.

Vermeulen, K. & Vansteelandt, S. (2015) Bias-reduced doubly robust estimation. *Journal of the American Statistical Association*, 110(511), 1024–1036.

Wang, H. & Yee, D. (2019) I-SPY 2: a neoadjuvant adaptive clinical trial designed to improve outcomes in high-risk breast cancer. *Current Breast Cancer Reports*, 11(4), 303–310.

Wang, L., Rotnitzky, A., Lin, X., Millikan, R.E. & Thall, P.F. (2012) Evaluation of viable dynamic treatment regimes in a sequentially randomized trial of advanced prostate cancer. *Journal of the American Statistical Association*, 107(498), 493–508.

Wason, J. (2019) Design of multi-arm, multi-stage trials in oncology. In: Halabi, S. & Michiels, S. (Eds.) *Textbook of clinical trials in oncology: a statistical perspective*. New York: Chapman and Hall/CRC Press, pp. 155–182.

Wu, L., Wang, J. & Wahed, A.S. (2021) Interim monitoring in sequential multiple assignment randomized trials. *Biometrics*, 46, 1–11.

Zhang, B., Tsiatis, A., Laber, E. & Davidian, M. (2013) Robust estimation of optimal dynamic treatment regimes for sequential treatment decisions. *Biometrika*, 100, 681–694.

SUPPORTING INFORMATION

Web Appendices referenced in Section 4, 5, 6, 7, and 8 and source code are available with this paper at the Biometrics website on Wiley Online Library.

Data S1

How to cite this article: Manschot, C., Laber, E. & Davidian, M. (2023) Interim monitoring of sequential multiple assignment randomized trials using partial information. *Biometrics*, 1–14. https://doi.org/10.1111/biom.13854