HHS Public Access

Author manuscript

Stat Med. Author manuscript; available in PMC 2026 January 14.

Published in final edited form as:

Stat Med. 2025 December ; 44(28-30): e70353. doi:10.1002/sim.70353.

Effects Among the Affected

Lina M. Montoya[1,2], **Elvin H. Geng**[3], **Michael Valancius**[2], **Michael R. Kosorok**[2,4], **Maya L. Petersen**[5]

[1]School of Data Science and Society, University of North Carolina at Chapel Hill, NC, U.S.A.

[2]Department of Biostatistics, University of North Carolina at Chapel Hill, NC, U.S.A.

[3]Division of Infectious Diseases, Washington University in St. Louis, MO, U.S.A.

[4]Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, NC, U.S.A.

[5]Division of Biostatistics, University of California, Berkeley, CA, U.S.A.

Abstract

We propose a novel causal estimand that elucidates how response to an earlier treatment (e.g., treatment initiation) modifies the effect of a later treatment (e.g., treatment discontinuation), thus learning if there are effects among the (un)affected. Specifically, we consider a working marginal structural model summarizing how the average effect of a later treatment varies as a function of the (estimated) conditional average effect of an earlier treatment. We define the estimand to be a data-adaptive causal parameter, allowing for estimation of the conditional average treatment effect using machine learning without making strong smoothness assumptions. We show how a sequentially randomized design can be used to identify this causal estimand, and we describe a targeted maximum likelihood estimator for the resulting statistical estimand, with influence curve-based inference. We present simulation studies that evaluate the performance of this estimator under various finite-sample scenarios. Throughout, we use the "Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care" trial (NCT02338739) as an illustrative example, showing that discontinuation of conditional cash transfers for HIV care adherence was most harmful among those who most had an increase in benefits from them initially.

Keywords

conditional average treatment effect; data-adaptive parameter; marginal structural model; sequential multiple assignment randomized trial; targeted maximum likelihood estimation

1 | INTRODUCTION

Many treatments appear to be beneficial to receive *and* harmful to stop – yet they must be delivered in a time-limited way. For example, long-term opioid use for chronic pain

Correspondence Corresponding author Lina M. Montoya. lmontoya@unc.edu.

CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

management can lead to dependence and tolerance. Even with proper tapering, patients who initially benefited may experience withdrawal symptoms and increased pain sensitivity, making pain management more challenging post-discontinuation[1]. Programs that promote healthy behaviors in workplace settings may lose their effectiveness once the program ends, as some employees initially helped return to their old habits if the motivation was primarily driven by program incentives[2]. Nutritional support initiatives can improve dietary habits, but discontinuation may lead to a return to previous, less healthy eating habits, especially in segments of the population that depend on these programs for access to healthier food options[3].

Whether or not to deem a time-limited intervention worthwhile for a population traditionally comes down to the simple question of whether, on average, it provides a net benefit; in other words, whether the expected benefit of starting outweighs the expected harm of stopping. However, such an apparently straightforward result risks missing nuances that can come into focus once examined through the lens of precision medicine[4] and, more generally, treatment effect heterogeneity – specifically, if the harms of discontinuing an intervention may occur in different (or the same) people as those for whom the intervention is initially helpful. Consider a net beneficial intervention. It may be the case that initiating the intervention has no effect on some but helps others; and, once removed, the harms of discontinuation only or primarily accrue among those who initially benefited from the intervention. In other words, the time-limited intervention may provide a temporary benefit for some, a sustained benefit to others, but harms no one. On the other hand, it may be the case that the harm of treatment discontinuation accrues in people who were not benefited to begin with; i.e., the persons harmed by treatment removal are distinct from those helped by starting.

Identifying and distinguishing between these two scenarios, and more generally quantifying the extent to which the harms of stopping an intervention are correlated with the benefits of starting it, is important both for understanding the impacts of deploying an intervention and for understanding its mechanism of action. An intervention that provides a time-limited benefit to some, and perhaps a sustained benefit to a smaller number, could be deployed to a full population (cost and resource constraints aside) without risking harm to any. In contrast, an intervention whose removal harms those not benefited to begin with requires additional tailoring before deployment to avoid doing harm.

Longitudinal data can allow us to quantify the extent to which the benefit of treatment initiation modifies the effect of subsequent treatment discontinuation. In particular, Sequential, Multiple Assignment Randomized Trials (SMARTs) permit asking such causal questions because of the conditional randomization that occurs at several stages. Often, a SMART's sequential conditional randomization is leveraged to identify treatment effect heterogeneity, by baseline and time-varying covariates, that informs point-treatment or longitudinal (optimal) dynamic treatment regimes (e.g., see Kidwell and Almirall, 2023[5]; Kosorok and Moodie, 2015[6]; and Murphy, 2005[7]). However, in this paper, we suggest that SMARTs are also particularly well-suited to answer questions about the heterogeneous distribution of the harms and benefits from time-limited treatments.

As an illustrative example, consider the SMART called Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial (NCT02338739), whose primary objective was to identify individualized strategies for helping patients remain retained in their HIV care (Figure 1)[8]. In this study, adults living with HIV in rural Kenya were sequentially randomized at two time-points to receive incentives that aimed to reduce the occurrence and/or recurrence of a lapse in their care (defined as ≥ 14 days late to a clinic visit). Specifically, at each stage, patients were randomized to receive either a conditional cash transfer (CCT) for on-time clinic appointments or a non-CCT based intervention (such as SMS text messages, peer navigators, or standard-of-care treatment). Under this design, it is straightforward to identify heterogeneous treatment effects of, for example, discontinuing CCTs at the second time-point after receiving them at the first time-point by baseline (e.g., age, sex, SES) and/or time-varying (e.g., viral load status, updated pregnancy status) variables. Indeed, there exist a range of machine learning algorithms to estimate such heterogeneous treatment effects (see, e.g., Kennedy, 2023,[9]; Semenova and Chernozhukov, 2021[10]; and Wager and Athey, 2018[11]).

In this paper, we explore this type of downstream, second-stage effect modification not by a set of variables themselves, but rather by a function of initial treatment response – namely, the conditional average treatment effect (CATE). This allows us to ask our original question: how do the effects of discontinuation differ between those for whom treatment initiation is expected to benefit or harm? Or if treatment was helpful for all, how does the effect of discontinuation differ between those with large versus attenuated benefits? For example, in ADAPT-R, it was found that CCTs were beneficial on average – marginally and conditional on baseline covariates. Further, it was found that, among people who did not have an initial lapse in care, continuing CCTs was beneficial, on average, compared to discontinuing CCTs. Here, we are interested in uncovering differential effects of continuing versus discontinuing CCTs as a function of participants' initial expected benefit from CCT initiation. In this way, we are able to investigate incentive discontinuation harms among those for whom the likelihood of initial benefits are large versus small. It may be the case, for example, that having an initial strong effect of CCTs leads to greater "withdrawal" of CCTs – which may be expected if CCTs are operating primarily via directly addressing structural barriers present only for a subset of the population. Or it may be the opposite: having initial strong effects of CCTs may be protective against the effects of their removal, which could occur if the CCTs' primary mechanism was via increasing intrinsic motivation for different segments of the population. Beyond mechanistic insights, a finding of maximal harm among persons with minimal initial benefit would highlight the risks of deploying a time-limited CCT intervention at scale to a subset of the population (without additional tailoring at either the initiation or discontinuation stages).

To this end, we introduce a causal parameter that helps us understand differential downstream consequences of initial, upstream treatment effects. In particular, we consider a working marginal structural model (MSM)[12,13] that is a function of an estimated CATE, and define a data-adaptive parameter of interest[14] defined by this working model. Following the causal roadmap[15], we then identify this parameter as a function of the observed data distribution, and present an estimation procedure for our target parameter using Targeted

Maximum Likelihood Estimation (TMLE)[16,17], with influence-curve based inference on estimates. This type of analysis can allow us to examine effects among the (un)affected.

The rest of this paper is organized as follows. In Section 2, we describe the data generated from a SMART and the causal model that generates this data, with a description of the ADAPT-R data, which motivates the development of these methods. In Section 3, we define our causal question of interest (are there effects among the [un]affected?), as well as the target causal parameter that corresponds to the specified causal question of interest (a parameter of a working MSM) that summarizes how the effect of treatment discontinuation varies as a function of initial treatment response. Specifically, we use machine learning to estimate the conditional average treatment effect (CATE) for treatment initiation, and the estimated CATE as a baseline effect modifier, resulting in a data-adaptive target parameter. In Section 4, we discuss the conditions needed for identification of the causal parameter of interest, and how data generated from a SMART can facilitate this identification process. In Section 5, we discuss a TMLE of the parameters of the identified working MSM, as well as influence-curve based inference for this estimator. In Section 6, we present two simulation studies illustrating performance of this estimator in finite samples. In Section 7, we apply these methods to the ADAPT-R study. We close with a discussion.

2 | DATA AND CAUSAL MODEL

The following data are generated from a K-stage SMART for a time $t \in \{0, \ldots K\}$, where K corresponds to the number of randomization stages: 1) categorical interventions $A(t) \in \mathscr{A}_t$ (which could include missingness or measurement variables, such as a time-varying indicator of right-censoring); and, 2) non-intervention variables $L(t) \in \mathscr{L}_t$ between interventions at time $t-1$ and t, which include baseline covariates $L(0)$, time-varying covariates $L(t)$, $t \neq \{0, K\}$ (which includes intermediate outcomes $Y(t) \in \mathbb{R},\ t \neq \{0, K\}$), and a final outcome $L(K) = Y(K)$. Here, $L(-1) = A(-1) = A(0) = \varnothing$ and overbars are used to denote a variable's past history, i.e., $\bar{A}(t) = (A(1), \ldots, A(t))$ and $\bar{L}(t) = (L(0), \ldots, L(t))$.

The subset of observed variables that are used to assign treatment at time t according to the SMART's design are denoted $\bar{Z}(t) \subseteq (\bar{A}(t-1), \bar{L}(t-1))$. With this, the following structural causal model (SCM, denoted $\mathscr{M}^F$) describes the process that gives rise to data generated from a SMART design[18], where the random variables in $\mathscr{M}^F$ follow the joint distribution $P_{U,X} \in \mathscr{M}^F$:

$$L(t) = f_{L(t)}\left(U_{L(t)}, \bar{L}(t-1), \bar{A}(t)\right)$$

$$A(t) = f_{A(t)}\left(U_{A(t)}, \bar{Z}(t)\right),$$

for $t = 0, \ldots, K$. Here, $U = \left(U_{L(t)}, U_{A(t)}\right)$ represents the unmeasured random input to the data generating system at time t, while $X = (\bar{L}(K), \bar{A}(K))$ represents the endogenous variables in the causal model. In a SMART, the functions $f_{A(t)}$ are known for all t, as they are the

randomization scheme used in the SMART. Further, $U_{A(t)}$ is known by design in a SMART to be independent of all other unobserved error terms.

3 | CAUSAL QUESTION AND PARAMETERS

The causal question of interest asks: do the initial effects of an earlier intervention modify the effects of a subsequent intervention? In this section, we will translate this causal question into a target causal parameter $\Psi^F(P_{U,X})$, where Ψ^F is a mapping from the causal model to the parameter space $\mathbb{R}$. Following are four causal estimands we list in an order that shows how each responds to the limitations of the previous one listed, illustrating the logic we used to arrive at our target causal estimand (Section 3.4) inspired by an ideal, though un-identifiable, causal estimand (Section 3.1).

3.1 | Causal Estimand 1: Later effects modified by earlier individualized effects (the ideal causal estimand)

To begin to answer whether initial effects of an earlier intervention modify the effects of a subsequent intervention, we first intervene on the SCM to generate individualized, initial effects of interest at time r; for $t = 0, \ldots, r-1, r+1, \ldots, K$:

$$L(t) = f_{L(t)}\left(U_{L(t)}, \bar{L}(t-1), \bar{A}(t)\right)$$

$$A(t) = f_{A(t)}\left(U_{A(t)}, \bar{Z}(t)\right)$$

$$A(r) = a(r) \in \{0, 1\}$$

$$L(r) = f_{L(r)}\left(U_{L(r)}, \bar{L}(r-1), \bar{A}(r)\right),$$

where the support of $A(r)$, i.e., $\mathscr{A}_r$, must be $\{0, 1\}$. In this way, we can define an individualized effect at time r as: $ITE_{P_{U,X}}(r) = Y(r)_{a(r)=1} - Y(r)_{a(r)=0}$, where $Y(r)_{a(r)}$ is the outcome at time r if, possibly contrary to fact, an individual had received treatment $a(r)$. At time r, individuals with $ITE_{P_{U,X}}(r) > 0$ benefited from the treatment, those with $ITE_{P_{U,X}}(r) < 0$ were harmed by the treatment, and those with $ITE_{P_{U,X}}(r) = 0$ experienced no benefit or harm from the treatment.

We note that, here, we define "individualized effects of an earlier intervention" as ones that occur at one timepoint before time s; namely, those that occur at only time r. Alternatively, one could define such individualized effects as the cumulative effect of receiving a treatment until and at time r versus never receiving a treatment, i.e., $Y(r)_{\bar{a}(r)=1} - Y(r)_{\bar{a}(r)=0}$. Similar modifications to the causal parameter would follow, sequential identification criteria would be needed for identification of the resultant causal estimand (i.e., Causal Estimand 4), and

additional longitudinal methods would be required for estimation. For simplicity, however, we focus on point-treatment individualized initial effects here.

Next, we intervene on the SCM to understand the effects of a later intervention as a function of the earlier individualized effects generated in the previous step. To do this, we consider an intervention on the SCM at time $s > r$; for $t = 0, \ldots, s-1, s+1, \ldots, K$:

$$L(t) = f_{L(t)}\left(U_{L(t)}, \bar{L}(t-1), \bar{A}(t)\right)$$

$$A(t) = f_{A(t)}\left(U_{A(t)}, \bar{Z}(t)\right)$$

$$A(s) = a(s) \in \mathscr{A}_s$$

$$L(s) = f_{L(s)}\left(U_{L(s)}, \bar{L}(s-1), \bar{A}(s)\right),$$

to define the expected counterfactual outcome at time s had all individuals received treatment $a(s)$, conditional on individual-level effects at time r, i.e.:

$$\mathbb{E}_{P_{U,X}}[Y(s)_{a(s)} \mid ITE_{P_{U,X}}(r)]. \tag{1}$$

This can be interpreted as the expected outcome $Y(s)$ under an intervention $a(s)$, conditional on individualized effects at time r – where individualized effects are defined as individual-level benefits and harms.

Let P_0 be the true distribution of our observed data O, an element of $\mathscr{M}$, the statistical model. We assume the observed data $O_i \equiv \left(\bar{L}(K)_i, \bar{A}(K)_i\right) \sim P_0 \in \mathscr{M}, i = 1, \ldots, n$ were generated by sampling n i.i.d. copies from a data-generating system contained in $\mathscr{M}^F$ above. Functions of the observed data distribution are denoted with a subscript 0.

It is well known that $ITE_{P_{U,X}}(r)$ is not identified as a function of P_0, as this is the fundamental problem of causal inference[19]: both $Y(r)_{a(r)=1}$ and $Y(r)_{a(r)=0}$ of $ITE_{P_{U,X}}(r)$ can never both be observed at a given time. Thus, this ideal causal estimand, Equation 1, cannot be identified, and $ITE_{P_{U,X}}(r)$ cannot be written as $ITE_0(r)$, which is the goal of identification. For this reason, in the next section, we focus on a causal estimand that, given sufficient covariate richness, aims to move close to Equation 1 while remaining identifiable.

3.2 | Causal Estimand 2: Later effects modified by earlier conditional average treatment effects

Since $ITE_{P_{U,X}}(r)$ cannot be identified as a function of the observed data distribution, the estimand in Equation 1 cannot be identified. For this reason, we seek to find an identifiable alternative to $I_{TE_{PU,X}}(r)$ and thus Equation 1 that answers our causal question.

When the covariates $\bar{L}(r-1)$ are sufficiently rich so that ITE can be predicted with great precision, the $ITE_{P_{U,X}}(r)$ is well approximated by the CATE. In other words, $ITE_{P_{U,X}}(r)$ could be approximated by the conditional average treatment effect (CATE) at time r:

$$B_{P_{U,X}}(H(r)) \equiv \mathbb{E}_{P_{U,X}}[Y(r)_{a(r)=1} - Y(r)_{a(r)=0} \mid H(r)], \tag{2}$$

where $H(r)$ is patient history prior and up to time r, i.e., $H(r) \equiv (\bar{L}(r-1), \bar{A}(r-1))$. Equation 2, the CATE, at $H(r) = h(r)$ can be interpreted as the expected treatment effect given the covariate/treatment history profile $h(r)$. At time r, if $B_{P_{U,X}}(h(r)) > 0$, individuals with a history profile of $h(r)$ benefited from the treatment in expectation; if $B_{P_{U,X}}(h(r)) < 0$, individuals with a history profile $h(r)$ were harmed by the treatment in expectation; and if $B_{P_{U,X}}(h(r)) = 0$, individuals with a history profile of $h(r)$ experienced no benefit or harm from the treatment in expectation. Discussed in more detail in Section 3.4, this CATE can be identified in a straightforward manner under a SMART design.

Equation 1 could then be used as a way to approach the following causal estimand:

$$\mathbb{E}_{P_{U,X}}[Y(s)_{a(s)} \mid B_{P_{U,X}}(h(r))], \tag{3}$$

which can be interpreted as the true average effect of $a(s)$ on $Y(s)$, conditional on individualized effects at time r – where individualized effects are defined as expected benefits and harms in patient history strata $h(r)$.

It may be the case that the inputs needed to define the full space of possible outputs defined by Equation 3 are finite and small; e.g., that $H(r)$ consists of a small number of discrete covariates, each with a small number of levels, and treatment histories that take on a small number of possible values. In this way, we could potentially evaluate Equation 3 for every treatment and covariate combination, the result being a list of expected counterfactual outcomes under every treatment level $a(s) \in \mathscr{A}_s$, given every possible CATE value.

There are limitations to this approach. First, given that we have required covariate richness to go after a causal parameter that answers our causal question, it is more likely that $H(r)$ will consist of several covariates, many of them continuous, as in ADAPT-R (see Section 7.1). In that case, Equation 3 would instead be a continuous curve and it would be impossible to evaluate and list its every possible value. Along the same lines, it may be

the case that there are small or empty cells in our data at hand for each of the necessary covariate and treatment values that define Equation 3, rendering this causal parameter potentially inestimable (i.e., the curse of dimensionality) without smoothing over $a(s)$ and $B_{P_{U,X}}(h(r))$. Second, even if we could list all possible outputs of Equation 3, it is more difficult to interpret a list of outputs rather than (parameters of a) single function that generates those outputs. For these reasons, rather than listing every possible output of Equation 3 by every $a(s)$ and $B_{P_{U,X}}(h(r))$, it would be more desirable to know the true underlying "effect among affected" curve that generates each possible value of Equation 3.

3.3 | Causal Estimand 3: A projection of the true causal function onto a working marginal structural model

We wish to know the true summary of the parameters $\{\mathbb{E}_{P_{U,X}}[Y(s)_{a(s)} \mid B_{P_{U,X}}(h(r))] : a(s), B_{P_{U,X}}(h(r))\}$. This true, "effect among affected" curve is described by the function $m^*: (\mathscr{A}_t, \mathscr{L}_t) \to \mathbb{R}$, where:

$$m^*\left(a(s), B_{P_{U,X}}(H(r))\right) \equiv \mathbb{E}_{P_{U,X}}[Y(s)_{a(s)} \mid B_{P_{U,X}}(H(r))]$$

which is the causal curve of interest.

One approach for understanding $m^*\left(a(s), B_{P_{U,X}}(H(r))\right)$ is to assume it follows a parametric MSM[12]; for example, if $Y(s) \in [0, 1]$, one could assume the following logistic model:

$$\text{logit } m^*\left(a(s), B_{P_{U,X}}(H(r))\right) = \beta_0 + \beta_1 a(s) + \beta_2 B_{P_{U,X}}(H(r)) + \beta_3 a(s) B_{P_{U,X}}(H(r)). \tag{4}$$

With such a parametric MSM, we could summarize the true "effect among affected" curve, and potentially smooth across covariate and treatment values for which there is no representation in the data, as desired. Here, β_3 could be interpreted causally as quantifying the degree of modification of the effect of $A(s)$ on $Y(s)$ by initial, individualized effects of $A(r)$ on $Y(r)$.

The issue here, however, is that we have no knowledge that Equation 4 holds true, and thus this assumption may be incorrect. We note that if we did think this equality were true, we would have explicitly stated it in our causal model, e.g., in the functional form for $f_{L(s)}$ in our SCM. But by making this parametric assumption at this point, we are not respecting our causal model, which could mean this assumption is incorrect and β_3 is not defined.

We can alternatively define a *working* MSM, as described by Neugebauer and van der Laan[13] and Rosenblum and van der Laan[20]. Instead of assuming that $m^*\left(a(s), B_{P_{U,X}}(H(r))\right)$ follows the functional form described in Equation 4, we instead define a working MSM, $m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)$, that follows the functional form:

$$\text{logit } m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right) \equiv \beta_0 + \beta_1 a(s) + \beta_2 B_{P_{U,X}}(H(r)) + \beta_3 a(s) B_{P_{U,X}}(H(r)).$$

We can then define our updated causal quantity of interest as the projection of $m^*\left(a(s), B_{P_{U,X}}(H(r))\right)$ onto $m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)$ according to:

$$\boldsymbol{\beta}_{P_{U,X}} = \arg\min_{\boldsymbol{\beta}} - \sum_{a(s) \in \mathscr{A}_s} \mathbb{E}_{P_{U,X}}\Big[h\left(a(s), B_{P_{U,X}}(H(r))\right)\Big(Y(s)_{a(s)} \log m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)\Big) + \left(1 - Y(s)_{a(s)}\right) \log\left(1 - m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)\right)\Big],$$

where $h\left(a(s), B_{P_{U,X}}(H(r))\right)$ is a user-specified stabilizing weight function. For example, $h\left(a(s), B_{P_{U,X}}(H(r))\right)$ could be $P(A(2) = 1 \mid B_n(H(r)))$ or simply set to 1 (the choice of weight stabilizing function in an MSM is discussed in, e.g., Petersen et al.[16]); we proceed with the latter, because our scientific question gives equal weight to all treatments at time s and blip values at time r. Then, $\boldsymbol{\beta}_{P_{U,X}}$ is the solution to the equation:

$$0 = \mathbb{E}_{P_{U,X}}\left[\frac{\frac{d}{d\boldsymbol{\beta}_{P_{U,X}}} m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s),\ B_{P_{U,X}}(H(r))\right)}{m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s),\ B_{P_{U,X}}(H(r))\right)\left(1 - m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s),\ B_{P_{U,X}}(H(r))\right)\right)} \times \left(Y(s)_{a(s)} - m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s), B_{P_{U,X}}(H(r))\right)\right)\right],$$

which, by the law of total expectation, can also be written as:

$$0 = \mathbb{E}_{P_{U,X}}\left[\frac{\frac{d}{d\boldsymbol{\beta}_{P_{U,X}}} m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s),\ B_{P_{U,X}}(H(r))\right)}{m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s),\ B_{P_{U,X}}(H(r))\right)\left(1 - m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s),\ B_{P_{U,X}}(H(r))\right)\right)} \times \left(\mathbb{E}_{P_{U,X}}\left[Y(s)_{a(s)} \mid H(s)\right] - m_{\boldsymbol{\beta}_{P_{U,X}}}\left(a(s),\ B_{P_{U,X}}(H(r))\right)\right)\right],$$

where $H(s) \equiv (\bar{L}(s-1), \bar{A}(s-1))$. Thus, we can instead say we are interested in the projection of $m^*\left(a(s), B_{P_{U,X}}(H(r))\right)$ onto $m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)$ according to:

$$\boldsymbol{\beta}_{P_{U,X}} = \arg\min_{\boldsymbol{\beta}} - \sum_{a(s) \in \mathscr{A}_s} \mathbb{E}_{P_{U,X}}\Big[\left(\mathbb{E}_{P_{U,X}}\left[Y(s)_{a(s)} \mid H(s)\right] \log m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)\right) + \left(1 - \mathbb{E}_{P_{U,X}}\left[Y(s)_{a(s)} \mid H(s)\right]\right) \log\left(1 - m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)\right)\Big]. \tag{5}$$

Here, we may be interested in $\beta_3 \in \boldsymbol{\beta}_{P_{U,X}}$ from this working MSM, representing the degree of modification of the effect of $A(s)$ on $Y(s)$ by initial, individualized effects of $A(r)$ on $Y(r)$. We reiterate that choosing a working MSM, such as the logistic model as we have done here, defines our causal parameter of interest, but it does not represent a causal or statistical assumption.

3.4 | Causal Estimand 4: A data-adaptive causal parameter (the target causal estimand)

To recap, we have now defined a potential causal parameter of interest as the projection of $m^*\left(a(s), B_{P_{U,X}}(H(r))\right)$ onto $m_{\boldsymbol{\beta}}\left(a(s), B_{P_{U,X}}(H(r))\right)$ according to Equation 5. The true causal curve, $m^*\left(a(s), B_{P_{U,X}}(H(r))\right)$, depends on the true CATE, $B_{P_{U,X}}(H(r))$. In practice, however, it may be the case that we are ultimately interested in the true causal curve that depends on the estimated CATE from a specific sample, not the true unknown CATE of the population. As we will discuss later, going after such a parameter may additionally be appealing, as it relaxes assumptions needed for consistent estimation of our ultimate causal parameter. To define such a data-adaptive parameter, we first need to identify the CATE as a function of P_0.

In a SMART, the randomization condition $Y(t)_{a(t)} \perp A(t) \mid \bar{Z}(t)$ holds for $t = 1, \ldots, K$, because $A(t)$ is conditionally randomized based on values of $\bar{Z}(t)$ (which may be $\varnothing$, in which case $A(t)$ is marginally randomized). Informally, this means there are no unmeasured common causes between $Y(t)$ and $A(t)$ at each timepoint, conditional on $\bar{Z}(t)$. We can extend this to state $Y(r)_{a(r)} \perp A(r) \mid H(r)$.

According to the SMART design, at certain $\bar{Z}(t) = \bar{z}(t)$ it is the case that $g_{A(t),0}(A(t) = a(t) \mid \bar{Z}(t)) > 0 - a.e.$ holds for each $t = 1, \ldots, K$ and $a(t) \in \mathscr{A}_t$, where $g_{A(t),0}(A(t) \mid \bar{Z}(t) = \bar{z}(t)) = P_0(A(t) \mid \bar{Z}(t) = \bar{z}(t))$, the true conditional probability of the treatment at time t given accrued treatment and covariate variables that defines the SMART's design. We consider CATEs at time r for which the relevant positivity condition is not violated, in which case we can say $g_{A(r),0}(A(r) = a(r) \mid H(r)) > 0 - a.e.$ for $a(r) \in \{0,1\}$. Informally, this positivity condition states that, conditional on patient history (including randomization triggers), there must be a positive probability of receiving treatment $A(t) = a(t)$. It is important to consider this positivity condition to understand which CATEs at time r, i.e., $B_{P_{U,X}}(H(r))$, are feasible to examine. For example, in ADAPT-R, the CATE of $A(1) \in \{\text{CCT}, \text{ no CCT}\}$ on $Y(1)$ is identifiable because $A(1)$ was marginally randomized and thus all patients could have theoretically received all possible treatments in $\mathscr{A}_1$ (see Section 7.2.1).

Both the randomization and positivity conditions permit identifying the initial CATE function (i.e., Equation 2) as the so-called "blip" function: $B_{P_{U,X}}(H(r)) = B_0(H(r)) \equiv \mathbb{E}_0\left[Y(r) \mid A(r) = 1, H(r)\right] - \mathbb{E}_0\left[Y(r) \mid A(r) = 0, H(r)\right]$. Let $B_n(H(r))$ be an estimate of $B_0(H(r))$, estimated on a finite sample (we discuss estimators for the blip in Section 5.1). This data set has empirical distribution P_n, giving each observation weight $\frac{1}{n}$.

In this way, we can define our *target* causal parameter as the projection of

$$m^*(a(s), B_n(H(r))) \equiv \mathbb{E}_{P_{U,X}}[Y(s)_{a(s)} \mid B_n(H(r))]$$

onto

$$\text{logit}\, m_{\boldsymbol{\beta}}(a(s), B_n(H(r))) \equiv \beta_0 + \beta_1 a(s) + \beta_2 B_n(H(r)) + \beta_3 a(s) B_n(H(r)) \tag{6}$$

according to:

$$\boldsymbol{\beta}_{P_{U,X}} = \arg\min_{\boldsymbol{\beta}} - \sum_{a(s) \in \mathscr{A}_s} \mathbb{E}_{P_{U,X}} \Big[\Big(\mathbb{E}_{P_{U,X}} \big[Y(s)_{a(s)} \mid H(s) \big] \log m_{\boldsymbol{\beta}}(a(s), B_n(H(r))) \Big) + \Big(1 - \mathbb{E}_{P_{U,X}} \big[Y(s)_{a(s)} \mid H(s) \big] \Big) \log \big(1 - m_{\boldsymbol{\beta}}(a(s), B_n(H(r))) \big) \Big] .$$

In particular, our target causal parameter of interest is $\Psi^F(P_{U,X}) = \beta_3 \in \boldsymbol{\beta}_{P_{U,X}}$ from this working MSM, quantifying the degree of modification of the effect of $A(s)$ on $Y(s)$ by initial, individualized effects of $A(r)$ on $Y(r)$ estimated on the finite sample in hand. We are interested in testing the null hypothesis that $\Psi^F(P_{U,X}) = 0$, i.e., there are no differential effects of discontinuing CCTs by on estimated individualized effects of initiating CCTs; the alternative hypothesis is $\Psi^F(P_{U,X}) \neq 0$. Here, we have kept $\mathscr{A}_s$ general such that $A(s)$ could take on more than two levels; however, in Appendix B we present a lower-dimensional working MSM that may be of interest if $\mathscr{A}_s \in \{0, 1\}$.

4 | IDENTIFICATION OF THE TARGET CAUSAL PARAMETER

Define the statistical target parameter of interest corresponding to the target causal parameter as a mapping $\Psi : \mathscr{M} \to \mathbb{R}$. With the point-treatment randomization $Y(s)_{a(s)} \perp A(s) \mid H(s)$ and positivity $g_{A(s),0}(A(s) = a(s) \mid H(s)) > 0 - a.e.$ conditions holding for all $a(s) \in \mathscr{A}_s$, we can identify the true causal effect of $A(s)$ on $Y(s)$ conditional on the estimated blip at time r as a function of the observed data distribution P_0; that is: $\mathbb{E}_{P_{U,X}}[Y(s)_{a(s)} \mid B_n(H(r))] = \mathbb{E}_0[\mathbb{E}_0[Y(s) \mid A(s) = a(s), H(s)] \mid B_n(H(r))]$. As before, both in SMARTs and observational studies, some covariate/treatment combinations may be infeasible to examine, so it is imperative that relevant positivity conditions hold for the causal parameter of interest. We note that this identification procedure did not require the sequential randomization and positivity conditions, but rather the point-treatment randomization and positivity conditions, as our causal question and thus parameter does not ask about simultaneous, joint effects of multiple $A(t)$s, but rather two interventions separately. If the causal question additionally requires intervening on other A nodes (as in the ADAPT-R example; see Section 7.2.2), sequential randomization $Y(t)_{\bar{a}(t)} \perp A(t) \mid \bar{L}(t-1), \bar{A}(t-1) = \bar{a}(t-1)$ and positivity $g_{A(t),0}(A(t) = a(t) \mid \bar{L}(t-1), \bar{A}(t-1) = \bar{a}(t-1)) > 0 - a.e.$ hold in a SMART, as well.

Supposing that $Y \in [0,1]$ and the causal parameter of interest is an element of a vector of coefficients in a logistic working MSM, then a target statistical parameter can be defined as the projection of $\mathbb{E}_0[\mathbb{E}_0[Y(s) \mid A(s) = a(s), H(s)] \mid B_n(H(r))]$ onto Equation 6 according to:

$$\beta_0 = \arg\min_{\beta} - \sum_{a(s) \in \mathcal{A}_s} \mathbb{E}_0 \Big[h(a(s), B_n(H(r))) \Big(\mathbb{E}_0 \big[Y(s) \mid A(s) = a(s), H(s) \big] \log m_\beta(a(s), B_n(H(r))) \Big) + (1 - \mathbb{E}_0[Y(s) \mid A(s) = a(s), H(s)]) \log \big(1 - m_\beta(a(s), B_n(H(r))) \big) \Big],$$

where β_0 solves the equation:

$$0 = \mathbb{E}_0 \left[\sum_{a(s) \in \mathcal{A}_s} h(a(s),\ B_n(H(r))) \frac{\frac{d}{d\beta_0} m_{\beta_0}(a(s),\ B_n(H(r)))}{m_{\beta_0}(a(s),\ B_n(H(r)))(1\ -\ m_{\beta_0}(a(s),\ B_n(H(r))))} \times \Big(\mathbb{E}_0 \big[Y(s) \mid A(s) = a(s), H(s) \big] - m_{\beta_0}(a(s),\ B_n(H(r))) \Big) \right].$$

The estimation goal is now defined: using our observed data, we aim to estimate $\beta_3 = \Psi(P_0)$, an element of β_0.

5 | ESTIMATION AND INFERENCE

This section is concerned with estimation and inference of $\Psi(P_0)$, the β_3 coefficient on the working MSM that signals whether the effects of downstream interventions on downstream outcomes are modified by conditional effects of initial interventions on initial outcomes that are estimated on the sample in hand. We propose using TMLE for estimation of the parameters of the working MSM, paired with influence curve-based inference to generate p-values and confidence intervals on those estimates.

5.1 | Estimation of the CATE/Blip Function

Many estimators exist for the blip function on (i.e., the identified CATE)[10,9,21,22]. For simplicity, we implement a simple plug-in/single-stage Q-learning estimator for the blip that requires estimating $Q_0^r(A(r), H(r)) \equiv \mathbb{E}_0[Y(r) \mid A(r), H(r)]$. Using any regression-based approach, such as SuperLearner[23], one can estimate Q_0^r, which we call Q_n^r. Importantly, the chosen approach should include algorithms that allow for interactions between $A(r)$ and $H(r)$, in order to capture any initial differential effects by $H(r)$. To generate an estimate of the CATE, predict at $A(r) = 1$ and $A(r) = 0$ to generate $Q_n^r(1, H(r))$ and $Q_n^r(0, H(r))$, respectively. This provides an estimate of the CATE/blip: $B_n(H(r)) = Q_n^r(1, H(r)) - Q_n^r(0, H(r))$.

5.2 | Estimation and Inference of the Target Parameter

We estimate our statistical parameter, $\Psi(P_0)$, using TMLE. We choose TMLE as our primary estimator because it is a double robust, substitution estimator. Other candidate estimators for $\Psi(P_0)$ include the g-computation[24] and inverse probability weighting (IPW)[25] estimators (which are not double robust), as well as the augmented IPW estimator[26] (which is not a substitution estimator). We refer the reader to van der Laan and Rose[27] for an introduction to TMLE and how its properties compare to other estimators.

In Appendix C, we provide a step-by-step procedure for implementing a TMLE for $\Psi(P_0)$, which we call $\hat{\Psi}(P_n) = \beta^*_{3,n}$, drawing from Petersen et al.[16] and Rosenblum and van der Laan[17]. The ltmle R package[28,16] also implements this TMLE, which we utilize for the simulations (Section 6) and ADAPT-R data analysis (Section 7).

We now elaborate on sufficient conditions under which the aforementioned TMLE procedure produces a $\sqrt{n}$ consistent and asymptotically normal estimator for our target parameter $\Psi(P_0)$. The general theory for data-adaptive estimators, which we apply to our setting, was originally presented in Hubbard et al.[14]. Let $g_{A(s),n}$ and $Q^{*,s}_n$ be estimators of $g_{A(s),0}$ and $Q^s_0(A(s), H(s)) \equiv \mathbb{E}_0[Y(s) \mid A(s), H(s)]$, respectively. When $g_{A(s),0}$ is known (as in a SMART design), it can be correctly estimated via a parametric model. We assume that $Q^{*,s}_n$ converges in probability to some Q^s (which may differ from Q^s_0). We require an additional smoothness assumption, namely that the efficient influence curve for $\boldsymbol{\beta}$ (called $D^*(P)$, presented in Appendix C) when applied to an empirical distribution is a P_0-Donsker class and converges in quadratic mean. Under such conditions, our estimator is $\sqrt{n}$ consistent and asymptotically normal with a consistent estimator of the variance of $\boldsymbol{\beta}_0$ given by $\Sigma_n = \mathbb{E}_n[D^*(Q^{*,s}_n, g_{A(s),n}, \boldsymbol{\beta}^*_n)^2]$. With this, we can construct asymptotically conservative 95% confidence intervals for $\Psi(P_0)$, i.e.: $\beta^*_{3,n} \pm \Phi^{-1}(0.975)\frac{\sqrt{\Sigma_n(3,3)}}{\sqrt{n}}$. We note that one could circumvent making these smoothness assumptions by going after a sample-split specific target parameter[14] estimated with a cross-validated TMLE, although at the posssible loss of interpretability.

6 | SIMULATION STUDIES

Using simulations, we evaluated the performance of the TMLE for the data-adaptive parameter $\Psi(P_0)$, the parameter representing whether there is modification of the effect of $A(s)$ on $Y(s)$ by initial estimated conditional effects $B_n(H(r))$. Specifically, we investigated performance (bias, variance, mean squared error, 95% confidence interval coverage, and type I error/power at the $\alpha = 0.05$ significance level) under 1,000 repetitions of sample size 1,815 (the original ADAPT-R sample size) under various generating processes: 1) a simple, two-stage SMART design intended to illustrate the differential effects this parameter could uncover (for the general case of any feasible downstream effects by initial effects – not necessarily with the same treatment options at both stages) and 2) a more complex data generating process that mimics the ADAPT-R study in which the treatment options are the same at both stages (i.e., initiate/not initiate treatment at the first stage and, among those who initiated treatment, discontinue/continue treatment at the second stage). All simulations were implemented in R[29], and the code, simulated data, and results can be found at https://github.com/lmmontoya/effects-among-affected. Specific data generating processes for both simulations can be found in Appendix D.

6.1 | Simulation 1

Consider three data generating processes (DGPs) from a simple two-stage SMART design. The average effect of the initial treatment $A(1)$ on the initial outcome $Y(1)$ for all DGPs is positive: $\mathbb{E}_{P_{U,X}}[Y(1)_{a(1)=1} - Y(1)_{a(1)=0}] \approx 0.14$, and the blip random variable for the DGPs can take on 4 possible values, of which 3 are positive, and 1 is negative. Let $Y^1(2)$ be the second-stage outcome for DGP 1, $Y^2(2)$ be the second-stage outcome for DGP 2, and $Y^3(2)$ be the second-stage outcome for DGP 3. This simulation is set up such that $\mathbb{E}_{P_{U,X}}[Y^q(2)_{a(2)=1} - Y^q(2)_{a(2)=0}] \approx 0.12$ for $q \in \{1,2,3\}$; i.e., the marginal second-stage risk differences are positive and the same for all data generating processes.

However, the first DGP illustrates a scenario in which the benefits of second-stage treatment accrue among those for whom the initial treatment is harmful. The risk difference conditional on a positive blip $\mathbb{E}_{P_{U,X}}[Y^1(2)_{a(2)=1} - Y^1(2)_{a(2)=0} \mid B_0(L(0)) > 0] \approx 0.08$, while the risk difference conditional on a negative blip $\mathbb{E}_{P_{U,X}}[Y^1(2)_{a(2)=1} - Y^1(2)_{a(2)=0} \mid B_0(L(0)) < 0] = 0.23$. The true β_3 coefficient corresponding to the working MSM in Equation 6 in a scenario where $B_n = B_0$ is approximately −1.92. We note that this value is not $\Psi(P_{U,X})$, our target causal quantity, as the latter is dependent on the finite sample at hand (a characteristic of data-adaptive parameters) and thus will be different for every simulation instance[14].

Conversely, the second DGP illustrates a scenario in which the benefits of second-stage treatment accrue among those for whom initial treatment is helpful. The risk difference conditional on a positive blip $\mathbb{E}_{P_{U,X}}[Y^2(2)_{a(2)=1} - Y^2(2)_{a(2)=0} \mid B_0(L(0)) > 0] \approx 0.15$, while the risk difference conditional on a negative blip $\mathbb{E}_{P_{U,X}}[Y^2(2)_{a(2)=1} - Y^2(2)_{a(2)=0} \mid B_0(L(0)) < 0] \approx 0.00$. Here, the true β_3 coefficient corresponding to the working MSM in Equation 6 in a scenario where $B_n = B_0$ is approximately 1.77.

Finally, the third DGP illustrates a null scenario in which there are no differential effects of the second-stage treatment by first-stage treatment effects. The risk difference conditional on a positive and negative blip is the same $\mathbb{E}_{P_{U,X}}[Y^3(2)_{a(2)=1} - Y^3(2)_{a(2)=0} \mid B_0(L(0)) > 0] = \mathbb{E}_{P_{U,X}}$. Here, the true β_3 coefficient is $[Y^2(2)_{a(2)=1} - Y^2(2)_{a(2)=0} \mid B_0(L(0)) < 0] \approx 0.12$ approximately 0.00.

Thus, the first-stage and second-stage treatment effects for all DGPs are identical, yet the differential downstream treatment effects by initial upstream treatment effects differ between the DGPs – a result that was revealed by our proposed methods and would not otherwise have been uncovered by simply looking at marginal effects of initiating or discontinuing treatment, nor by examining differential effects of that treatment by each baseline or time-varying covariate in turn.

For this simulation example, we estimated Q_0^2, $g_{A(2),0}$ and $B_0(L(0))$ using correctly specified parametric models. Detailed performance results can be found in Table 1. Under this estimator configuration, and with a sample size of 1,815, we saw 95.40% and 95.30% confidence interval coverage and were powered at 87.90% and 84.50% to test the null hypothesis that $\Psi(P_0) = 0$, corresponding, under assumptions, to a null hypothesis that the interaction between the estimated CATE for initial treatment and the effect of subsequent treatment is 0, for the DGPs corresponding to $Y^1(2)$ and $Y^2(2)$, respectively. For the third DGP, the null scenario, we saw 95.4% confidence interval coverage and the null was falsely rejected 4.60% of the time.

6.2 | Simulation 2

The second simulation study mimics the ADAPT-R trial. Under this DGP, $\mathbb{E}_{P_{U,X}}[Y(1)_{a(1)=1} - Y(1)_{a(1)=0}] = 0.0554$ and $\mathbb{E}_{P_{U,X}}[Y(2)_{a(1)=1,a(2)=1} - Y(2)_{a(1)=1,a(2)=0}] = 0.0564$. As in the question asked with ADAPT-R data, here we are interested in summarizing $\mathbb{E}_{P_{U,X}}[Y(2)_{a(2),a(1)=1} \mid B_n(L(0))]$ using a working MSM (discussed in more detail in Section 7), where our target causal parameter $\Psi(P_{U,X})$ is the β_3 coefficient according to Equation 6 and 7, below. Additionally, the true β_3 coefficient corresponding to this working MSM in a scenario where $B_n = B_0$ is approximately 2.36 – meaning that there is modification of CCT discontinuation effects by CCT initiation effects. Specifically, harms of discontinuation accrue among those for whom CCT initiation is most beneficial *and* benefits of discontinuation accrue among those for whom CCT initiation is least beneficial.

Here, the Q factors used to estimate the marginal structural parameters were estimated using mis-specified parametric logistic regressions, while the g factors were estimated using correctly specified parametric logistic regressions. The blip function was estimated with single-stage Q-learning, where the Q factors were estimated using SuperLearner with a diversity of candidate algorithms ranging from the mean to a simple, parametric models to tree-based methods, with many allowing for interactions between $A(1)$ and $Y(1)$ (a detailed list of candidate algorithms included in the library can be found in Appendix D). As with the first simulation study, detailed performance results can be found in Table 1. Under this DGP, for this estimator configuration, and with ADAPT-R's sample size of 1,815, we saw 95.90% confidence interval coverage and were powered at 87.20% for detecting effect modification by initial effects.

7 | APPLICATION TO ADAPT-R

In the context of ADAPT-R, our scientific question is the following causal question: does the initial, individualized effectiveness of initial CCT modify its discontinuation effects? Is it the case that incentive discontinuation (i.e., stopping CCT administration) harms those for whom the probability of initial benefit is small versus large? In this section, we apply the methods described in this article to ADAPT-R data to answer this scientific question.

7.1 | ADAPT-R data

We refer the reader to Geng et al., 2023[8] and Montoya et al., 2023[30] for details on the ADAPT-R study. Briefly, adults living with HIV in Kisumu, Kenya were initially randomized with equal probability to interventions intended to prevent lapses in HIV care (short message service [SMS] messages, conditional cash transfers [CCTs], or standard-of-care [SOC] counseling). If patients had a lapse in HIV care, they were re-randomized with equal probability to a more intensive intervention (SMS and CCTs, peer navigator, or SOC outreach), intended to re-engage them back into care. If patients succeeded in their initial care and were initially randomized to an active arm, they were re-randomized (with equal probability) to either continue or discontinue the initial intervention. Those who succeeded and were initially given SOC remained in SOC.

Letting $t = 1$ be the time of first randomization and $t = 2$ be the time of second randomization (either date of first retention lapse or one year after initial randomization, whichever occurs first), the observed ADAPT-R data consisted of the following:

- Baseline covariates $L(0)$ (the full list of baseline covariates are in Appendix A);
- Stage 1 prevention intervention $A(1)$, which consisted of either receiving CCTs (coded as $A(1) = 1$), or not receiving CCTs (i.e., receiving SOC or SMS; coded as $A(1) = 0$);
- Time-varying covariates assessed between randomization to Stage 1 and Stage 2 interventions, $L(1) = (Y(1), W(1))$, which included: 1) $Y(1)$, an indicator of whether there was a lapse in care ($\geq$ 14 days late to a clinic visit) within the first year after enrollment, where $Y(1) = 1$ is no lapse in care. $Y(1)$ in this case is also the first-stage outcome; 2) $W(1)$, other time-varying covariates unrelated to the randomization scheme, such as time from first randomization to second randomization;
- Stage 2 retention intervention $A(2)$, which consisted of either a CCT-based intervention (coded as $A(2) = 1$) or a non-CCT based intervention (coded as $A(2) = 0$);
- The outcome of interest $L(2) = Y(2)$, an indicator of no lapse in care (i.e., no visits made $\geq$ 14 days late) within the year following re-randomization to second-stage treatment, where $Y(2) = 1$ is no lapse in care.

A schematic of the study and how the data were generated is presented in Figure 1. The SCM relating these variables is covered by the general SCM above, for $t = 2$.

7.2 | Causal parameters in the context of ADAPT-R

7.2.1 | Conditional average treatment effects of preventative CCTs: We can intervene on the ADAPT-R's SCM to define a model on the CATE. Letting $r = 1$, we define counterfactuals that quantify the initial effects of giving initial CCTs on year 1 outcomes, where if $a(1) = 1$, this would set $A(1)$ equal to everyone receiving CCTs; conversely, setting $A(1)$ to $a(1) = 0$ means everyone receives an intervention that is not CCTs. In this first stage,

$Z(1) = \varnothing$, because patients were marginally randomized to their prevention intervention. In this way, we can define the CATE relative to ADAPT-R outcomes at time 1 as follows: $B_{P_{U,X}}(L(0)) \equiv \mathbb{E}_{P_{U,X}}[Y(1)_{a(1)=1} - Y(1)_{a(1)=0} \mid L(0)]$. In words, this is the average treatment effect of receiving CCTs $A(1)$ on HIV care adherence $Y(1)$, conditional on patients' baseline covariates $L(0)$.

Using our knowledge of how data were collected in the ADAPT-R study, we assume the observed data $O_i \equiv (L(0)_i, A(1)_i, L(1)_i, A(2)_i, L(2)_i) \sim P_0 \in \mathcal{M},\ i = 1, \ldots, n$ were generated by sampling 1,815 independent and identically distributed copies from a data-generating system contained in ADAPT-R's causal model, above.

Further, in ADAPT-R, individuals were marginally randomized to $A(1)$, implying $Y(1)_{a(1)} \perp A(1)$ and thus $Y(1)_{a(1)} \perp A(1) \mid L(0)$ hold. Further, the positivity conditions hold, as well, since $g_{A(1),0}(A(1) = 1 \mid L(0)) = 1/3$ and $g_{A(1),0}(A(1) = 0 \mid L(0)) = 2/3$. Using this, we can identify the initial CATE as a function of ADAPT-R's observed distribution, i.e., $B_{P_{U,X}}(L(0)) = B_0(L(0)) \equiv \mathbb{E}_0\big[Y(1) \mid A(1) = 1, L(0)\big] - \mathbb{E}_0\big[Y(1) \mid A(1) = 0, L(0)\big]$. An estimator of the effect of receiving initial CCTs within subgroups of $L(0)$ based on ADAPT-R's empirical distribution P_n is $B_n(L(0))$, an estimate of the true first-stage blip function $B_0(L(0))$.

7.2.2 | The target causal parameter: application to ADAPT-R: In the context of ADAPT-R, we are interested in understanding the differential effects of discontinuing CCTs by estimated conditional effects of initiating CCTs. Exploiting ADAPT-R's sequential randomization scheme, we could additionally intervene on the model such that everyone receives CCTs, i.e., set $A(1) = 1$ for all in the structural equations. In this way, we could examine effects in a world where everyone would have received initial CCTs, thereby including all participants in the evaluation of this parameter. The causal estimand is thus $m^*(a(2),\ B_n(L(0))) = \mathbb{E}_{P_{U,X}}\big[Y(2)_{a(1)=1,\,a(2)} \mid B_n(L(0))\big]$, interpreted as the effect of discontinuing CCTs after receiving them, within strata of estimated initial effects.

The true causal function $m^*(a(2),\ B_n(L(0)))$ corresponds to the counterfactual probabilities of treatment adherence under an intervention on (dis)continuation of CCTs that would have been observed for the population for each initial, individualized effect of receiving CCTs. We specify the following working MSM to summarize how these counterfactual probabilities under continuation and discontinuation of CCTs vary as a function of initial, baseline covariate-specific CCT effects:

$$\text{logit}\, m_{\boldsymbol{\beta}}(a(2),\ B_n(L(0))) = \beta_0 + \beta_1 a(2) + \beta_2 B_n(L(0)) + \beta_3 a(2) B_n(L(0))\,. \tag{7}$$

Letting the weights $h(a(2),\ B_n(L(0))) = 1$, we can derive our target causal parameter from the projection of $m^*(a(2),\ B_n(L(0)))$ onto $m_{\boldsymbol{\beta}}(a(2),\ B_n(L(0)))$ according to:

$$\boldsymbol{\beta}_{P_{U,X}} = \arg\min_{\boldsymbol{\beta}} - \sum_{a(2)\in\{0,1\}} \mathbb{E}_{P_{U,X}}\Big[Y(2)_{a(1)=1,a(2)} \log m_{\boldsymbol{\beta}}(a(2),\ B_n(L(0))) + \big(1 - Y(2)_{a(1)=1,a(2)}\big) \log\big(1 - m_{\boldsymbol{\beta}}(a(2),\ B_n(L(0)))\big)\Big].$$

Our causal parameter of interest is $\Psi^F(P_{U,X}) = \beta_3 \in \boldsymbol{\beta}_{P_{U,X}}$, which, in this case, quantifies how the effect of discontinuing CCTs on treatment adherence is modified by initial, individualized responses to initiating CCTs estimated on the ADAPT-R trial data. Our null hypothesis is that $\Psi^F(P_{U,X}) = 0$, i.e., there are no differential effects of discontinuing CCTs by on individualized effects of initiating CCTs; the alternative hypothesis is $\Psi^F(P_{U,X}) \neq 0$.

7.3 | Identifying the target causal parameter in ADAPT-R

With the ADAPT-R design, it is known that randomization to Stage 2 treatment, conditional on $A(1) = 1$, in ADAPT-R was based on whether or not a person had a lapse in year 1, implying $Y(2)_{a(2)} \perp A(2) \mid A(1) = 1, Y(1)$ and thus $Y(2)_{a(2)} \perp A(2) \mid \bar{L}(1), A(1) = 1$ hold. Additionally, it is known that the following probabilities are positive; thus, positivity holds:

- $g_{A(2),0}(A(2) = 1 \mid L(0), A(1) = 1, W(1), Y(1) = 0) = 1/2$
- $g_{A(2),0}(A(2) = 1 \mid L(0), A(1) = 1, W(1), Y(1) = 1) = 1/3$
- $g_{A(2),0}(A(2) = 0 \mid L(0), A(1) = 1, W(1), Y(1) = 1) = 2/3$
- $g_{A(2),0}(A(2) = 0 \mid L(0), A(1) = 1, W(1), Y(1) = 0) = 1/2$

We note that the probabilities $g_{A(2),0}(A(2) = 1 \mid L(0), A(1) = 0, W(1), Y(1) = 1)$ and $g_{A(2),0}(A(2) = 1 \mid L(0), A(1) = 0, W(1), Y(1) = 0$ are not examined, as they are neither viable treatment paths nor do they correspond to our causal parameter or scientific question of interest (i.e., we are not interested in the probabilities of continuing or discontinuing a CCT-based treatment given that patients initially did not receive a CCT-based treatment).

With this, we can identify the causal parameter examining the effect of discontinuing CCTs – had everyone started them – on later treatment adherence outcomes, conditional on the individualized effect of starting CCTs on earlier treatment adherence outcomes. This is because patients were sequentially randomized based on measured information $\bar{Z}(2)$, so the sequential randomization (i.e., $Y(2)_{\bar{a}(2)} \perp A(2) \mid \bar{L}(1), A(1) = 1$ and $Y(1)_{\bar{a}(1)} \perp A(1) \mid L(0)$)) and positivity (i.e., $g_{A(1),0}(A(1) = a(1) \mid L(0)) > 0 - a.e.$ and $g_{A(2),0}\big(A(2) = a(2) \mid \bar{L}(1), A(1) = 1\big) > 0 - a.e.$) assumptions hold, for $a(1) \in \{0,\ 1\}$ and $a(2) \in \{0,\ 1\}$. We identify this longitudinal causal parameter as the following statistical parameter[31]:

$$\mathbb{E}_{P_{U,X}}[Y(2)_{a(1)=1,a(2)} \mid B_n(L(0))] = \mathbb{E}_0[\mathbb{E}_0[\mathbb{E}_0[Y(2) \mid L(0), A(1) = 1, L(1), A(2) = a(2)] \mid L(0), A(1) = 1] \mid B_n(L(0))].$$

Because $Y(2)$ is binary, then:

$$\beta_0 = \arg\min_{\beta} - \sum_{a(2) \in \{0,1\}} \mathbb{E}_0 \Big[\mathbb{E}_0[\mathbb{E}_0[Y(2) \mid L(0), A(1) = 1, L(1), A(2) = a(2)] \mid L(0), A(1) = 1] \times \log m_\beta(a(2),\ B_n(L(0))) + (1 - \mathbb{E}_0[\mathbb{E}_0[Y(2) \mid L(0), A(1) = 1, L(1), A(2) = a(2)] \mid L(0), A(1) = 1]) \times \log \left(1 - m_\beta(a(2),\ B_n(L(0)))\right) \Big],$$

which solves the equation:

$$0 = \mathbb{E}_0 \left[\sum_{a(2) \in \{0,1\}} \frac{\frac{d}{d\beta_0} m_{\beta_0}(a(2),\ B_n(L(0)))}{m_{\beta_0}(a(2),\ B_n(L(0)))(1 - m_{\beta_0}(a(2),\ B_n(L(0))))} \times \left(\mathbb{E}_0[\mathbb{E}_0[Y(2) \mid L(0), A(1) = 1, L(1), A(2) = a(2)] \mid L(0), A(1) = 1] - m_{\beta_0}(a(2),\ B_n(L(0))) \right) \right].$$

The target statistical parameter of interest is identified as $\Psi(P_0) = \beta_3$ in β_0, which is what we wish to estimate with the ADAPT-R data.

7.4 | ADAPT-R results

7.4.1 | Overall effects: Average effects of 1) starting CCTs and 2) stopping CCTs had everyone started were estimated using TMLE[27], where nuisance parameters were estimated using parametric models (g factors were correctly specified and Q factors were possibly mis-specified, where Q factors were adjusted for sex and age for precision[32,30]). The blip function was estimated as in Section 5.1 and 6.2 – single-stage Q-learning was implemented, where the Q factors were estimated using a SuperLearner that included all baseline covariates.

Among the ADAPT-R sample (i.e., a sample of adults living with HIV in rural Kenya), it was shown that the probability of having no lapse in care during year 1 of the study was higher under a CCT-based treatment versus a non-CCT-based treatment (risk difference [RD]: 7.48%, 95% confidence interval [CI]: 3.66%, 11.30%). In addition, there was variability in the distribution of the average effect of receiving CCTs on remaining in care during year 1, conditional on covariates, and most of these estimates were positive (Figure E1; range of estimated blip values: [−0.01, 0.17]). At the same time, the probability of remaining in care was significantly higher had everyone received CCTs initially and then continued versus discontinued receiving a CCT-based intervention (RD: 22.96%, 95% CI: 16.37%, 29.55%). In other words, discontinuing CCTs was harmful, on average, after initiating them. Further, the increase in benefits was significantly lower than the increase in harm (estimated difference in benefits minus harms: −15.48 %, 95% CI: −22.82%, −8.13%), suggesting that benefits of initiating CCTs do not outweigh the harms, overall.

7.4.2 | Effects among the affected: Based on our results on average effects, we aimed to understand the differential effects of discontinuing CCTs by the effects of initiating CCTs – in other words, CCT discontinuation effects among the (un)affected by initating CCTs. Specifically, we are interested in the effects on HIV care adherence. To do this,

using the methods introduced above, we tested the null hypothesis that the coefficient on the interaction term between the estimated CATE for the initial treatment and the effect of the second stage treatment β_3 on the MSM described in Equation 7 was 0. Parameters of the working MSM were estimated using the TMLE described in Section 5.2, where nuisance parameters were estimated using parametric models: g factors were estimated using correctly specified parametric models and Q factors were estimated using possibly mis-specified parametric models, adjusting for baseline (age and sex) and time-varying covariates (time from randomization to re-reandomization) for precision[32,30].

We estimated our target parameter $\Psi(P_0)$ to be $\hat{\Psi}(P_n) = \beta^*_{3,n} = 10.71$ (95 % CI: −0.74, 22.16; p-value: 0.0668), implying effect modification on the multiplicative scale at the $\alpha = 0.1$ level. If one accepts this result as a true signal (using the acceptable Type 1 error rate of $\alpha = 0.1$ often employed in evaluation of interactions[33]), this indicates that the effect of discontinuing CCTs on treatment adherence is modified by initial, individualized responses to initiating CCTs estimated on the ADAPT-R trial data. This suggests that the effect of CCT discontinuation on HIV care adherence differs by initial responses to CCT assignment. In particular, Figure 2 illustrates that the second stage benefits of continuing CCTs were larger for those with larger initial CCT benefits compared to negligible initial CCT benefits.

8 | DISCUSSION

The purpose of this paper was to present a causal parameter that examines how downstream effects may be modified by initial effects of particular interventions. Specifically, we proposed a working MSM summarizing how the effects of a later stage treatment are modified by earlier, estimated individualized effects; our target causal parameter was the interaction coefficient on this working MSM. In addition, we presented a TMLE for estimating the corresponding target statistical parameter, in addition to the efficient influence curve of this parameter, which allowed us to construct valid inference, including 95% confidence intervals, as shown numerically via simulations. To the best of our knowledge, this causal estimand has not yet been presented or described in the causal inference or precision medicine literature.

Importantly, this method was applied to the ADAPT-R study, a SMART intended to identify strategies for helping people living with HIV remain in their HIV care. In the primary analysis for ADAPT-R, it was shown that CCTs were an effective strategy for preventing lapses in HIV care and, overall, maintenance of CCTs was better than their discontinuation[8]. However, from these primary results it was not clear whether this maintenance was uniform across all participants who initially received CCTs, and further, if any (dis)continuation effects of CCTs may have been moderated by initial responses to CCT treatment. Using the methods we developed, we showed that continuation of CCTs is most effective among those most likely to benefit from CCTs initially. This is consistent with the notion that time-limited benefits were present for a subset of participants while they were receiving CCTs, but these benefits were reversed for those participants once CCT administration stopped.

In general, the proposed methods are useful for gaining a more nuanced understanding of the effects of time-limited interventions, which could begin to inform decisions about how to deploy them. For example, knowing that CCT discontinuation is only worse among those for whom CCTs initially most helps could suggest one should continue CCTs among those for whom it initially most helps. And, among those for whom CCTs does not help initially, it does not make a difference if CCTs are maintained and thus resources can be saved by not administering CCTs to the segment of the population for whom CCTs did not make much of an initial impact. Moving towards the best policy recommendation, however, could require having more decision support than just the first-stage blip (as in the MSMs presented in this paper). Thus, future work could extend these methods by using the initial response to treatment as input to a personalized recommendation of who should continue treatment. Concretely, this would mean using the first-stage CATE/blip of initiating the intervention as an additional tailoring variable "covariate" (more accurately, a dimension reduction/summary of baseline covariates) to the input of a second-stage optimal dynamic treatment rule for discontinuing the intervention.

Additionally, other future work may consider examining how one can move toward the ideal causal parameter – specifically, from causal estimand 4 (the target causal parameter of interest in this manuscript) to 3. The difference between the two estimands is that causal estimand 4 examines effects conditional on B_n, an estimate of the identified CATE, whereas causal estimand 3 examines effects conditional on B_0, the true identified CATE. The former is a data-adaptive parameter, which has practical advantages but may be more challenging to interpret and generalize to a target population, whereas the latter is not. Future work could examine the statistical conditions necessary to estimate causal parameter 3, which will likely rely on certainty that B_n converges to B_0 at an appropriate rate.

Finally, as with all causal inference analyses, it is important to ensure that the target causal parameter one is interested in estimating is answering a relevant and feasible causal question of interest. For example, the relevancy of a target causal parameter may be in question when defining the notion of "harm" for the corresponding research question. In our setting, a treatment was defined to be "harmful" if it was worse than the other treatment option, because of our interest in detecting when non-initiation and discontinuation were worse than initiation and continuation, respectively. As such, we defined a treatment as harmful if the expected value of the relevant outcome under that treatment was less than the expected value of the outcome under the comparator treatment – if higher outcomes are better. However, there may be alternative definitions to "harm." For example, a treatment could be defined as harmful if the expected value of the outcome under such a treatment is greater than or less than an *a priori* defined threshold for the expected outcome (e.g., if the probability of HIV care adherence [or an outcome that captures side effects] under discontinuation of CCTs is less than $k\%$). Alternatively, "harm" may be present when the expected value of a later outcome under a later treatment option is less than the expected value of an earlier outcome under an earlier treatment option – again, if higher outcomes are better (e.g., the probability of HIV care adherence at year 2 under discontinuation of CCTs is less than the probability of HIV care adherence at year 1 under initiation of CCTs). Second, the feasibility of the target causal parameter should be examined. In this case, it is imperative

that the longitudinal (e.g., SMART) data in hand are generated in such a way that it can lend itself to answering the kinds of questions asked in this paper. For example, it may not be feasible to ask about differential effects of a treatment at a later stage by its earlier effects if the same treatment options do not appear in both stages of a SMART. We recommend interest in these kinds of causal questions be considered at the design stage of a study.

ACKNOWLEDGMENTS

The authors would like to thank Drs. Daniel Almirall and David Benkeser for helpful discussions related to this work. Research reported in this publication was supported by NINR award R01NR018801 and NIMH award K99MH133985/R00MH133985. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

Funding Information

This research was supported by the Research reported in this publication was supported by NIH awards R01AI074345 and R00MH133985.

APPENDIX

A: ADAPT-R BASELINE COVARIATES

The following is the full list of ADAPT-R's baseline covariates $L(0)$: participant sex, age, CD4 count, depression, social support, food insecurity, treatment supporter, distance to clinic, wealth index, total work hours, total income, total likely discretionary spending, total fixed income, casual/wage work, work on house farm, self-employment, body mass index, pregnancy status, education level, occupation, marital status, and WHO clinical staging for HIV.

B: PARAMETER FOR BINARY TREATMENT

For a binary treatment at time s, one could instead define the effect among affected parameter as follow: $\mathbb{E}_{P_{U,X}}[Y(s)_{a(s)=1} - Y(s)_{a(s)=0} \mid B_n(H(r))]$. Thus, the following lower-dimensional linear working model could be used:

$$\text{logit } m_{\boldsymbol{\beta}}(B_n(H(r))) = \beta_0 + \beta_1 B_n(H(r))$$

which defines the vector (letting stabilization weights = 1):

$$\boldsymbol{\beta}_{P_{U,X}} = \arg\min_{\boldsymbol{\beta}} \mathbb{E}_{P_{U,X}}\left[\left([Y(s)_{a(s)=1} - Y(s)_{a(s)=0}] - m_{\boldsymbol{\beta}}(B_n(H(r)))\right)^2\right].$$

Then, $\boldsymbol{\beta}_{P_{U,X}}$ is the solution to the equation:

$$0 = \mathbb{E}_{P_{U,X}}\left\{\frac{d}{d\boldsymbol{\beta}_{P_{U,X}}} m_{\boldsymbol{\beta}_{P_{U,X}}}(B_n(H(r)))\left[(Y(s)_{a(s)=1} - Y(s)_{a(s)=0}) - m_{\boldsymbol{\beta}_{P_{U,X}}}(B_n(H(r)))\right]\right\},$$

which can be written as:

$$0 = \mathbb{E}_{P_{U,X}}\left\{\frac{d}{d\boldsymbol{\beta}_{P_{U,X}}} m_{\boldsymbol{\beta}_{P_{U,X}}}(B_n(H(r)))\left[\mathbb{E}_{P_{U,X}}[Y(s)_{a(s)=1} - Y(s)_{a(s)=0} \mid H(s)] - m_{\boldsymbol{\beta}_{P_{U,X}}}(B_n(H(r)))\right]\right\},$$

Our causal parameter of interest is then $\Psi^F(P_{U,X}) = \beta_1 \in \boldsymbol{\beta}_{P_{U,X}}$ from this working MSM.

Under the point-treatment randomization $Y(s)_{a(s)} \perp A(s) \mid H(s)$ and positivity conditions $g_{A(s),0}(A(s) = a(s) \mid H(s)) > 0 - \ a.e.$ hold for all $a(s) \in \{0,\ 1\}$ discussed in the main text:

$$\begin{aligned}
&\mathbb{E}_{P_{U,X}}\left[Y(s)_{a(s)=1} - Y(s)_{a(s)=0} \mid H(s)\right] \\
&= \mathbb{E}_{P_{U,X}}\left[Y(s)_{a(s)=1} \mid H(s)\right] - \mathbb{E}_{P_{U,X}}\left[Y(s)_{a(s)=0} \mid H(s)\right] \\
&= \mathbb{E}_0[Y(s) \mid A(s)=1, H(s)] - \mathbb{E}_0[Y(s) \mid A(s)=0, H(s)].
\end{aligned}$$

Then, let

$$\tilde{D}(O) = \frac{2A(s) - 1}{g_{A(s)}(A(s) \mid \bar{L}(s-1), \bar{A}(s-1))}[Y(2) - Q^S(A(s), H(s))] + Q^S(1, H(s)) - Q^S(0, H(s))$$

where $Q^s(A(s), H(s)) \equiv \mathbb{E}\left[Y(s) \mid A(s), H(s)\right]$. Then $\mathbb{E}_0\left[\tilde{D}(O) \mid H(s)\right] = \mathbb{E}_0\left[Y(s) \mid A(s)=1, H(s)\right] - \mathbb{E}_0\left[Y(s) \mid A(s)=0, H(s)\right]$ if $Q^s = Q_0^s$ or $g_{A(s)} = g_{A(s),0}$.

We can identify the true effect of $A(s)$ on $Y(s)$ conditional on the estimated blip at time r:

$$\begin{aligned}
\mathbb{E}_{P_{U,X}}[Y(s)_{a(s)=1} - Y(s)_{a(s)=0} \mid B_n(H(r))] &= \mathbb{E}_{P_{U,X}}\left[\mathbb{E}_{P_{U,X}}\left[Y(s)_{a(s)=1} - Y(s)_{a(s)=0} \mid H(s)\right] \mid B_n(H(r))\right] \\
&= \mathbb{E}_0\left[\mathbb{E}_0[\tilde{D}(O) \mid H(s)] \mid B_n(H(r))\right]
\end{aligned}$$

Again, supposing that the causal parameter of interest is an element of a vector of coefficients in a linear working MSM, then:

$$\boldsymbol{\beta}_0 = \arg\min_{\boldsymbol{\beta}} \mathbb{E}_0\left[\left(\mathbb{E}_0[\tilde{D}(O) \mid H(s)] - m_{\boldsymbol{\beta}}(B_n(H(r)))\right)^2\right].$$

where $\boldsymbol{\beta}_0$ solves the equation:

$$0 = \mathbb{E}_0\left\{\frac{d}{d\boldsymbol{\beta}_0} m_{\boldsymbol{\beta}_0}(B_n(H(r)))\left[\mathbb{E}_0[\tilde{D}(O) \mid H(s)] - m_{\boldsymbol{\beta}_0}(B_n(H(r)))\right]\right\},$$

Using our observed data, our goal would then be to estimate $\beta_1 = \Psi(P_0)$, an element of $\boldsymbol{\beta}_0$. We note that this method is similar to the best linear predictor of the CATE concept in[34].

C: TMLE FOR STATISTICAL PARAMETER

We describe the procedure for implementing a TMLE for $\Psi(P_0)$, the target statistical parameter of interest, and constructing confidence intervals around the resulting estimate, using procedures previously described by[16] and[17]. We illustrate this under the scenario where $\mathscr{A}_s = \{0,1\}$, although this could be generalized to second-stage treatments with more than two treatment options:

1. First, generate an estimator of $g_{A(s),0}$ called $g_{A(s),n}$ (e.g., a maximum likelihood estimator for g_0 according to a correctly specified parametric model for g_0, such as a logistic regression of $A(s)$ on $H(s)$). The estimated predicted probability of being treated and untreated at time s is then $g_{A(s),n}(1 \mid H(s))$ and $g_{A(s),n}(0 \mid H(s))$, respectively.
2. Generate an initial estimate of $Q_0^s(A(s), H(s)) \equiv \mathbb{E}_0[Y(s) \mid A(s), H(s)]$ by regressing $Y(s)$ on $A(s)$ and $H(s)$; call this estimator Q_n^s. Using this regression fit, predict at each $A(s) = 1$ and $A(s) = 0$ value to generate $Q_n^s(1, H(s))$ and $Q_n^s(0, H(s))$.
3. Among those who received treatment $A(s) = 1$, create a $n_1 \times 6$ matrix with the columns having the following vectors of length n_1 (where n_1 is the number of people who received $A(s) = 1$). Call this matrix M_1:
   - **a.** $Y(s)$
   - **b.** $A(s)$ set to 1 for all, i.e., the constant 1 repeated n_1 times
   - **c.** $B_n(H(r))$
   - **d.** $A(s) \times B_n(H(r))$, which in this case is the same as $B_n(H(r))$ among those with $A(s) = 1$
   - **e.** $\text{logit}(Q_n^s(1, H(s)))$
   - **f.** Weight $1/g_{A(s),n}(1 \mid H(s))$
4. Among those who did not receive treatment $A(s) = 0$, create a $n_0 \times 6$ matrix with the columns having the following vectors of length n_0 (where n_0 is the number of people who received $A(s) = 0$). Call this matrix M_0:
   - **a.** $Y(s)$
   - **b.** $A(s)$ set to 0 for all, i.e., the constant 0 repeated n_0 times
   - **c.** $B_n(H(r))$
   - **d.** $A(s) \times B_n(H(r))$, which in this case is 0 repeated n_0 times
   - **e.** $\text{logit}(Q_n^s(0, H(s)))$
   - **f.** Weight $1/g_{A(s),n}(0 \mid H(s))$

5. Stack M_1 and M_0 together to create the $n \times 6$ matrix M.

6. Using the data from the matrix M, fit a pooled logistic regression of $Y(s)$ on $A(s)$, $B_n(H(r))$, and $A(s) \times B_n(H(r))$ with offset $\text{logit}(Q_n^s(A(s), H(s)))$ and weights $1/g_{A(s),n}(A(s) \mid H(s))$. This gives a fit $\varepsilon_n = (\varepsilon_{0,n}, \varepsilon_{1,n}, \varepsilon_{2,n}, \varepsilon_{3,n})$.

7. Generate $Q_n^{*,s}(1, H(s))$ and $Q_n^{*,s}(0, H(s))$ by using the fit from the previous step and predicting at $A(s) = 1$ and $A(s) = 0$, respectively. That is:

$$Q_n^{*,s}(1, H(s)) = \text{logit}^{-1}\left(\text{logit}(Q_n^s(1, H(s))) + \varepsilon_{0,n} + \varepsilon_{1,n} + \varepsilon_{2,n} B_n(H(r)) + \varepsilon_{3,n} B_n(H(r))\right)$$

and

$$Q_n^{*,s}(0, H(s)) = \text{logit}^{-1}\left(\text{logit}(Q_n^s(0, H(s))) + \varepsilon_{0,n} + \varepsilon_{2,n} B_n(H(r))\right)$$

8. Stack $Q_n^{*,s}(1, H(s))$ and $Q_n^{*,s}(0, H(s))$ to create a single vector $\bar{Q}_n^{*,s}$ of length $2 \times n$. Fit a pooled logistic regression of $\bar{Q}_n^{*,s}$ on $a(s) \in \{0,1\}$ and $B_n(H(r))$ according to the model in Equation 6 in the main text. This gives a TMLE of β_0, called β_n^*. The β_3 coefficient in β_n^*, which we will call $\hat{\Psi}(P_n) = \beta_{3,n}^*$, gives a TMLE of the target parameter $\Psi(P_0)$.

For a working MSM according to Equation 6 and with $\mathscr{A}_s = \{0, 1\}$ (noting, again, that this could be generalized to second-stage treatments with more than two treatment options), the efficient influence function for β at a distribution P is $D^*(P) \equiv C^{-1} D(P)$, where C is a normalizing matrix, $C = -\mathbb{E}\left[\frac{d}{d\beta} D(P)\right]$, and

$$\begin{aligned} D(P) = & \frac{h(a(s), B_n(H(r)))\left(Y(s) - Q^S(A(s), H(s))\right)}{g(A(s) \mid H(s))}(1, A(s), B_n(H(r)), A(s)B_n(H(r)))^T \\ & + \sum_{a(s) \in \mathscr{A}_S} h(a(s), B_n(H(r)))\left(Q^S(a(s), H(s))\right. \\ & \left. - m_\beta(a(s), B_n(H(r)))\right)(1, a(s), B_n(H(r)), a(s)B_n(H(r)))^T. \end{aligned}$$

D: SIMULATION SPECIFICATIONS

This section provides the exact data generating processes (DGPs) for simulation 1 and 2.

D.1 Simulation 1

In Simulation 1, we simulated data according to the following three DGPs:

$$L(0) = \begin{cases} L^1(0) \sim \ Bernoulli(p = 0.5), \\ L^2(0) \sim \ Bernoulli(p = 0.5), \end{cases}$$

$$A(1) \sim Bernoulli(p = 0.5),$$

$$Y(1) \sim Bernoulli\left(p = Q_0^1\left(L^1(0), L^2(0), A(1)\right)\right)$$

$$A(2) \sim Bernoulli(p = 0.5),$$

$$Y^q(2) \sim \text{Bernoulli}\ (p = Q_0^{2,q}(L^1(0), L^2(0), A(1), Y(1), A(2))\,.$$

Here, baseline covariates measured on all subjects $L(0)$ comprises $L^1(0)$ and $L^2(0)$. Additionally, $Q_0^1(L(0), A(1))\big) = \text{logit}^{-1}(L^1(0) + L^2(0) + A(1) + L^1(0)A(1) + 2L^2(0)A(1) - 5A(1)L^1(0)L^2(0))$ so that the true first-stage blip function is:

$$B_0(L(0)) = \text{logit}^{-1}\left(L^1(0) + L^2(0) + 1 + L^1(0) + 2L^2(0) - 5L^1(0)L^2(0)\right) - \text{logit}^{-1}\left(L^1(0) + L^2(0)\right).$$

Here, $Y^q(2)$ is indexed by $q \in \{1,\ 2,\ 3\}$ such that the first process generates the outcome $Y^1(2)$ according to $Q_0^{2,1}(L(0), A(1), Y(1), A(2))) = \text{logit}^{-1}\left(L^1(0)A(2)\right)$, the second generates $Y^2(2)$ according to $Q_0^{2,2}(L(0), A(1), Y(1), A(2))) = 1 - \text{logit}^{-1}\left((1 - A(2))\left(1 - L^1(0)\right)\right)$, and the third generates $Y^3(2)$ according to $Q_0^{2,2}(L(0), A(1), Y(1), A(2))) = \text{logit}^{-1}(A(2) + 1.381)$.

D.2 Simulation 2

This section provides the exact DGP for Simulation 2, the simulation study that mimics the ADAPT-R trial. Specifically, we simulated data according to the following DGP:

$$L^1(0) \sim Normal\left(\mu = 0, \sigma^2 = 1\right)$$

$$L^2(0) \sim Bernoulli(p = 0.5)$$

$$L^3(0) \sim Normal\left(\mu = L^1(0) + L^2(0), \sigma^2 = 1\right)$$

$$A(1) \sim Bernoulli(p = 1/3)$$

$$W(1) \sim Normal\left(\mu = L^1(0) + L^3(0) + A(1), \sigma^2 = 1\right)$$

$$Y(1) \sim Bernoulli\left(p = Q_0^1\left(A(1), L^1(0), L^2(0), L^3(0)\right)\right)$$

$$A(2) \sim \begin{cases} Bernoulli(p = 1/3) & \text{if } A(1) = 1, Y(1) = 1 \\ Bernoulli(p = 1/2) & \text{if } A(1) = 1, Y(1) = 0 \\ Bernoulli(p = 1/3) & \text{if } A(1) = 0, Y(1) = 1 \\ Bernoulli(p = 0) & \text{if } A(1) = 0, Y(1) = 0 \end{cases}$$

$$Y(2) \sim Bernoulli\left(p = Q_0^2\left(L^1(0), L^2(0), L^3(0), A(1), Y(1), A(2)\right)\right),$$

Here, baseline covariates measured on all subjects
$L(0)$ comprises $L^1(0)$, $L^2(0)$, and $L^3(0)$. Additionally,
$Q_0^1\left(A(1), L^1(0), L^2(0), L^2(0)\right) = 1 - (0.5\ \text{logit}^{-1}(1 - L^1(0)^2 + 3L^2(0) + 5L^3(0)^2A(1) - 4.45A(1)),$
$+ 0.5\ \text{logit}^{-1}\left(0.5 + L^3(0) + 2L^1(0)L^2(0) + 3\left|L^2(0)\right|A(1) - 1.5A(1)\right))$
a function drawn and modified from[21], and
$Q_0^2(L^1(0), L^2(0), L^3(0), A(1), Y(1), A(2))) = 1 - \text{logit}^{-1}$.
$(|L^1(0)|^{Y(1)} + Y(1) - 2A(2) + 5A(2)L^2(0) + sin(L^3(0) - 4))$

SuperLearner was used to estimate the Q factors needed to estimate the blip function. The library of candidate algorithms used in the SuperLearner were: 1) `SL.mean`, the mean of the outcome, 2) three logistic regressions, each with main term for the treatment $A(1)$ and each covariate $L^j(0)$, plus an interaction $L^j(0)$ times $A(1)$, for $j \in \{1, \ldots, 3\}$, 3) a single logistic regression with all $L^j(0)$ and $A(1)$ as main terms, plus all two-way interactions between $L^j(0)$ times $A(1)$, for all $j \in \{1, \ldots, 3\}$, 4) `SL.glmnet`[35,36], 5) `SL.bayesglm`[37], 6) `SL.earth`[38], and 6) `SL.randomForest`[39]. In the simulations, the Q factors were estimated using mis-specified parametric regressions in the sense that the SuperLearner library does not include estimators that correctly capture the absolute value, sine function, and quadratic function in Q_0^2. This same SuperLearner library was used for estimating the Q factors for the blip in the the ADAPT-R analysis.

E: HISTOGRAM OF THE STAGE 1 CATE/BLIP

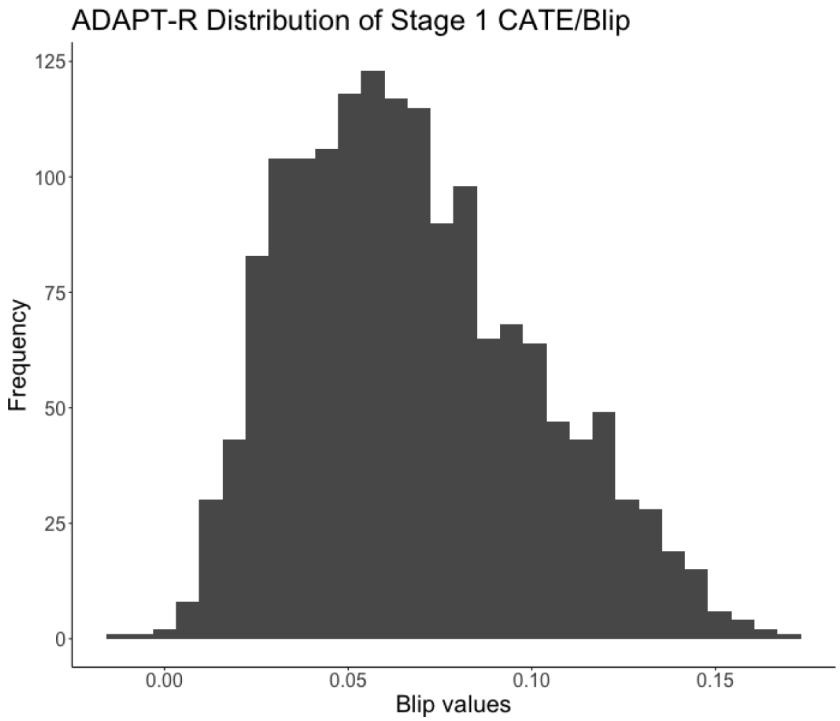

FIGURE E1:
Distribution of conditional average treatment effect (CATE)/blip estimates of initial conditional cash transfers (CCTs) on one-year HIV care retention for the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial.

Abbreviations:

ADAPT-R	Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care
CATE	conditional average treatment effect
CCT	conditional cash transfer
ITE	individualized treatment effect
MSM	marginal structural model
SCM	structural causal model
SMART	sequential multiple assignment randomized trial
SMS	short message service
SOC	standard of care
TMLE	targeted maximum likelihood estimator

REFERENCES

1. Kuntz JL, Dickerson JF, Schneider JL, et al. Factors associated with opioid-tapering success: A mixed methods study. Journal of the American Pharmacists Association. 2021;61(3):248–257. [PubMed: 33485815]
2. Cancelliere C, Cassidy JD, Ammendolia C, Côté P. Are workplace health promotion programs effective at improving presenteeism in workers? A systematic review and best evidence synthesis of the literature. BMC Public Health. 2011;11:1–11. [PubMed: 21199570]
3. Atoloye AT, Savoie-Roskos MR, Guenther PM, Durward CM. Effectiveness of Expanded Food and Nutrition Education Program in changing nutrition-related outcomes among adults with low income: A systematic review. Journal of Nutrition Education and Behavior. 2021;53(8):691–705. [PubMed: 33931296]
4. Kosorok MR, Laber EB. Precision medicine. Annual Review of Statistics and its Application. 2019;6:263–286.
5. Kidwell KM, Almirall D. Sequential, multiple assignment, randomized trial designs. JAMA. 2023;329(4):336–337. [PubMed: 36692577]
6. Kosorok MR, Moodie EE. Adaptive Treatment Strategies in Practice: Planning Trials and Analyzing Data for Personalized Medicine. Society for Industrial and Applied Mathematics, 2015.
7. Murphy SA. An experimental design for the development of adaptive treatment strategies. Statistics in Medicine. 2005;24(10):1455–1481. [PubMed: 15586395]
8. Geng EH, Odeny TA, Montoya LM, et al. Adaptive strategies for retention in care among persons living with HIV. NEJM Evidence. 2023;2(4):EVIDoa2200076.
9. Kennedy EH. Towards optimal doubly robust estimation of heterogeneous causal effects. Electronic Journal of Statistics. 2023;17(2):3008–3049.
10. Semenova V, Chernozhukov V. Debiased machine learning of conditional average treatment effects and other causal functions. The Econometrics Journal. 2021;24(2):264–289.
11. Wager S, Athey S. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. Journal of the American Statistical Association. 2018;113(523):1228–1242. doi: 10.1080/01621459.2017.1319839
12. Robins JM. Marginal structural models. 1997 Proceedings of the American Statistical Association. 1998a:1–10.
13. Neugebauer R, van der Laan M. Nonparametric causal effects based on marginal structural models. Journal of Statistical Planning and Inference. 2007;137(2):419–434.
14. Hubbard AE, Kherad-Pajouh S, van der Laan MJ. Statistical inference for data adaptive target parameters. The International Journal of Biostatistics. 2016;12(1):3–19. [PubMed: 27227715]
15. Petersen ML, van der Laan MJ. Causal models and learning from data: integrating causal modeling and statistical estimation. Epidemiology. 2014;25:418–426. [PubMed: 24713881]
16. Petersen M, Schwab J, Gruber S, Blaser N, Schomaker M, van der Laan M. Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. Journal of Causal Inference. 2014;2(2):147–185. [PubMed: 25909047]
17. Rosenblum M, van der Laan MJ. Targeted maximum likelihood estimation of the parameter of a marginal structural model. The International Journal of Biostatistics. 2010;6(2).
18. Pearl J Causality: Models, Reasoning, and Inference. Cambridge University Press, 2000.
19. Holland PW. Statistics and causal inference. Journal of the American statistical Association. 1986;81(396):945–960.
20. Rosenblum M, van der Laan MJ. Using regression models to analyze randomized trials: asymptotically valid hypothesis tests despite incorrectly specified models. Biometrics. 2009;65(3):937–945. [PubMed: 19210739]
21. Luedtke AR, van der Laan MJ. Super-Learning of an optimal dynamic treatment rule. The International Journal of Biostatistics. 2016;12(1):305–332. [PubMed: 27227726]
22. Montoya LM, van der Laan MJ, Luedtke AR, Skeem JL, Coyle JR, Petersen ML. The optimal dynamic treatment rule superlearner: considerations, performance, and application to criminal

justice interventions. The International Journal of Biostatistics. 2023;19(1):217–238. [PubMed: 35708222]

23. van der Laan MJ, Polley EC, Hubbard AE. Super learner. Statistical Applications in Genetics and Molecular Biology. 2007;6(1):1–21.
24. Robins JM. 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. 1986;7:9–12.
25. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. Epidemiology. 2000:550–560. [PubMed: 10955408]
26. Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. In: American Statistical Association. 1999:6–10.
27. van der Laan MJ, Rose S. Targeted Learning: Causal Inference for Observational and Experimental Data. Springer, 2011.
28. Lendle SD, Schwab J, Petersen ML, van der Laan MJ. ltmle: An R Package Implementing Targeted Minimum Loss-Based Estimation for Longitudinal Data. Journal of Statistical Software. 2017;81(1):1–21. doi: 10.18637/jss.v081.i01
29. R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing; Vienna, Austria: 2024.
30. Montoya LM, Kosorok MR, Geng EH, Schwab J, Odeny TA, Petersen ML. Efficient and robust approaches for analysis of sequential multiple assignment randomized trials: Illustration using the ADAPT-R trial. Biometrics. 2023;79(3):2577–2591. [PubMed: 36493463]
31. Bang H, Robins JM. Doubly robust estimation in missing data and causal inference models. Biometrics. 2005;61(4):962–973. [PubMed: 16401269]
32. Moore KL, van der Laan MJ. Covariate adjustment in randomized trials with binary outcomes: targeted maximum likelihood estimation. Statistics in Medicine. 2009;28(1):39–64. [PubMed: 18985634]
33. Selvin S Statistical analysis of epidemiologic data. 35. Oxford University Press, 2004.
34. Chernozhukov V, Demirer M, Duflo E, Fernandez-Val I. Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in India. tech. rep., National Bureau of Economic Research; 2018.
35. Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. Journal of statistical software. 2010;33(1):1. [PubMed: 20808728]
36. Simon N, Friedman J, Hastie T, Tibshirani R. Regularization paths for Cox's proportional hazards model via coordinate descent. Journal of statistical software. 2011;39(5):1.
37. Gelman A, Jakulin A, Pittau MG, Su YS. A weakly informative default prior distribution for logistic and other regression models. 2008.
38. Friedman JH. Multivariate adaptive regression splines. The Annals of Statistics. 1991;19(1):1–67.
39. Breiman L Random forests. Machine Learning. 2001;45(1):5–32.

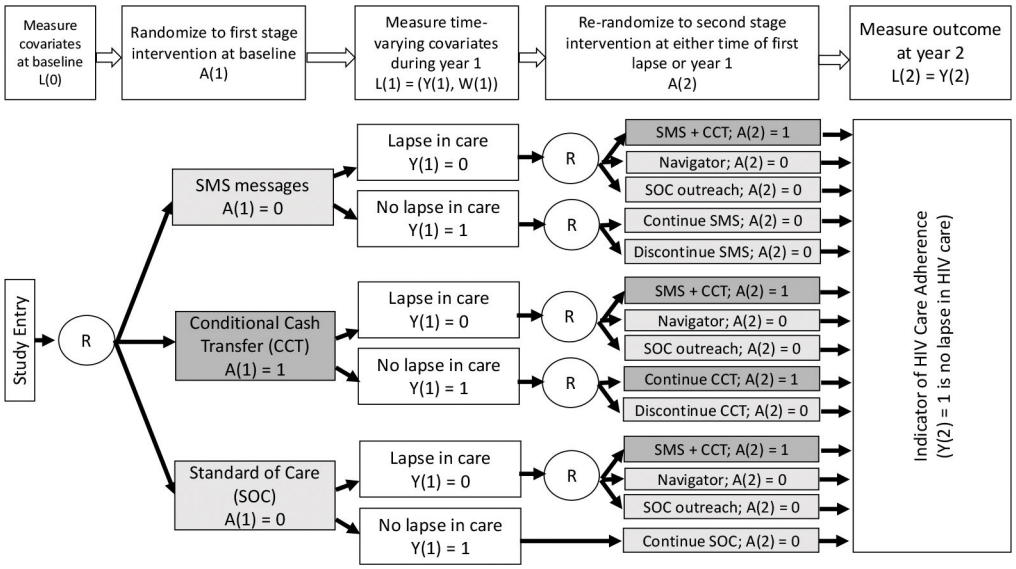

FIGURE 1:

Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial design. The shaded boxes are the treatments to which patients were randomized. Boxes shaded with light gray are non-CCT-based treatments; boxes shaded with dark gray are CCT-based treatments.

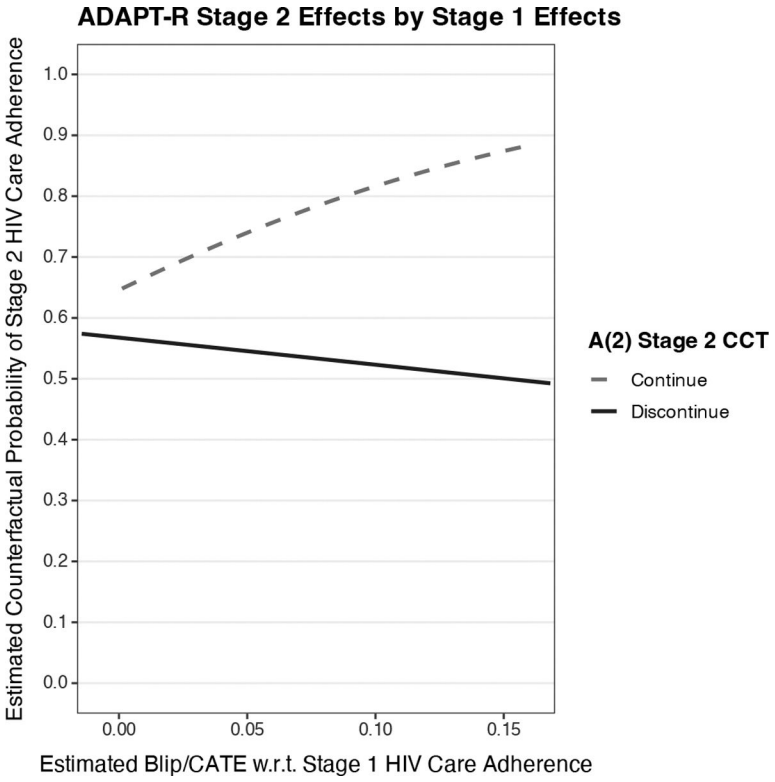

FIGURE 2:

Differential effects of Stage 2 interventions (i.e., effect of continuing versus discontinuing conditional cash transfer [CCT]-based interventions on two-year HIV care adherence) by Stage 1 effects (i.e., conditional average effects [CATE]/blip of initial, preventative CCT on one-year HIV care adherence) for the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial.

TABLE 1:

Performance (bias, variance, mean squared error [MSE], 95% confidence interval [CI] coverage [cov.], and power) for two simulation studies: Simulation 1 and Simulation 2. Simulation 1 illustrates three “simple" sequentially randomized experiments (DGP 1, 2, and 3) in which the differential downstream treatment effects by initial upstream treatment effects differ. Simulation 2 mimics the Adaptive Strategies for Preventing and Treating Lapses of Retention in HIV Care (ADAPT-R) trial.

	Bias	Variance	MSE	95% CI Cov.	Power	Type I Error
Simulation 1, DGP 1	−0.0150	0.4699	0.4702	94.7%	86.3%	-
Simulation 1, DGP 2	−0.0013	0.3694	0.3694	96.1%	85.7%	-
Simulation 1, DGP 3	−0.0155	0.8566	0.8568	95.4%	-	4.6%
Simulation 2	0.1043	4.9055	4.9163	95.9%	87.2%	-