Title

Efficient and robust approaches for analysis of sequential multiple assignment randomized trials: Illustration using the ADAPT-R trial.

Journal

Biometrics, 79(3)

Authors

Montoya, Lina
Kosorok, Michael
Geng, Elvin
et al.

Publication Date

2023-09-01

DOI

10.1111/biom.13808

Peer reviewed

HHS Public Access

Author manuscript

Biometrics. Author manuscript; available in PMC 2023 September 28.

Published in final edited form as:

Biometrics. 2023 September ; 79(3): 2577–2591. doi:10.1111/biom.13808.

Efficient and robust approaches for analysis of sequential multiple assignment randomized trials: Illustration using the ADAPT-R trial

Lina M. Montoya[1], **Michael R. Kosorok**[1,2], **Elvin H. Geng**[3], **Joshua Schwab**[4], **Thomas A. Odeny**[5,6], **Maya L. Petersen**[4]

[1]Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA

[2]Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina, USA

[3]Division of Infectious Diseases, Washington University in St. Louis, St. Louis, Missouri, USA

[4]Division of Biostatistics, School of Public Health, University of California, Berkeley, California, USA

[5]Center for Microbiology Research, Kenya Medical Research Institute, Nairobi, Kenya

[6]Division of Oncology, Department of Medicine, Washington University in St. Louis, St. Louis, Missouri, USA

Abstract

Personalized intervention strategies, in particular those that modify treatment based on a participant's own response, are a core component of precision medicine approaches. Sequential multiple assignment randomized trials (SMARTs) are growing in popularity and are specifically designed to facilitate the evaluation of sequential adaptive strategies, in particular those embedded within the SMART. Advances in efficient estimation approaches that are able to incorporate machine learning while retaining valid inference can allow for more precise estimates of the effectiveness of these embedded regimes. However, to the best of our knowledge, such approaches have not yet been applied as the primary analysis in SMART trials. In this paper, we present a robust and efficient approach using targeted maximum likelihood estimation (TMLE) for estimating and contrasting expected outcomes under the dynamic regimes embedded in a SMART, together with generating simultaneous confidence intervals for the resulting estimates. We contrast this method with two alternatives (G-computation and inverse probability weighting estimators). The precision gains and robust inference achievable through the use of TMLE to evaluate the effects of embedded regimes are illustrated using both outcome-blind simulations and a real-data

Correspondence Lina M. Montoya, Department of Biostatistics, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA. lmontoya@unc.edu.

SUPPORTING INFORMATION

Web Appendix A and B, referenced in Section 4, Web Appendix C and D, referenced in Section 5, and Web Appendix E, referenced in Section 6 are available with this paper at the Biometrics website on Wiley Online Library. The code, simulations, and results for this paper can be found at https://github.com/lmmontoya/SMART-sims and in the Supporting Information.

analysis from the Adaptive Strategies for Preventing and Treating Lapses of Retention in Human Immunodeficiency Virus (HIV) Care (ADAPT-R) trial (NCT02338739), a SMART with a primary aim of identifying strategies to improve retention in HIV care among people living with HIV in sub-Saharan Africa.

Keywords

dynamic treatment regimes; precision medicine; sequential multiple assignment randomized trial; targeted maximum likelihood estimation

1 | INTRODUCTION

One question central to precision medicine and public health asks: "who should get which intervention, and in what sequence?" For example, a wide class of sequenced strategies start with an initial intervention, and then switch to a new, often higher intensity intervention based on participant response. These strategies are personalized because both the decision to switch interventions and the timing of the switch depend on an individual's own response. Data generated from a sequential multiple assignment randomized trial (SMART) provide a straightforward way of evaluating the causal effects of such sequenced adaptive strategies (or dynamic treatment regimes). Often, participants are given treatment (either randomly or deterministically) at pre-specified decision points based on their measured information (e.g., past treatments and/or intermediate covariates) up to that point. Assigning treatment sequentially based on a participant's measured past—including commonly, a patient's own response to earlier treatment—defines a SMART's embedded dynamic treatment regimes (or simply, embedded regimes). These embedded regimes correspond to adaptive personalized strategies for assigning treatment, thus contributing to the goals of precision health. Critically, by design, SMARTs allow the effects of these embedded regimes (and others, such as optimal dynamic treatment regimes based on covariates beyond those that define the trial design; Kosorok & Laber 2019) to be identified and estimated without risk of bias.

SMART designs are increasingly growing in popularity. For example, a recent review by Bigirumurame et al. (2022) cites 24 SMART protocol papers published since 2014. While primary analyses for SMARTs sometimes aim to examine the single timepoint static effects of the treatment options in the SMART's nested trials, they increasingly (in either primary or secondary aims) aim to evaluate the effects of embedded regimes (e.g., Kasari et al. 2014; Karp et al. 2019) or additionally tailored individual interventions (e.g., Sherwood et al. 2016). When evaluating the SMART's embedded regimes, common approaches for estimating the expected counterfactual outcome (or "value") of a given embedded regime use inverse probability weighting (IPW) estimators, including weighting and replicating approaches (introduced in Robins 2002; van der Laan & Petersen 2007; Bembom and van der Laan 2007; see also Nahum-Shani et al. 2012) and G-computation approaches (introduced in Robins 1986, 1987; Lavori & Dawson 2000, 2004). IPW estimators, and some G-computation estimators (depending on how the sequential regressions are estimated) will generally provide unbiased estimates of the value of the embedded regime; however,

they are inefficient in that they do not make full use of baseline and time-updated covariates to improve estimator precision. Advances in semiparametric efficient substitution estimators, such as longitudinal targeted maximum likelihood estimation (TMLE), allow for the integration of machine learning in the estimation process, enabling more precise estimates while retaining valid inference (see Petersen et al. 2015 for a review in the context of SMARTs). Recent work has documented the potential of flexible covariate adjustment using machine learning, and TMLE in particular, to improve precision in single timepoint individually randomized trials (e.g., Benkeser et al. 2021) and cluster randomized trials (e.g., Balzer et al. 2021). Simulations used to inform the design of SMARTs (see, e.g., Petersen et al. 2015; Benkeser et al. 2020) further support the potential benefits of longitudinal TMLE for the primary analysis of embedded regimes in SMART studies. However, to the best of our knowledge, neither longitudinal TMLE nor other semiparametric efficient estimators have been implemented or reported as the primary analysis method of a published SMART.

In this paper, we first review, using the "Causal Roadmap" (Petersen & van der Laan, 2014), how SMART designs can be used to identify the effects of embedded regimes, including the expected counterfactual outcome (or value) of each regime had all participants in the population followed it. We then describe an efficient and robust approach to estimating these counterfactual quantities without reliance on model assumptions, beyond the assumption of sequential randomization known by design. Specifically, we describe a longitudinal TMLE (Bang & Robins, 2005; van der Laan & Gruber, 2012) for estimating the values of these embedded regimes. TMLE is a double robust, semi-parametric, efficient, plug-in estimator that incorporates machine learning to improve efficiency without sacrificing reliable inference. We review the assumptions needed for valid statistical inference using this estimator, and we show how to construct individual and simultaneous confidence intervals to evaluate multiple embedded regimes within a SMART. Specifically, we illustrate the use of longitudinal TMLE as the primary pre-specified analysis in the recently completed adaptive strategies for preventing and treating lapses of retention in Human Immunodeficiency Virus (HIV) Care (ADAPT-R) trial (NCT02338739). We provide simulations to demonstrate the robustness of the approach, including an illustration of how outcome-blind simulations based on real trial data can be used to inform key decisions that must be pre-specified in a trial's analysis plan, such as specification of the machine learning methods employed for nuisance parameter estimation. We further provide a comparison to the commonly used IPW estimator. Using simulations and analyses of the trial data, we illustrate how the pre-specified use of TMLE integrating machine learning, in the analysis of the ADAPT-R trial, resulted in substantial improvements in efficiency, and thereby trial power, and discuss the interpretation of trial results.

The paper is organized as follows. In Section 2, we provide background on the ADAPT-R trial. In Section 3, we describe the causal model, define the causal parameters corresponding to the value of each embedded regime, and identify statistical parameters. In Section 4, we discuss estimation and inference of the identified statistical parameters. In Section 5, we present two simulation studies, with the dual objectives of illustrating the performance of these estimators and demonstrating how outcome blind simulations can be used to fully pre-specify a machine learning-based primary trial analysis using TMLE. In Section 6, we apply these methods to the ADAPT-R study. We close with a discussion.

2 | THE ADAPT-R TRIAL

The ADAPT-R trial was a SMART carried out to evaluate individualized sequenced behavioral interventions to optimize successful HIV care outcomes in Kenya. Up to 30% of persons receiving HIV care in this population experience at least one lapse in HIV care; these lapses in retention can result in loss of viral suppression. Importantly, patients that experience a retention lapse have a diversity of characteristics and needs (Geng et al., 2015). As a result, there is no "one-size-fits-all" incentive or strategy to help patients stay in care and achieve virologic suppression, demonstrating the need for effective personalized treatment regimes to increase successful HIV care outcomes.

In ADAPT-R, 1,809 persons living with HIV and initiating antiretroviral treatment (ART) in the Nyanza region of Kenya were randomized to one of three initial interventions to prevent a lapse in care (short message service [SMS] text messages, conditional cash transfers [CCTs] in the form of transportation vouchers for on-time visits, or standard of care [SOC] education and counseling). Patients who had a lapse in care within the first year of follow-up were re-randomized to a more intensive intervention to facilitate return to care (SMS text messages paired with CCTs, peer navigation, or SOC outreach); patients who did not have a lapse in care during the first year and who received SMS or CCTs in the first randomization were re-randomized to either continue or discontinue that intervention (study design shown in Figure 1).

Thus, in ADAPT-R there were 15 embedded regimes (see Table 1 for the complete list) that would have initially administered either SMS, CCTs, or SOC to all patients starting ART, and then either (a) SMS with CCTs, peer navigators, or SOC in the second stage should a lapse occur, or (b) for those on active first line treatment, a decision to continue or discontinue first stage treatment should no lapse occur. This paper describes how to estimate the counterfactual probability of having suppressed viral replication (plasma HIV RNA level < 500 copies/ml) two years after initial randomization, if a given embedded regime had been used for the full study population.

3 | CAUSAL ROADMAP

3.1 | Causal models

The following structural causal model (SCM, denoted as $\mathscr{M}^F$) will be used to describe the longitudinal process that gives rise to variables that are observed (endogenous) and not observed (exogenous) (Pearl, 2000). The random variables in $\mathscr{M}^F$ follow the joint distribution $P_{U,X}$,; the SCM describes the set of possible distributions for $P_{U,X}$,. For a time t, the endogenous variables are (1) categorical interventions $A(t) \in \mathscr{A}_t$ (which could include right-censoring); (2) covariates $X(t) \in \mathscr{X}_t$, which include baseline covariates and time-varying covariates between interventions at time $t-1$ and t (which could include indicators of time-dependent processes, such as death), and; (3) an outcome $Y \in \mathbb{R}$. Overbars are used to denote a variable's past history, for example, $\overline{A}(t) = (A(1), \ldots, A(t))$ and $\overline{X}(t) = (X(1), \ldots, X(t))$, and $A(0) = X(0) = \varnothing$. Let $\overline{Z}(t) \subseteq (\overline{A}(t-1), \overline{X}(t))$ denote the subset of endogenous tailoring variables used by design in the SMART to assign treatment at time

t. Then, for observation time $t = 1, \ldots, K$, the following structural equations can describe a SMART's longitudinal data-generating process:

$$\begin{aligned} X(t) &= f_{X(t)}\left(U_{X(t)}, \overline{X}(t-1), \overline{A}(t-1)\right) \\ A(t) &= f_{A(t)}\left(U_{A(t)}, \overline{Z}(t)\right) \\ Y &= f_Y\left(U_Y, \overline{X}(K), \overline{A}(K)\right), \end{aligned} \tag{1}$$

where exogenous variables are denoted as $U = (U_{X(t)}, U_{A(t)}, U_Y)$ and represent the unmeasured random input to the data-generating system. Importantly, $f_{A(t)}$ represent known parametric functions (specifically, the randomization scheme used in the SMART). Further, $U_{A(t)}$ is known by design in a SMART to be independent of all other exogenous factors. We note that in a SMART, t need not be the observation time; here, t is the time of treatment assignment. Thus, in a K-stage SMART (without intervening on non-randomized intervention nodes, such as censoring), K corresponds to the number of randomization stages.

3.1.1 | Data and models: application to ADAPT-R study—The ADAPT-R study provides an illustration of a SMART, where $t = 1$ is time of first randomization and $t = 2$ is time of second randomization (either date of first retention lapse or one year after initial randomization, whichever occurs first). In ADAPT-R, baseline covariates $X(1)$ included participant sex, age, WHO disease stage, CD4+ T cell count, an alcohol consumption measure, pregnancy status, and clinic site. The Stage 1 prevention intervention $A(1)$ consisted of either SMS text messages, CCTs, or SOC, each assigned with equal probability. Covariates assessed between randomization to Stage 1 and Stage 2 interventions, $X(2) = (L(2), S(2))$, included $L(2)$, an indicator of whether there was a lapse in care (≥ 14 days late to a clinic visit) within the first year after enrollment, and $S(2)$, which included death, transfer to another clinic, time from first randomization to second randomization, pregnancy status at second randomization, plasma HIV RNA level at second randomization, and whether a participant could be successfully contacted prior to randomization. The Stage 2 retention intervention $A(2)$ consisted of either (a) SOC outreach, SMS and CCT combined (SMS+CCT), or a peer navigator (Nav), each assigned with equal probability if there was a lapse in care $(L(2) = 1)$; or, (b) continuing or discontinuing the Stage 1 intervention, each assigned with equal probability, if there was no lapse in care $(L(2) = 0)$ and the initial intervention was either SMS or CCT; or, (c) continuing SOC if there was no lapse in care and the initial intervention was SOC. The outcome of interest Y was an indicator of remaining alive and with viral suppression at year 2.

The SCM for ADAPT-R can be written as follows, where $\overline{Z}(2) = (A(1), L(2))$:

$$\begin{aligned} X(1) &= f_{X(1)}(U_{X(1)}) \\ A(1) &= f_{A(1)}(U_{A(1)}, Z(1)) \\ X(2) &= f_{X(2)}(U_{X(2)}, X(1), A(1)) \\ A(2) &= f_{A(2)}\left(U_{A(2)}, \overline{Z}(2)\right) \\ Y &= f_Y\left(U_Y, \overline{X}(2), \overline{A}(2)\right), \end{aligned} \tag{2}$$

where $A(1)$ is drawn from a multinomial distribution with $n = 1$, $k = 3$, $p_{\text{SMS}} = p_{\text{CCT}} = p_{\text{SOC}} = 1/3$ and $A(2)$ is drawn from a Bernoulli distribution with $p_{\text{discont.}} = p_{\text{cont.}} = 0.5$ if $L(2) = 0$ and

$A(1) \in \{\text{SMS}, \text{CCT}\}$, deterministically equal to SOC if $L(2) = 0$ and $A(1) = \text{SOC}$, and drawn from a multinomial distribution with $n = 1$, $k = 3$ and $p_{\text{SMS+CCT}} = p_{\text{Nav.}} = p_{\text{SOC}} = 1/3$ if $L(2) = 1$. Additionally, here $(U_{A(1)}, U_{A(2)})$ are independent of each other and all other Us.

3.2 | Causal questions and parameters

The focus of this paper is on evaluating outcomes under the dynamic treatment regimes embedded in a SMART. In other words, our causal questions take the form: what are the expected outcomes at the end of follow up if all members of the target population had followed each of the dynamic regimes embedded in the SMART design?

Let $\phi_t(\bar{A}(t-1), \bar{X}(t)) \subseteq \mathscr{A}_t$ denote the set of allowable treatments for a participant presenting with $(\bar{A}(t-1), \bar{X}(t)) = (\bar{a}(t-1), \bar{x}(t))$ at time t. Then, let a decision rule d_t be a function that takes as input the information accrued on a participant up to time t and outputs a single treatment level from among the set of possible treatment levels to which a participant could be randomly assigned, for all covariate and treatment histories, that is, $d_t : (\mathscr{X}_1 \times \cdots \times \mathscr{X}_t, \mathscr{A}_1 \times \cdots \times \mathscr{A}_{t-1}) \to \phi_t(\bar{A}(t-1), \bar{X}(t))$. Denote $\mathscr{D}_t$ as the set of all such decision rules at time t. Let $d = (d_1, d_2, \ldots, d_K)$ be a dynamic treatment regime (i.e., a sequence of rules for assigning a treatment level at each randomization stage), and let $\bar{d}_t : (\mathscr{X}_1 \times \cdots \times \mathscr{X}_t, \mathscr{A}_1 \times \cdots \times \mathscr{A}_{t-1}) \to \phi_1(X(1)) \times \cdots \times \phi_t(\bar{A}(t-1), \bar{X}(t))$ denote a regime sequence until time t. Let $\mathscr{D}$ be the set of all such dynamic treatment regimes.

We focus here on the *embedded* dynamic treatment regimes in a SMART, which are particular sequences of rules with input $\bar{Z}(t)$, the tailoring variables used for assigning treatment at time t in the actual randomization scheme employed by the trial, and output in $\phi_t(\bar{A}(t-1), \bar{X}(t))$, for all t. The set of embedded regimes $\widetilde{\mathscr{D}}$ are subset of the entire set of dynamic regimes, and we denote $\tilde{\bar{d}}_t(\bar{Z}(t))$ as an element of $\widetilde{\mathscr{D}}$, where $\tilde{\bar{d}}_t(\bar{Z}(t)) = \left\{\tilde{d}_1(Z(1)), \tilde{d}_2(\bar{Z}(2)), \ldots, \tilde{d}_t(\bar{Z}(t))\right\}$ is a SMART's embedded regime until time t.

A counterfactual outcome under an embedded dynamic treatment regime $\tilde{d}$ is an individual's outcome if, possibly contrary to fact, the individual had been assigned treatment according to the embedded regime $\tilde{d}$. This counterfactual outcome, denoted as $Y_{\tilde{d}}$, can be derived under an intervention on the above SCM, in which at each randomization stage in the SMART, the randomized treatment assignment mechanism used in the SMART is replaced with a deterministic assignment of a single treatment level based on observed history; that is, for $t = 1, \ldots, K$:

$$\begin{aligned} X(t) &= f_{X(t)}\left(U_{X(t)}, \bar{X}(t-1), \bar{A}(t-1)\right) \\ A(t) &= \tilde{d}_t(\bar{Z}(t)) \\ Y_{\tilde{d}} &= f_Y\left(U_Y, \bar{X}(K)\right). \end{aligned} \tag{3}$$

The target causal parameters that answer our aforementioned causal queries are summary measures of the post-intervention distribution contained within the SCM. Here, the relevant causal parameters are the expected counterfactual outcomes had all participants received each of the SMART's embedded dynamic regimes; that is, for one $\tilde{d} \in \widetilde{D}$:

$$\Psi_d^F(P_{U,X}) = \mathbb{E}_{P_{U,X}}[Y_d], \tag{4}$$

and the vector of the counterfactual values of the D embedded regimes is denoted $\Psi^F(P_{U,X}) = \left\{\Psi_{d^{(1)}}^F(P_{U,X}), \ldots, \Psi_{d^{(D)}}^F(P_{U,X})\right\}$.

Of note, although in the current paper we focus on evaluating the particular regimes embedded within a SMART, we are not limited to asking the above causal questions when analyzing a SMART; by design, SMARTs easily allow for answering many causal questions corresponding to alternative aims of the study, such as:

1. Point treatment static regimes for the embedded nested trials. For example, in ADAPT-R: what is the counterfactual probability of either experiencing a lapse in retention by one year or viral non-suppression at one year (an interim outcome not affected by the second line intervention assignment) had everyone received each of the initial interventions (SMS, CCT, and SOC)?
2. Point treatment optimal dynamic treatment rule. For example, in ADAPT-R: what is the optimal way to assign initial SMS, CCT, or SOC to participants based on their measured baseline characteristics to minimize the probability of a retention lapse by year one or viral non-suppression at year one?
3. Longitudinal optimal dynamic treatment regime. For example, in ADAPT-R: what is the optimal way to assign Stage 1 and Stage 2 treatments, in sequence, based on the observed baseline and time-varying covariates to minimize viral suppression at year two?

We refer the reader to Kosorok and Laber (2019) for an overview of possible methods for answering these questions, particularly those that estimate optimal dynamic treatment rules.

Further, it could also be of interest to contrast pairs of embedded regimes; for example, for a pair of embedded regimes numbered $i, j \in \{1, \ldots, D\}, i \neq j$, one possible causal parameter that contrasts the efficacy between the two strategies is $\mathbb{E}_{P_{U,X}}[Y_{d^{(i)}} - Y_{d^{(j)}}]$. Such contrasts follow naturally from the approach described in the paper to estimate the regime-specific mean outcomes. These contrasts could be specified a priori, or omnibus tests could be employed, such as comparing the best embedded regime versus the worst (without knowing in advance which is which) or whether there are any significant differences in any of the regime values.

3.2.1 Causal parameters—application to ADAPT-R study—Within ADAPT-R, the set of allowable treatments at each timepoint given past participant information is as follows: $\phi_1(X(1)) = \{\text{SMS}, \text{CCT}, \text{SOC}\}$, $\phi_2(X(1), A(1), S(2), L(2) = 1) = \{\text{SMS} + \text{CCT}, \text{Nav.}, \text{SOC}\}$, $\phi_2(X(1), A(1) \in \{\text{SMS}, \text{CCT}\}, S(2), L(2) = 0) = \{\text{continue}, \text{discontinue}\}$, and $\phi_2(X(1), A(1) = \{\text{SOC}\}, S(2), L(2) = 0) = \{\text{continue}\}$.

Then, let $d = (d_1, d_2)$ be a dynamic treatment regime that uses participant information to assign the allowed treatments $\phi_1(X(1))$ and $\phi_2(A(1), \overline{X}(2))$ at Stages 1 and 2, respectively; that

is, d assigns $A(1)$ and $A(2)$ based on baseline covariates $X(1)$ and time-varying covariates and initial treatment $\{\overline{X}(2), A(1)\}$, respectively.

Specifically, we are interested in evaluating the particular sequence of rules that were used for assigning treatment within the SMART, $\tilde{d} \in \widetilde{\mathscr{D}}$, that is, the embedded regimes within ADAPT-R. The decision rules within the embedded regimes are characterized as follows: (1) at time 1, treat with either SMS, CCT, or SOC, regardless of baseline covariates, that is, $\tilde{d}_1 : Z(1) \rightarrow \phi_1(X(1))$, where $Z(1) = \emptyset$ and, (2) at time 2, treat with either SMS+CCT, Nav., SOC outreach, continue, or discontinue, depending on the initial treatment decision and whether there is a lapse in care in year 1, that is, $\tilde{d}_2 : Z(2) \rightarrow \phi_2(\overline{X}(2), A(1))$, where $Z(2) = \{A(1), L(2)\}$. For example, one embedded dynamic treatment regime $\tilde{d}$ assigns treatment via the following strategy: (1) $\tilde{d}_1$ = assign SMS to everyone; (2) $\tilde{d}_2$ = assign SMS+CCT if $L(2) = 1$ (lapse in care), continue SMS otherwise (succeed in care).

For one embedded regime, $\Psi^F_{\tilde{d}}$ answers the causal question: what is the probability of viral suppression at two years of follow-up had everyone been assigned the same Stage 1 intervention, then each person assigned a Stage 2 intervention based on the participant's Stage 1 intervention and whether or not that person had a lapse in care? Further, we are interested in the vector Ψ^F, which contains the counterfactual probabilities of 2 year viral suppression had everyone received each of the 15 strategies listed in Table 1.

Finally, one may be interested in contrasting two adaptive strategies for preventing lapses in HIV care. For example, to compare the efficacy of the first two regimes in Table 1, let $\tilde{d}^{(1)}$ be regime #1 from Table 1 (SOC, then SOC outreach if lapse and continue if no lapse), and $\tilde{d}^{(2)}$ be regime #2 (SMS, then SOC outreach if lapse and continue if no lapse). Then, the causal parameter that contrasts these two strategies is $\mathbb{E}_{P_{U,X}}[Y_{\tilde{d}^{(1)}} - Y_{\tilde{d}^{(2)}}]$.

3.3 | Statistical model, identification, and statistical target parameter

We assume that the observed data $O_i \equiv \left(\overline{X}(K)_i, \overline{A}(K)_i, Y_i\right) \sim P_0 \in \mathscr{M}, i = 1, \ldots, n$ were generated by sampling n independent and identically distributed copies from a data-generating system contained in $\mathscr{M}^F$ above. Here, P_0 is the observed data distribution, an element of $\mathscr{M}$, the statistical model.

Two conditions are necessary for identification; that is, for determining that the causal parameter (i.e., Equation (4), a function of the counterfactual distribution, $P_{U,X}$) is equivalent to a statistical parameter (a function of the observed data distribution P_0) for all distributions $P_{U,X}$ contained in $\mathscr{M}^F$. For $t = 1, \ldots, K$ and $\tilde{d} \in \tilde{D}$, we consider the (1) sequential randomization assumption (SRA): $Y_{\tilde{d}_t} \perp A(t) \mid \overline{X}(t), \overline{A}(t-1) = \tilde{d}_{t-1}(\overline{Z}(t-1))$; and, (2) sequential positivity assumption: $g_0\left(A(t) = \tilde{d}_t(\overline{Z}(t)) \mid \overline{X}(t), \overline{A}(t-1) = \tilde{\tilde{d}}_{t-1}(\overline{Z}(t-1))\right) > 0 - a.e.$, where $g_{A(t),0}(A(t) \mid \overline{X}(t), \overline{A}(t-1)) = P_0(A(t) \mid \overline{X}(t), \overline{A}(t-1))$ is the true conditional probability of the treatment at time t given measured time-varying variables used in the study design. Informally, the SRA states that there are no unmeasured common causes between assignment of $A(t)$ and Y, given that the individual has followed the regime up to t and

information accrued up to t. The sequential positivity assumption states that among subjects who have followed the regime up to t, there must be a positive probability of continuing to follow that regime at t, regardless of a participant's past information.

A SMART, by design, ensures that both conditions are met. For example, in ADAPT-R, individuals are completely randomized to $A(1)$ and are randomized based on measured, accrued information (i.e., $\overline{Z}(2)$) to $A(2)$. The probability of receiving any of the embedded decision rule treatments at Stage 1 given baseline covariates $X(1)$ is 1/3; the probability of receiving any of the possible embedded decision rule treatments at Stage 2 is 1/3 among people who had a lapse, 1/2 among people who succeeded in care and were initially given SMS or CCT, and 1 among people who succeeded in care and were initially given SOC. The general statistical parameter corresponding to $\Psi^F_{\tilde{d}}(P_{U,X})$ for one embedded regime $\tilde{d} = \tilde{\overline{d}}_K(\overline{Z}(K))$ is the G-computation formula (Robins, 1986):

$$\begin{aligned}\Psi_{\tilde{d}}(P_0) = &\sum_{x(1),\ldots,x(K)} \mathbb{E}_0\Big[Y \mid \overline{X}(K) = \bar{x}(K), \overline{A}(K) = \tilde{\overline{d}}_K(\overline{Z}(K))\Big] \\ &\times \prod_{t=1}^{K} P_0(X(t) = x(t) \mid \overline{X}(t-1) = \bar{x}(t-1), \overline{A}(t-1) \\ &= \tilde{\overline{d}}_{t-1}(\overline{Z}(t-1))\Big),\end{aligned} \tag{5}$$

where the summation generalizes to an integral for continuous $X(t)$. Equation (5) can also be re-written as a series of iterated conditional expectations (ICEs; or sequential regressions) (Bang & Robins, 2005):

$$\begin{aligned}\Psi_{\tilde{d}}(P_0) = \mathbb{E}_0\Big[\mathbb{E}_0\Big[\cdots\mathbb{E}_0\Big[\mathbb{E}_0\Big[Y \mid \overline{X}(K), \overline{A}(K) = \tilde{\overline{d}}_K(\overline{Z}(K))\Big] \mid \\ \overline{X}(K-1), \overline{A}(K-1) = \tilde{\overline{d}}_{K-1}(\overline{Z}(K-1))\Big]\cdots \mid X(1), A(1) = \tilde{d}_1(Z(1))\Big]\Big],\end{aligned} \tag{6}$$

or as the following IPW estimand: $\Psi_{\tilde{d}}(P_0) = \mathbb{E}_0\left[\dfrac{\mathbb{I}\left[\overline{A}(K) = \tilde{\overline{d}}_K(\overline{Z}(K))\right]}{\prod_{t=1}^{K} g_0(A(t) \mid \overline{X}(t), \overline{A}(t-1))} Y\right]$.

The observed data for an ADAPT-R participant are $O = \{\overline{X}(2), \overline{A}(2), Y\}$; the observed dataset consists of 1,809 i.i.d. observations of O generated by a process described by the aforementioned causal model. The statistical target parameter corresponding to the value of (i.e., the expectation of the counterfactual outcome under) an embedded regime within ADAPT-R is:

$$\begin{aligned}\Psi_{\tilde{d}}(P_0) = &\sum_{x(1),x(2)} \mathbb{E}_0\Big[Y \mid \overline{X}(2) = \bar{x}(2), \overline{A}(2) = \tilde{\overline{d}}_2(\overline{Z}(2))\Big] \\ &\times P_0\Big(X(2) = x(2) \mid x(1), A(1) = \tilde{d}_1(Z(1))\Big) P_0(X(1) = x(1)).\end{aligned} \tag{7}$$

The vector of all embedded regime values is identified as $\Psi(P_0) = (\Psi_{\tilde{d}^{(1)}}(P_0), \ldots, \Psi_{\tilde{d}^{(D)}}(P_0))$; in ADAPT-R, $D = 15$. Finally, if one were interested in comparing the value of two embedded regimes $\tilde{d}^{(i)}$ and $\tilde{d}^{(j)}$, the statistical parameter corresponding to this contrast would be $\Psi_{\tilde{d}^{(i)}}(P_0) - \Psi_{\tilde{d}^{(j)}}(P_0)$.

4 | ESTIMATION AND INFERENCE FOR THE VALUES OF EMBEDDED REGIMES

We are interested in estimators for the statistical parameter identified in Section 3.3—that is, estimators for evaluating a SMART's embedded dynamic regimes. We focus on a longitudinal TMLE, and compare this with the IPW and G-computation estimator based on ICEs. All of these estimators can be implemented with the *ltmle* package (Lendle et al., 2017; Petersen et al., 2014). We briefly describe the longitudinal TMLE employed in the ADAPT-R analysis here, and we refer the reader to the Web Appendix A for a detailed description of the steps for implementing the three estimators.

The longitudinal TMLE employed here is a flexible and robust approach that estimates the value of a sequential regime by fitting initial estimates of the series of ICEs (Equation (6)) and updating these using either the known or estimated treatment mechanisms (Bang&Robins, 2005; van der Laan & Gruber, 2012). Critically, TMLE allows for the use of flexible machine learning methods, such as SuperLearner (van der Laan et al., 2007), to generate the initial estimates of the ICEs. Once updated using the treatment mechanisms, these estimates are then used to implement a plug-in estimator of the target parameter, as defined in Equation (6). In contrast, the G-computation estimators exclusively rely on initial (untargeted) estimates of the ICEs, while the IPW estimator relies on either the estimated or true (and known, in a SMART) treatment mechanisms to estimate the value of an embedded regime.

Inference for the TMLE estimates of the embedded regime values can be based on estimates of the efficient influence curve for the target statistical parameter (Bang &Robins, 2005), which can be used to construct Wald-type 95% confidence intervals that, under assumptions, provide nominal to conservative coverage for the value of the one embedded regime. Further, because one goal is to evaluate the multiple dynamic regimes embedded in a SMART at the same time, one can also use an estimate of the efficient influence curve to construct simultaneous confidence intervals (Cai & van der Laan, 2020). For example, in ADAPT-R, one goal might be to evaluate 15 embedded regimes simultaneously; simultaneous confidence intervals aim to ensure that all estimated confidence intervals contain the true values of the embedded regimes at the nominal level, thus providing one approach to account for multiplicity. The same approach can be easily extended to handle multiple comparisons of these regimes. We refer the reader to Web Appendix B for technical details on how to construct these confidence intervals.

If implemented carefully, longitudinal TMLE has the potential to substantially improve efficiency (both asymptotically and in finite samples), and thereby increase study power. In a SMART, the treatment mechanism is known; thus, if there is no censoring, one could use the true conditional treatment probabilities $g_0(A(t) \mid \overline{X}(t), \overline{A}(t-1)) \equiv g_0(A(t) \mid \overline{Z}(t))$ in either the IPW or TMLE estimators. Estimator precision can be improved, however, by estimating the treatment mechanism using a maximum likelihood estimate of the parameters of a correctly specified parametric model (such as a generalized linear model including either $Z(t)$ alone, or including additional covariates in $X(t)$) (van der Laan & Robins, 2003). Either of these estimator specifications (either usage of the true treatment

mechanism probabilities or estimates via correctly specified parametric models) will result in IPW and TMLE estimators that are consistent; however, the use of TMLE allows for further efficiency gains through additional estimation of the ICEs. Informally, the resulting TMLE will be efficient as long as: (1) either the ICE initial estimates are not overfit, or sample splitting is incorporated in the estimator such that the targeted update is fit on data independent of that used in the initial fit (Zheng & van der Laan, 2010); and, (2) the ICEs are estimated consistently. Importantly, when using TMLE (or other double robust semiparametric efficient estimators), the iterated conditional expectations can be estimated using machine learning, increasing the chance that they are estimated consistently and potentially further improving finite sample variance.

We contrast these efficiency properties of TMLE with those of the G-computation and IPW estimators. In particular, there is no valid theory for inference on the G-computation estimator if the ICEs are estimated using either flexible machine learning algorithms or with mis-specified parametric models. IPW estimators do allow for conservative or nominal inference with either consistently estimated or true values of the treatment mechanisms, but they are not efficient.

5 | SIMULATIONS

Using simulations, we evaluated the performance of various estimators for the values of a SMART's embedded regimes. We did this for two data generating processes (DGPs) corresponding to SMART designs (i.e., in which the true treatment mechanism $g_0(A(t) \mid \overline{Z}(t))$ is assumed known): (1) a simple, hypothetical DGP in which re-randomization is based only on intermediate covariates (and not initial treatment; DGP 1); and (2) an outcome-blind simulation based on data from ADAPT-R in which the actual ADAPT-R empirical covariate distribution and known treatment mechanism was used, but for which the outcomes themselves were simulated (DGP 2). Importantly, the latter represents a powerful tool for fine-tuning choices for pre-specification of an estimator.

For both DGPs, we first implemented the IPW estimator using the true treatment mechanisms, the $g_0(A(t) \mid \overline{Z}(t))$ factors, which are known in a SMART (denoted as "Min. adj. IPW w/g_0"). Second, we implemented an IPW estimator, where $g_0(A(t) \mid \overline{Z}(t))$ was estimated using the empirical proportions of each covariate and treatment history strata (i.e., a saturated model, denoted as "Min. adj. IPW w/g_n"). We note that the G-computation estimator and TMLE will generate equivalent estimates to this IPW estimator if the ICE factors are also estimated with saturated regression models. Third, we implemented an IPW estimator, where $g_0(A(t) \mid \overline{Z}(t))$ was estimated using main-terms logistic regression models of treatment on all past treatments and covariates (including covariates not used in the randomization scheme, but predictive of the outcome). In subsequent results, this is denoted as "Full adj. IPW." Fourth, we implemented the G-computation estimator that adjusted for all covariates through ICE factors estimated with SuperLearner (see Web Appendix C for specifications including the library of algorithms used). Finally, we implemented a TMLE in which ICEs were estimated adjusting for all covariates using SuperLearner, and g was estimated using a correctly specified logistic regression model noting that this

parametric model specification was known to contain the true treatment mechanism. For TMLE, in DGP 1 all baseline and time-varying covariates were used in the estimation of the g factors, while in DGP 2 estimates of g used the minimal adjustment set (i.e., $\bar{Z}(2)$) to avoid overfitting.

We evaluated estimator performance in terms of absolute bias, variance, confidence interval width, and 95% confidence interval coverage (for both individual and simultaneous confidence intervals). Inference was based on the influence-curve procedures described in Section 4. If IPW was employed, the estimated IPW influence curve was used for inference; if TMLE was employed, the estimated efficient influence curve was used for inference. As noted in Section 4, we do not provide inference results for the G-computation estimator.

We refer the reader to Web Appendix C for details on how the simulations were implemented, including specific DGPs and algorithm configurations. We used the *ltmle* R package for estimation and inference (Lendle et al., 2017; Petersen & van der Laan, 2014). Each simulation consisted of 1,000 iterations of $n = 1{,}692$ observations (the sample size for the ADAPT-R's analysis dataset after excluding 117 persons for a missing outcome measure). We additionally implemented these simulations for DGP 1 with a reduced sample size of $n = 750$ and various different SuperLearner library configurations for comparison.

5.1 | Simulation results

The results described below are shown in Figures 2 and 3, tables in Web Appendix C, and additional figures in Web Appendix D.

The untargeted G-computation estimator, in which the ICEs were fit using machine learning, exhibited the highest bias among the estimators across the embedded regimes evaluated. Specifically, for DGP 1, the hypothetical SMART with minimal covariates, the mean absolute difference between the G-computation estimate and the truth ranged from 0.09% to 0.69% (1.70–91.01 times that of the bias of any IPW or TMLE estimator). For DGP 2, which incorporated covariates re-sampled from the empirical distribution of ADAPT-R data, the absolute mean difference between the G-computation estimate and the truth ranged from 0.20% to 1.41%; across all embedded regimes except for 2 and 6, bias of the G-computation estimator was 1.62–138.64 times that of the bias of IPW or TMLE. For the remaining embedded regimes 2 and 6, bias of the G-computation estimator was minimal and similar to the IPW estimator and TMLE, likely because the SuperLearner consistently chose saturated regression models to estimate the ICEs, thus generating equivalent estimates to TMLE and IPW.

As expected, the IPW that used the known treatment mechanism ("Min. adj. IPW w/g_0") was unbiased with close to nominal coverage (93.1%–95.8% range across both DGPs and both confidence interval types), but exhibited a higher variance than other estimators. For example, the relative variance of the IPW estimator that used the known, true probabilities of treatment versus the IPW estimator that used empirical proportions given the minimal adjustment set $\bar{Z}(2)$ to estimate the g factors (i.e., "Min. adj. IPW w/g_0" vs "Min. adj. IPW w/g_n") was 2.10–6.14 for DGP 1 and 2.87–3.37 for DGP 2. Finally, the IPW estimator in which the g factors were estimated adjusting for additional baseline and time-varying

covariates (i.e., “Full adj. IPW”) resulted in some variance reduction compared to the IPW in which g was estimated using only the minimal adjustment set $\bar{Z}(2)$ (i.e., the “Min. adj. w/g_n” variance was up to 1.14 that of the “Full adj. IPW” variance). Although estimation of the treatment mechanism reduced the variance of the point estimator, reductions in the 95% confidence interval widths (and by implication, power for contrasting regimes) were limited by the fact that influence curve-based inference for IPW estimators in which g was estimated yielded conservative inference (i.e., 99.3%–100.0% confidence interval coverage across all DGPs and confidence interval types).

The TMLE with the estimated treatment mechanisms and ICEs estimated using machine learning (SuperLearner) were unbiased with close to nominal confidence interval coverage (93.4%–96.0% across both DGPs and both confidence interval types). In addition, TMLE showed variance gains relative to IPW, particularly for DGP 2 (e.g., the relative variance of the fully adjusted IPW versus TMLE was 1.01–1.12 for DGP 1 and 1.36–1.58 for DGP 2), through its ICE estimation using machine learning (i.e., beyond estimation of the g factors, as in IPW estimation). TMLE resulted in substantially narrower mean 95% confidence intervals across both DGPs (1.57–2.62 times that of fully adjusted IPW confidence interval widths), due both to a slightly lower variance of the estimator, and the less conservative influence curve-based variance estimation compared to IPW.

In further simulations (see results in Web Appendix D), a reduction in sample size to $n = 750$ resulted in similar comparative performance across the estimators. Inclusion of a tree-based method in the SuperLearner library (namely, recursive partitioning and regression trees; Breiman et al. 1984) or highly adaptive lasso (HAL; Benkeser & van der Laan 2016) yielded similar patterns; though, notably, when including a tree-based method the bias increased significantly for the G-computation estimator.

6 | ADAPT-R STUDY RESULTS

Of the 1,809 ADAPT-R participants, 117 did not have a viral load outcome; these patients were excluded from the analytic dataset ($n = 1{,}692$; noting that if an outcome variable has substantial missingness, one could incorporate this into the causal model and thus adjust for it). Using this sample, we conducted two analyses for the paper, described below. A full report and interpretation of ADAPT-R’s main results, including the clinical and public health implications, will be published in a separate manuscript.

First, we estimated (using TMLE, as described for the “Full adj. TMLE” estimator in simulations for DGP 2) and obtained inference on (using influence curve-based single and simultaneous confidence intervals) the value of each of ADAPT-R’s 15 HIV care retention strategies (embedded regimes). The results of this analysis are shown in Figure 4 and Web Appendix E (Web Table 5, which includes the number of patients who contributed to each of the regimes). A point estimate reflects the estimated probability of viral suppression had the study population followed one of ADAPT-R’s embedded regimes. For example, for the second embedded regime, we estimate that an intervention to deliver SMS to the full target population at a time of ART initiation, followed by a transition from SMS to SOC outreach if a lapse in retention occurred, or the continuation of SMS if a lapse did not occur, would

have resulted in 78.34% (95% simultaneous CI: [70.84%, 85.83%]) of the population alive and with a suppressed viral load after two years.

Second, we evaluated five pre-specified contrasts (i.e., differences in values of pre-specified pairs of rules) between the regimes, shown in Figure 5 and Web Appendix E (Web Table 6). The first four regime pairs were chosen to compare the "fully active" regime arms versus SOC throughout; a fifth pre-specified contrast evaluated the effectiveness of a strategy of time-limited CCT (initial CCT, with SOC outreach if a lapse occurred and discontinuation if no lapse occurred) versus SOC throughout. We implemented all three IPW estimators and the fully adjusted TMLE, described in the above simulations. Due to the small number of pre-specified hypotheses tested, multiplicity correction was not employed in these tests. TMLE estimates suggest that an active first line therapy (such as SMS or CCT) followed by a tailored peer-navigator was effective in improving viral suppression compared to the current HIV care standard throughout. In contrast, a single time limited CCTs did not result in significantly different viral suppression compared to receiving SOC throughout.

We note the important variance reduction in the TMLE that used the full adjustment set versus all other candidate estimators. In particular, confidence interval widths for non-TMLE estimates were wider (i.e., 1.92 to 2.40 times wider) than the full adjustment set. Critically, had we used an IPW estimator using either the known treatment mechanism or an estimated treatment mechanism—a version of weighting and replicating—instead of the presented TMLE, the results would not have detected any statistically significant contrasts.

7 | DISCUSSION

The purpose of this paper was to illustrate implementation of longitudinal TMLE, integrating machine learning, for the evaluation of the embedded dynamic regimes in a SMART study. While previously described, to the best of our knowledge, this approach has not yet been applied in the primary published analysis of a SMART. We additionally illustrated how to obtain simultaneous confidence intervals on the values of multiple dynamic treatment regimes embedded in a SMART. In simulations and real data, we found substantial precision benefits from using this double robust, efficient estimator, especially in conjunction with adjustment for time-varying and baseline covariates using flexible machine learning approaches.

Specifically, in simulations, TMLE and IPW showed reduced bias compared to a G-computation estimator that utilized non-targeted machine learning-based iterated outcome regressions. Estimation of the treatment mechanism (compared to using the known, true probabilities of treatment), reduced the variance of the IPW estimator, with some further reduction in variance achieved through adjustment for covariates in addition to the minimal set needed for satisfying the SRA. However, the width of the IPW confidence intervals remained higher than that of TMLE using the corresponding adjustment set. Precision gains in the TMLE were driven both by reduced variance of the point estimator and by less conservative influence curved-based variances' estimates. When estimating both its treatment mechanism (via a correctly specified, baseline and time-varying covariate adjusted logistic regression) and iterated outcome regressions (via SuperLearner, adjusting

for all covariates) TMLE maintained close to nominal confidence interval coverage. This is analogous to the efficiency gains seen when adjusting for baseline covariates predictive of the outcome in a standard, single time-point randomized trials (Moore & van der Laan, 2009).

In addition, we showed how to evaluate the embedded regimes of ADAPT-R, a SMART carried out in Kenya to prevent lapses in HIV care. ADAPT-R's embedded regimes consisted of a menu of strategies that adapted to patients' responses to HIV care. The purpose of evaluating these regimes was to see the effect of each of these tailored strategies on viral suppression among this population. Using TMLE (with and without simultaneous confidence intervals), we obtained estimates of the probability of viral suppression for each of ADAPT's 15 embedded regimes. Further, with all IPW and TMLE estimators presented in this paper, we contrasted the efficacy between pre-specified strategies. Notably, results of these analyses illustrated the precision benefits associated with the longitudinal TMLE. In particular, had we used any of a set of common IPW estimators—estimators that do not fully leverage covariate data, machine learning, and semi-parametric efficiency theory—we would not have been able to learn that active sequence strategies tailored to having a lapse in care improve viral suppression, compared to the current HIV care standard throughout. This emphasizes the importance of the described efficiency gains within our HIV care research, and also given the recent increase in "small n SMARTs" (e.g., Chao et al. 2020).

The results of the analysis of ADAPT-R's embedded regimes present a menu of individualized strategies to help patients remain in HIV care. In particular, these results demonstrate the necessity for following HIV patients longitudinally in their treatment program, in order to escalate their treatment when needed, and shed light on effective escalation strategies. Critically, this kind of insight is not one we could have gleaned without a SMART. The current work aims to contribute to help uncover the potential of SMART designs so that—with more precision and certainty—we are more equipped to learn which treatments work better for whom, and when.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGMENTS

Research reported in this publication was supported by NIAID awards R01AI074345, K24AI134413, and F31AI140962. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

Funding information

National Institute of Allergy and Infectious Diseases, Grant/Award Numbers: F31AI140962, K24AI134413,

DATA AVAILABILITY STATEMENT

The data that support the findings in this paper are available from the corresponding author upon reasonable request.

REFERENCES

Balzer LB, van der Laan M, Ayieko J, Kamya M, Chamie G, Schwab J, et al. (2021) Two-stage TMLE to reduce bias and improve efficiency in cluster randomized trials. Biostatistics. kxab043, 1–16.

Bang H & Robins JM (2005) Doubly robust estimation in missing data and causal inference models. Biometrics, 61(4), 962–973. [PubMed: 16401269]

Bembom O & van der Laan MJ (2007) Statistical methods for analyzing sequentially randomized trials. Journal of the National Cancer Institute, 99(21), 1577–1582. [PubMed: 17971533]

Benkeser D, Díaz I, Luedtke A, Segal J, Scharfstein D & Rosenblum M (2021) Improving precision and power in randomized trials for covid-19 treatments using covariate adjustment, for binary, ordinal, and time-to-event outcomes. Biometrics, 77(4), 1467–1481. [PubMed: 32978962]

Benkeser D, Horvath K, Reback CJ, Rusow J & Hudgens M (2020) Design and analysis considerations for a sequentially randomized HIV prevention trial. Statistics in Biosciences, 12(3), 446–467. [PubMed: 33767798]

Benkeser D & van der Laan M (2016) The highly adaptive lasso estimator. In: 2016 IEEE international conference on data science and advanced analytics (DSAA), pp. 689–696. IEEE.

Bigirumurame T, Uwimpuhwe 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. [PubMed: 34763037]

Breiman L, Friedman JH, Olshen RA & Stone CJ (1984) Classification and Regression Trees. Routledge.

Cai W & van der Laan MJ (2020) One-step targeted maximum likelihood estimation for time-to-event outcomes. Biometrics, 76(3), 722–733. [PubMed: 31729005]

Chao Y-C, Trachtman H, Gipson DS, Spino C, Braun TM & Kidwell KM (2020) Dynamic treatment regimens in small *n*, sequential, multiple assignment, randomized trials: an application in focal segmental glomerulosclerosis. Contemporary Clinical Trials, 92, 105989.

Geng EH, Odeny TA, Lyamuya RE, Nakiwogga-Muwanga A, Diero L, Bwana M, et al. (2015) Estimation of mortality among HIV-infected people on antiretroviral treatment in East Africa: a sampling based approach in an observational, multisite, cohort study. Lancet HIV, 2(3), e107–e116. [PubMed: 26424542]

Karp JF, Zhang J, Wahed AS, Anderson S, Dew MA, Fitzgerald GK, et al. (2019) Improving patient reported outcomes and preventing depression and anxiety in older adults with knee osteoarthritis: results of a sequenced multiple assignment randomized trial (SMART) study. The American Journal of Geriatric Psychiatry, 27(10), 1035–1045. [PubMed: 31047790]

Kasari C, Kaiser A, Goods K, Nietfeld J, Mathy P, Landa R, et al. (2014) Communication interventions for minimally verbal children with autism: a sequential multiple assignment randomized trial. Journal of the American Academy of Child & Adolescent Psychiatry, 53(6), 635–646. [PubMed: 24839882]

Kosorok MR & Laber EB (2019) Precision medicine. Annual Review of Statistics and its Application, 6, 263–286.

Lavori PW & Dawson R (2000) A design for testing clinical strategies: biased adaptive within-subject randomization. Journal of the Royal Statistical Society: Series A (Statistics in Society), 163(1), 29–38.

Lavori PW & Dawson R (2004) Dynamic treatment regimes: practical design considerations. Clinical trials, 1(1), 9–20. [PubMed: 16281458]

Lendle SD, Schwab J, Petersen ML & van der Laan MJ (2017) ltmle: an R package implementing targeted minimum loss-based estimation for longitudinal data. Journal of Statistical Software, 81(1), 1–21.

Moore KL & van der Laan MJ (2009) Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. Statistics in Medicine, 28(1), 39–64. [PubMed: 18985634]

Nahum-Shani I, Qian M, Almirall D, Pelham WE, Gnagy B, Fabiano GA, et al. (2012) Experimental design and primary data analysis methods for comparing adaptive interventions. Psychological Methods, 17(4), 457. [PubMed: 23025433]

Pearl J (2000) Causality: models, reasoning, and inference. Cambridge University Press.

Petersen M, Schwab J, Gruber S, Blaser N, Schomaker M & van der Laan M (2014) Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. Journal of Causal Inference, 2(2), 147–185. [PubMed: 25909047]

Petersen ML, Schwab J, Geng EH & van der Laan MJ (2015) Chapter 10: Evaluation of longitudinal dynamic regimes with and without marginal structural working models. In: Kosorok Michael R. & Moodie Erica E. M. (Eds.) Adaptive treatment strategies in practice: planning trials and analyzing data for personalized medicine (pp. 157–186). Society for Industrial and Applied Mathematics.

Petersen ML & van der Laan MJ (2014) Causal models and learning from data: integrating causal modeling and statistical estimation. Epidemiology, 25, 418–426. [PubMed: 24713881]

Robins JM (1986) A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. Mathematical Modelling, 7, 9–12.

Robins JM (1987) Addendum to "a new approach to causal inference in mortality studies with a sustained exposure period-application to control of the healthy worker survivor effect". Computers & Mathematics with Applications, 14(9–12), 923–945.

Robins JM (2002) Analytic methods for estimating hiv-treatment and cofactor effects. Methodological issues in AIDS behavioral research. Springer, pp. 213–288.

Sherwood NE, Butryn ML, Forman EM, Almirall D, Seburg EM, Crain AL, et al. (2016) The bestfit trial: a smart approach to developing individualized weight loss treatments. Contemporary Clinical Trials, 47, 209–216. [PubMed: 26825020]

van der Laan MJ & Gruber S (2012) Targeted minimum loss based estimation of causal effects of multiple time point interventions. The International Journal of Biostatistics, 8(1), 1557–4679.

van der Laan MJ & Petersen ML (2007) Causal effect models for realistic individualized treatment and intention to treat rules. The International Journal of Biostatistics, 3(1), 3. [PubMed: 19122793]

van der Laan MJ, Polley EC & Hubbard AE (2007) Super learner. Statistical Applications in Genetics and Molecular Biology, 6(1), 1–21.

van der Laan MJ. & Robins JM. (2003) Unified methods for censored longitudinal data and causality. Springer.

Zheng W & van der Laan MJ (2010) Asymptotic theory for cross-validated targeted maximum likelihood estimation. U.C. Berkeley Division of Biostatistics, Working Paper Series.

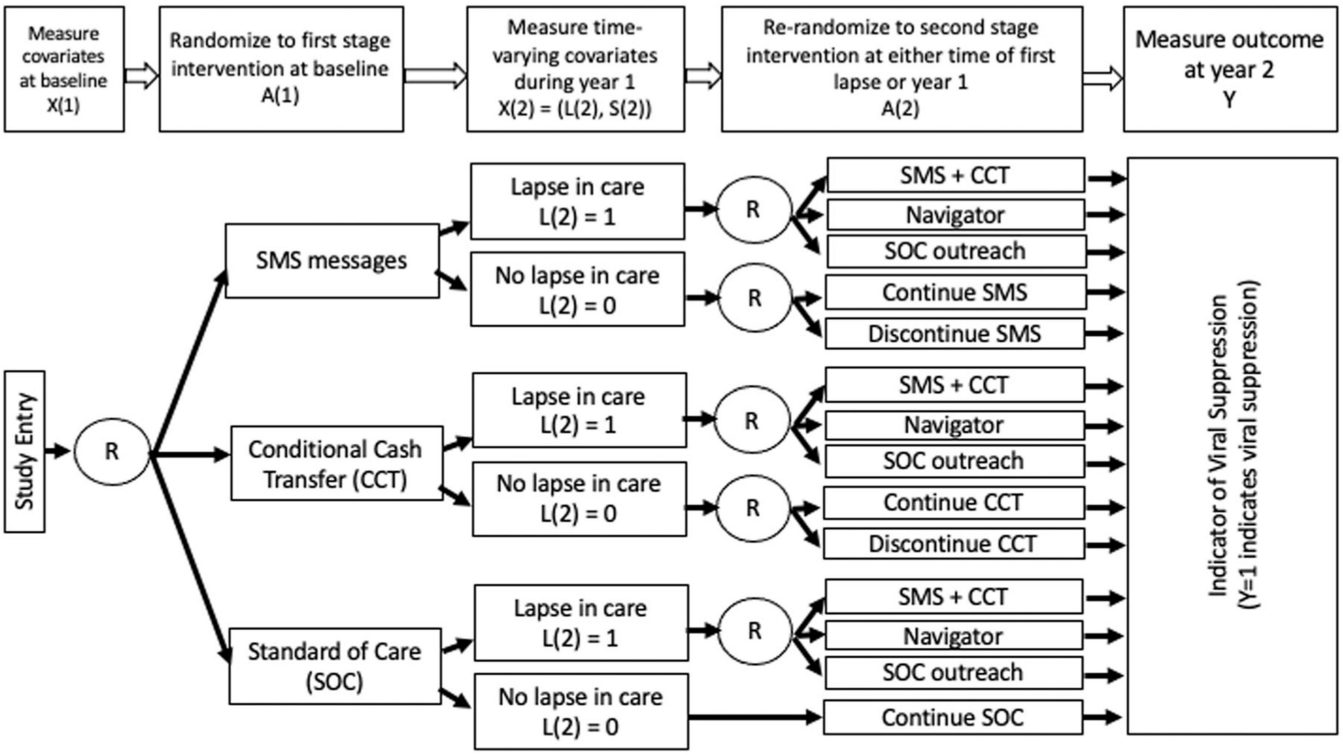

FIGURE 1.

The Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) study design, a sequential multiple assignment randomized trial (SMART). The circles with an "R" denote points of randomization

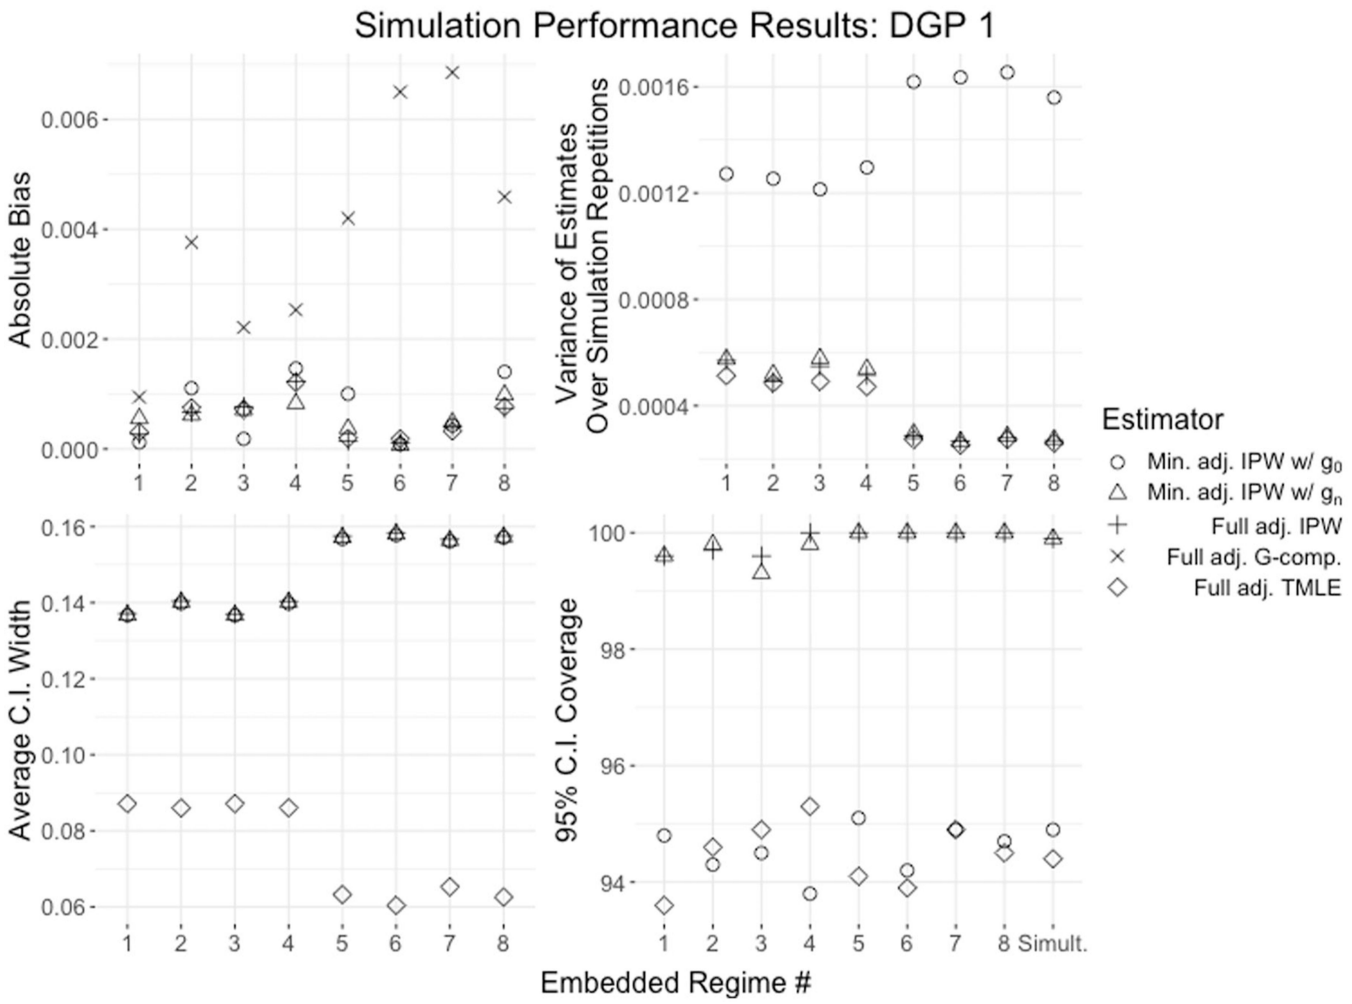

FIGURE 2.

DGP 1. Performance (top left panel is absolute bias, top right panel is Monte Carlo variance over simulation repetitions, bottom left panel is mean confidence interval [CI] width across simulation repetitions, and bottom right panel is 95% CI coverage) of candidate estimators of the value of each of the 8 embedded regimes within the simple Sequential Multiple Assignment Randomized Trial (SMART) generated from DGP 1. The five estimators evaluated are: (1) an inverse probability weighted (IPW) estimator with weights based on the true, known probability of receiving treatment given the initial treatment and lapse response ("Min. adj IPW w/g_0"); (2) an IPW estimator with estimated weights based on the empirical proportion of receiving treatment given the initial treatment and lapse response, which is equivalent to a TMLE or G-computation estimator where iterated ICE factors are estimated with saturated regression models ("Min. adj IPW w/g_n"); (3) an IPW estimator with estimated weights that adjust for all covariates ("Full adj. IPW"); (4) a G-computation estimator based on ICEs estimated with machine learning that adjust for all covariates ("Full adj. G-comp."); and (5) a targeted maximum likelihood estimator (TMLE) that adjusts for all covariates ("Full adj. TMLE"). Both individual and simultaneous CI coverage is shown under the regime numbers 1–8 and "Simult.," respectively

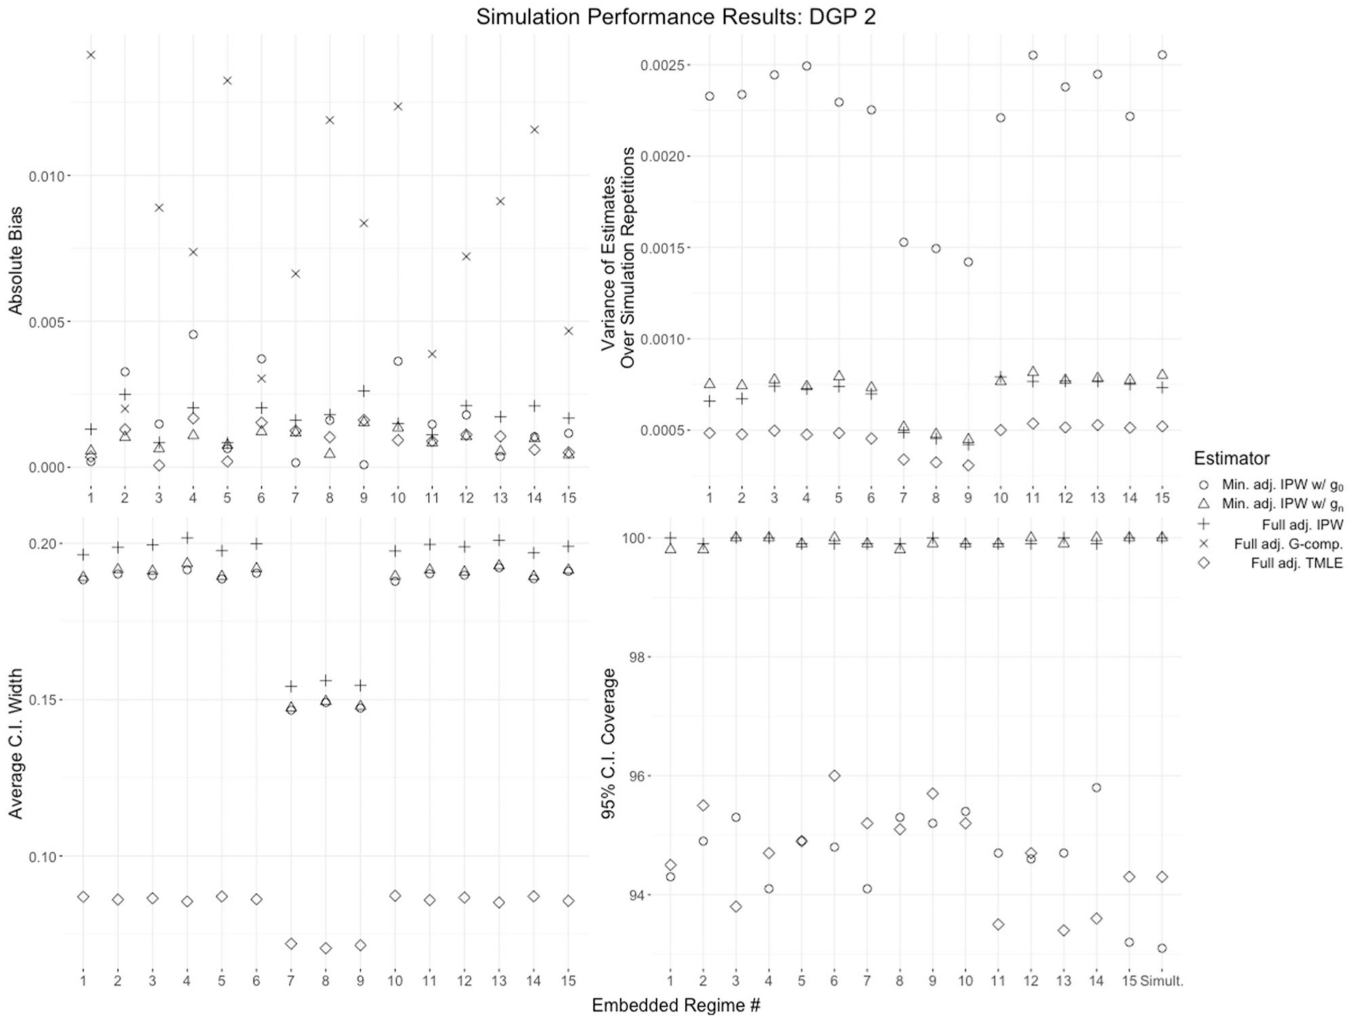

FIGURE 3.

DGP 2. Performance (top left panel is absolute bias, top right panel is Monte Carlo variance over simulation repetitions, bottom left panel is mean confidence interval [CI] width across simulation repetitions, and bottom right panel is 95% CI coverage) of candidate estimators of the value of each of the 15 embedded regimes within the outcome-blind simulation of the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial (DGP 2). The five estimators evaluated are: (1) an inverse probability weighted (IPW) estimator with weights based on the true, known probability of receiving treatment given the initial treatment and lapse response ("Min. adj IPW w/g_0"); (2) an IPW estimator with estimated weights based on the empirical proportion of receiving treatment given the initial treatment and lapse response, which is equivalent to a TMLE or G-computation estimator where iterated conditional expectation (ICE) factors are estimated with saturated regression models ("Min. adj IPW w/g_n"); (3) an IPW estimator with estimated weights that adjust for all covariates ("Full adj. IPW"); (4) a G-computation estimator based on ICEs estimated with machine learning that adjust for all covariates ("Full adj. G-comp."); and (5) a targeted maximum likelihood estimator (TMLE) that adjusts for all covariates ("Full adj.

TMLE"). Both individual and simultaneous CI coverage is shown under the regime numbers 1–15 and "Simult.," respectively

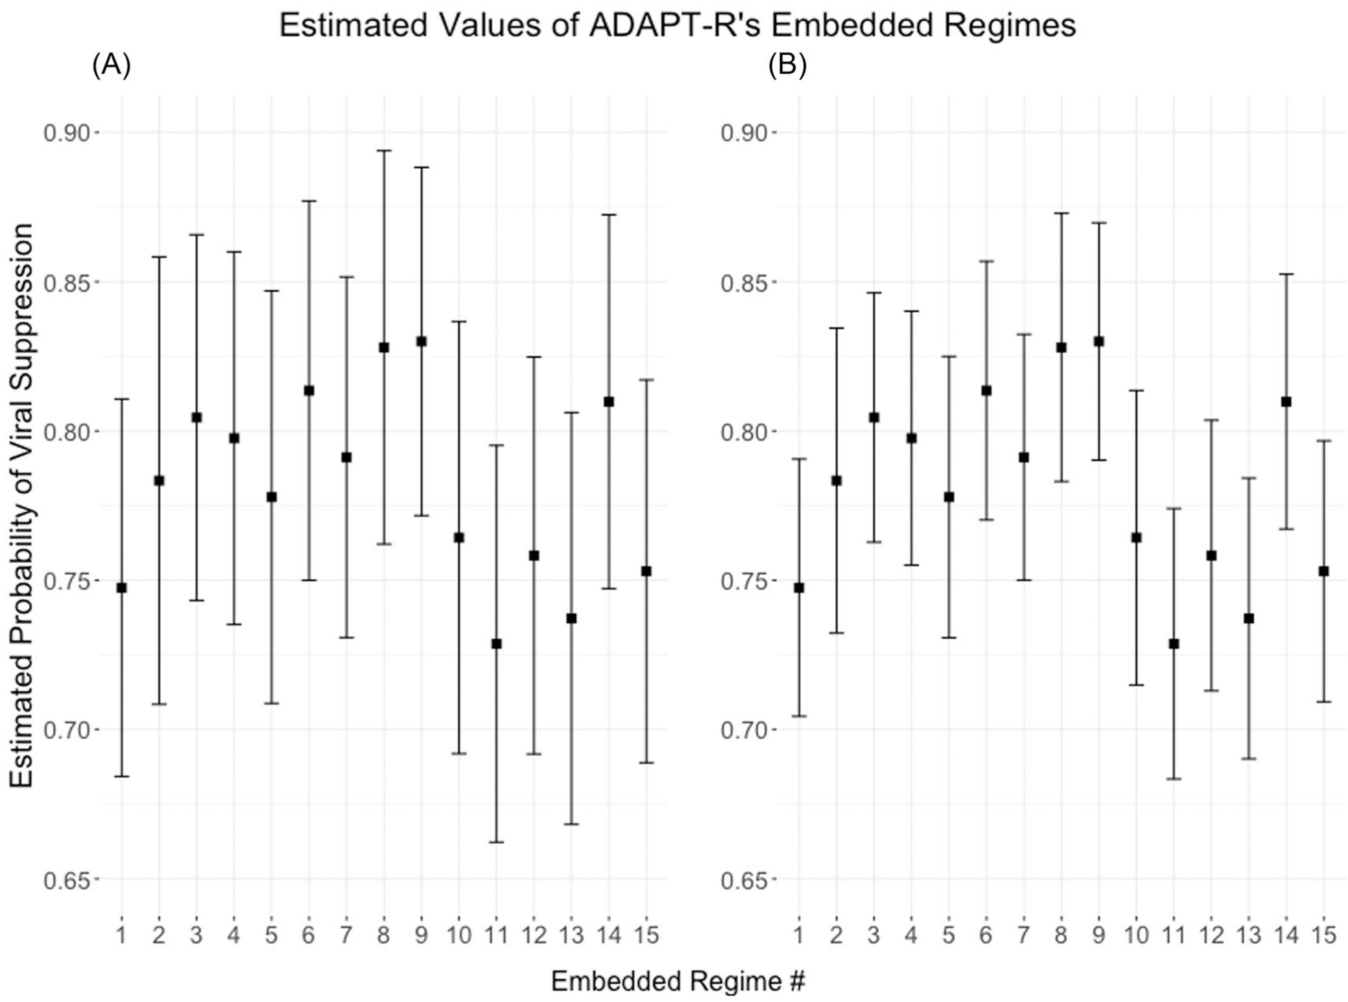

FIGURE 4.

Analysis of the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) study. Estimates of the probability of viral suppression under each of ADAPT-R's 15 embedded regimes are listed in Table 1. The squares are targeted maximum likelihood estimator (TMLE) point estimates and the error bars are 95% confidence intervals on these point estimates (simultaneous and single confidence intervals in panels A and B, respectively). We note that these point estimates vary slightly from pre-specified analyses in ADAPT-R in that the latter used SuperLearner rather than logistic regressions for estimation of the treatment mechanism

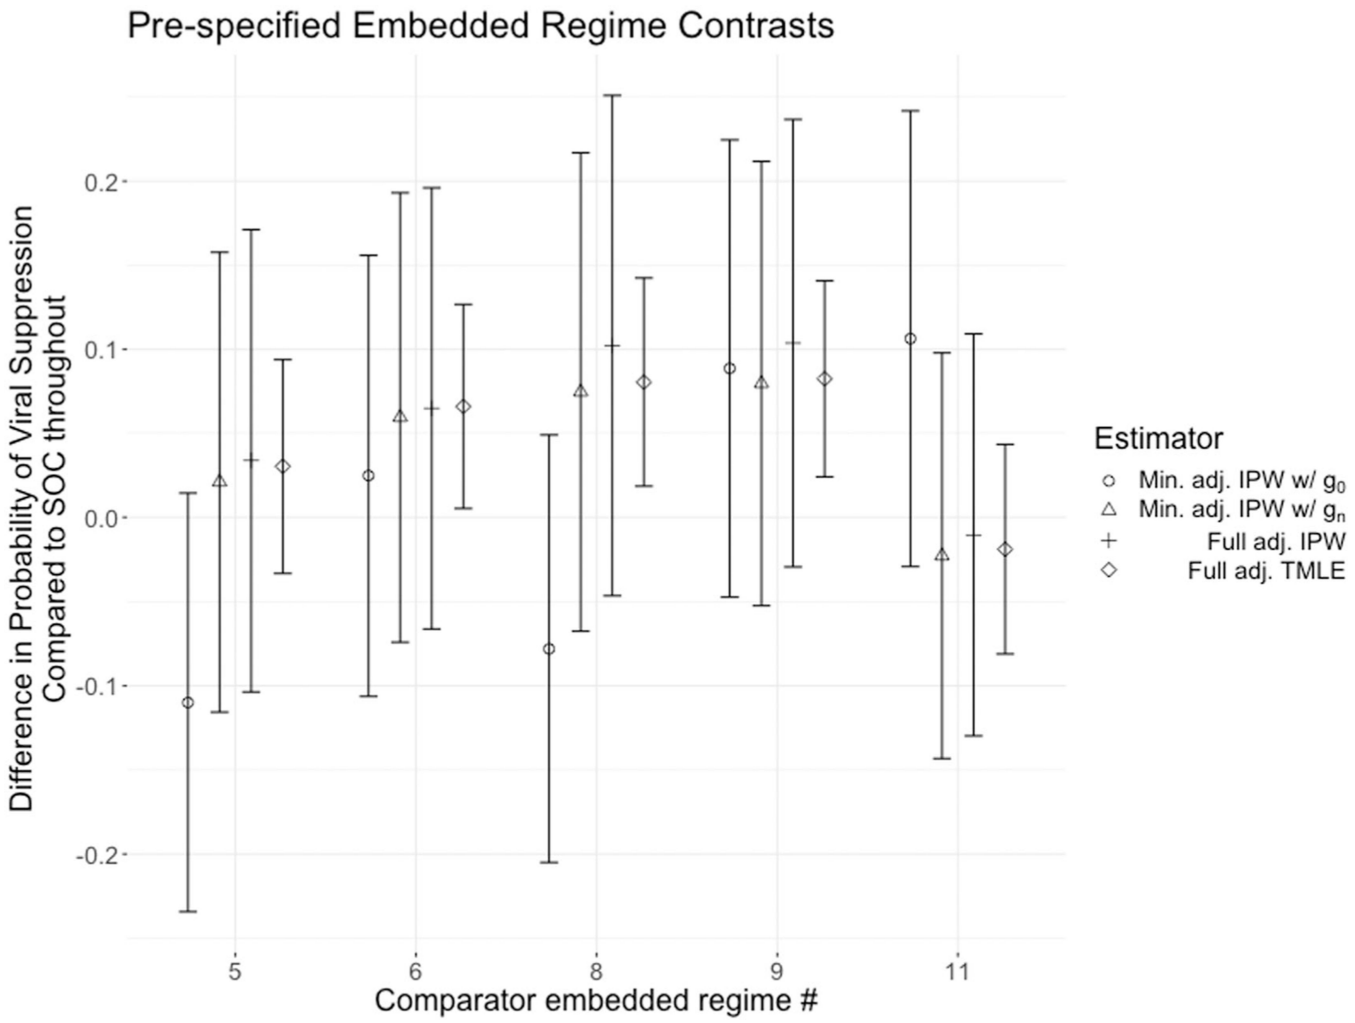

FIGURE 5.

Pre-specified contrast analysis of the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) study. Estimates of the difference in probability of viral suppression for the following pre-specified rules compared to standard of care (SOC) throughout: (1) short message service (SMS) with continuation if no lapse and addition of conditional cash transfer (CCT) if a lapse occurred (embedded regime number 5); (2) CCT with continuation if no lapse and addition of SMS if lapse occurred (embedded regime number 6); (3) SMS with continuation if no lapse and replacement with navigator if lapse occurred (regime number 8); (4) CCT with continuation of no lapse and replacement of navigator if lapse occurred (regime number 9); and (5) initial CCT, with SOC outreach if a lapse occurred and discontinuation if no lapse occurred (regime number 11). Shapes are point estimates (and error-bars are influence curve-based individual confidence intervals), which were generated with: (1) an inverse probability weighted (IPW) estimator with weights based on the true, known probability of receiving treatment given the initial treatment and lapse response ("Min. adj IPW w/g_0"); (2) an IPW estimator with estimated weights based on the empirical proportion of receiving treatment given the initial treatment and lapse response ("Min. adj IPW w/g_n"); (3) an IPW estimator with estimated weights that

adjust for all covariates (“Full adj. IPW”); and (4) a targeted maximum likelihood estimator (TMLE) that adjusts for all covariates (“Full adj. TMLE”)

TABLE 1

List of 15 dynamic treatment regimes embedded within the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) study (i.e., ADAPT-R's 15 embedded regimes)

Embedded regime $(\tilde{d})$	**Stage 1**	**Stage 2 if lapse**	**Stage 2 if no lapse**
1	SOC	SOC outreach	Continue
2	SMS	SOC outreach	Continue
3	CCT	SOC outreach	Continue
4	SOC	SMS + CCT	Continue
5	SMS	SMS + CCT	Continue
6	CCT	SMS + CCT	Continue
7	SOC	Navigator	Continue
8	SMS	Navigator	Continue
9	CCT	Navigator	Continue
10	SMS	SOC outreach	Discontinue
11	CCT	SOC outreach	Discontinue
12	SMS	SMS + CCT	Discontinue
13	CCT	SMS + CCT	Discontinue
14	SMS	Navigator	Discontinue
15	CCT	Navigator	Discontinue