

Individualized Treatment Rules: Generating Candidate Clinical Trials

Maya L. Petersen*

Steven G. Deeks†

Mark J. van der Laan‡

*Division of Biostatistics, School of Public Health, University of California, Berkeley, mayaliv@berkeley.edu

†Department of Medicine, San Francisco General Hospital, University of California, San Francisco

‡Division of Biostatistics, School of Public Health, University of California, Berkeley, laan@berkeley.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/ucbbiostat/paper207>

Copyright ©2006 by the authors.

Individualized Treatment Rules: Generating Candidate Clinical Trials

Maya L. Petersen, Steven G. Deeks, and Mark J. van der Laan

Abstract

Statistical methods have rarely been applied to learn individualized treatment rules, or rules for altering treatments over time in response to changes in individual covariates. Termed dynamic treatment regimes in the statistical literature, such individualized treatment rules are of primary importance in the practice of clinical medicine. History-Adjusted Marginal Structural Models (HA-MSM) estimate individualized treatment rules that assign, at each time point, the first action of the future static treatment plan that optimizes expected outcome given a patient's covariates. However, as we discuss here, the optimality of these rules can depend on the way in which treatment was assigned in the data from which the rules were derived. In this article we discuss the conditions sufficient for treatment rules identified by HA-MSM to be statically optimal, or in other words, to select the optimal future static treatment plan at each time point, regardless of the way in which past treatment was assigned. The resulting treatment rules form appropriate candidates for evaluation using randomized controlled trials. We demonstrate that a history-adjusted individualized treatment rule is statically optimal if it depends on a set of covariates that are sufficient to control for confounding of the effect of past treatment history on outcome. Methods and results are illustrated using an example drawn from the antiretroviral treatment of patients infected with HIV. Specifically, we focus on rules for deciding when to modify the treatment of patients infected with resistant virus.

1. INTRODUCTION

Many pressing clinical questions involve strategies or rules for deciding how treatments should be assigned and changed over time. To be effective, such strategies must be individualized. In other words, treatment must be assigned and modified in response to individual changes in disease progression, side effects, and other patient characteristics. Examples of such individualized treatment rules, known as dynamic treatment regimes in the statistical literature, include the decision to start anti-hypertensive medication in response to repeated measurements of hypertension and the decision to modify the dose of an antidepressant medication in response to adverse effects. Individualized treatment rules can be contrasted with static treatment regimens, in which treatment can change over time, but not in response to individual covariates. An example of this would be the design of most current randomized controlled trials, in which participants are assigned to take a drug of interest for a specified time interval. However, despite the fact that individualized treatment rules, often in the form of clinical guidelines, are used ubiquitously in clinical care, such rules are rarely derived rigorously from data.

Recent work has aimed to learn individualized treatment rules statistically. Structural nested mean models, presented by Murphy and Robins, aim to estimate dynamic regimes that optimize expected patient outcome [1, 2]. However, the programming required to implement these methods is substantial. In addition, this and related approaches estimate rules that optimize expected outcome at a fixed time point. Clinical practice, however, is often aimed at optimizing a moving target (e.g., a patient's outcome 6 months

in the future). Perhaps for these reasons, in practice estimation of optimal dynamic treatment regimes remains rare.

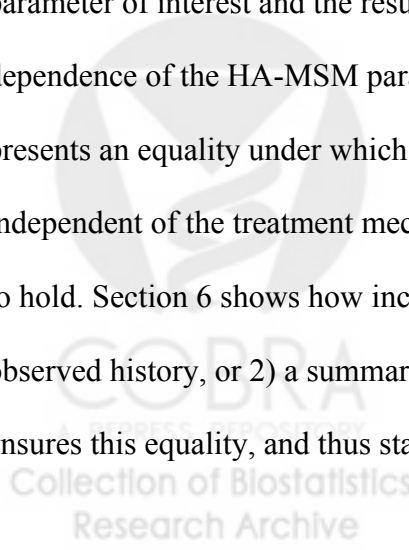
In this article, we present a novel statistical method for using observational data to estimate individualized treatment rules. Our method estimates rules that are not truly optimal among all dynamic rules (as estimated by Murphy, et. al.), but rather rules that are statically optimal [3]. As a result, our method can be implemented using standard software. We argue that static optimality is both an interpretable and interesting property, and point out several advantages of our approach compared to previous approaches aimed at estimating dynamic rules.

Our method employs history-adjusted marginal structural models (HA-MSM) [3-6], a generalization of the marginal structural model (MSM) approach to causal inference [5-8]. HA-MSM estimate individualized treatment rules by identifying, at each time point, the future static treatment plan that optimizes expected outcome given individual covariates of interest. The first action of this static plan is then assigned, and the optimal future static plan recalculated at the subsequent time point. The resulting treatment regimen is a dynamic rule, in that the treatment assigned can change in response to individual covariates [3, 4].

However, as we explain, the treatment rules identified by HA-MSM can depend on the way in which treatment was assigned in the study population (the treatment mechanism). As a result, if these rules were applied at a given time point to an identical population in which treatment up till that time point had been assigned differently, they would not necessarily identify the future treatment plan that would optimize expected

outcome. In this paper, we present conditions sufficient to ensure that the treatment rule estimated by HA-MSM is truly “statically optimal”, or in other words, that the treatment rule identifies at each time point the first action of a static treatment regimen that is optimal *regardless* of treatment assignment up till that time point. Using the directed acyclic graph (DAG) framework [9], we illustrate that specific causal structures are sufficient to ensure static optimality. Finally, we show that when these conditions are not met HA-MSM can still be used to derive statically optimal treatment rules, by incorporating a summary of patient history, such as that provided by the treatment mechanism, into the individualized rule.

Throughout the article, we rely on an example drawn from the treatment of HIV infection: when antiretroviral therapy (ART) fails to completely suppress the virus, what treatment rule should be used to decide when to modify a patient’s regimen? In Section 2, we present some background on this clinical question and the data used to address it. (For a full description of the sample and corresponding HA-MSM models, see Petersen et. al. [4].) Section 3 uses this data example to review HA-MSM, including key assumptions, the parameter of interest and the resulting dynamic treatment regimen. Section 4 discusses the dependence of the HA-MSM parameter on the observed treatment mechanism. Section 5 presents an equality under which HA-MSM treatment rules are statically optimal (i.e. independent of the treatment mechanism) and gives a sufficient condition for the equality to hold. Section 6 shows how including in the HA-MSM model either: 1) the entire observed history, or 2) a summary of the history provided by the treatment mechanism, ensures this equality, and thus static optimality. Section 7 uses Directed Acyclic Graph



theory to demonstrate that the necessary equality also holds under specific causal structures. Thus, in certain contexts, incorporation of the treatment mechanism may not be necessary. Section 8 presents examples of statically optimal treatment rules derived from the data example. The results illustrate the implementation of the treatment mechanism method to ensure static optimality, as well as a causal structure under which inclusion of the treatment mechanism as a covariate is not needed. Finally, we conclude with a discussion in Section 9, including a comparison of statically optimal regimens to the truly optimal dynamic regimes.


2. HIV DATA EXAMPLE: WHEN TO SWITCH ANTIRETROVIRAL THERAPY?

Combination antiretroviral therapy can successfully suppress viral replication in many HIV-infected individuals. As a result, the immune system recovers, CD4 T cell counts increase, and clinical prognosis improves dramatically [10]. Unfortunately, HIV frequently develops resistance to the drugs being used to treat it, allowing the virus to resume replication and resulting in an increase in the amount of virus in the patient's blood (measured as plasma HIV RNA level or viral load). When this loss of virologic suppression occurs, clinicians are faced with an important treatment decision: how to decide when to modify the patient's antiretroviral regimen. Waiting too long to switch to a new regimen may result in the accumulation of additional resistance mutations. Also, increased levels of viral replication will lead over time to progressive loss of CD4+ T cell counts. However, switching regimens too early can deplete future treatment options, and may result in increased toxicity. (For a review of this issue, see Deeks [11].) Importantly,

the decision as to “when to switch” is often made by clinicians based in part on the number of previous treatment regimens (with fewer prior regimens leading to a desire to modify therapy earlier), and the rate at which patients appear to be progressing clinically. Hence, the optimal strategy for deciding when to switch is likely to be based on the evolution of patient and virologic characteristics over the course of non-suppressive therapy.

We used HA-MSM to estimate the effect of time until modifying antiretroviral therapy regimen on CD4 T cell count 8 months in the future. Data were drawn from the Study on the Consequences of the Protease Inhibitor Era (SCOPE) cohort, an observational cohort of HIV-infected individuals followed between 2000 and 2004 in San Francisco, California. Data were collected on socioeconomic status (housing, income, employment), antiretroviral medication use and adherence, occurrence of opportunistic infection and malignancy, recreational drug use, plasma HIV RNA levels, and CD4/CD8 T cell counts. Subjects qualified for the current analysis ($t=0$) if they experienced loss of virologic suppression on an antiretroviral regimen. Loss of virologic suppression (“virologic failure”) was defined using either of the following criteria: 1) at least two detectable viral loads, and no undetectable viral loads in a 4 month period while on a stable antiretroviral regimen; or, 2) at least two undetectable viral loads and no detectable viral loads within the first 6 months of starting a new regimen. One hundred subjects experiencing a total of 116 episodes of virologic failure were identified from the SCOPE cohort and contributed to the analyses described here. For a full description of the sample, see [4].

The treatment of interest (time until modification of the antiretroviral regimen, or “switching”) was defined using a vector of binary covariates $\bar{A}(K) = (A(0), \dots, A(K))$,



where 0 denoted the time virologic failure occurred, and $K+1$ denoted the end of follow up. $A(t)$ remained equal to one as long as a patient remained on his original non-suppressive therapy, jumped to zero as soon as a subject switched therapy, and remained zero thereafter. Longitudinal covariates measured on a subject were defined $\bar{L}(K+1) = (L(0), \dots, L(K+1))$, where $L(t)$ was measured before $A(t)$. The outcome of interest for a given time point j was CD4 T cell count $m=8$ months in the future ($Y(j+m)$). The observed data for a given subject thus consisted of $O = (\bar{A}(K), \bar{L}(K+1))$.

3. REVIEW OF HA-MSM

HA-MSM rely on the counterfactual framework for causal inference, under which each individual has a set of counterfactual covariate processes corresponding to the paths each covariate would have followed under each possible treatment history. This set of counterfactuals is the full data: $X^{Full} = (\bar{L}_a(K+1), a \in A)$, where A denotes the set of possible treatment regimens, and $\bar{L}_a(K+1)$ denotes the counterfactual covariates that would have been observed over the course of follow-up if an individual had followed treatment regimen $\bar{A}(K) = \bar{a}(K)$. Using the counterfactual framework, a set of counterfactual CD4 T cell counts (and other covariates) existed for each individual under each possible switch time.

Under the counterfactual framework, the observed covariate values for an individual are assumed to be equal to the individual's counterfactual covariate values under her observed treatment history (the consistency assumption):

$$O = (\bar{A}(K+1), \bar{L}_{\bar{A}}(K+1)) \quad (1)$$

Under this assumption (1), the distribution of the observed data is indexed by 1) the distribution of the full data (F_x), and 2) the distribution of the treatment history given the full data ($g(\bar{a} | X^{Full}) = g$): $O \sim P_{F_x, g}$, where $g(\bar{a} | X^{Full}) = P(\bar{A} = \bar{a} | X^{Full})$ is the treatment mechanism, which acts as a missingness variable.

HA-MSM further assume the existence of no unmeasured confounders (the Sequential Randomization Assumption or SRA);

$$A(t) \perp\!\!\!\perp X^{Full} | \bar{A}(t-1), \bar{L}(t), \text{ for } t = 0, \dots, K \quad (2)$$

In the HIV application, the SRA assumes that no unmeasured variables affect the decision to switch as well as future CD4 T cell count. Under the SRA, the treatment mechanism can be written as:

$$g(\bar{A} | X^{Full}) = \prod_{t=0}^K P(A(t) | \bar{L}(t), \bar{A}(t-1)) \quad (3)$$

Standard MSM model the expectation (or some other parameter) of the counterfactual outcome, conditional on baseline covariates of interest: $E(Y_a(m) | V)$, where V denotes the values of effect modifiers of interest at baseline ($t=0$). HA-MSM can be understood as fitting a standard MSM at each time point during follow up, which models counterfactual outcomes indexed by treatment after that time point conditional on a subset of observed treatment and covariate history up till that time point. As an alternative to fitting a separate model at each time point, a common model can be fit across time points.

Let $\underline{a}(j)$ denote a future treatment regimen beginning at time point j until the outcome is measured m time points later ($\underline{a}(j) = a(j), \dots, a(j+m-1)$). In general, HA-MSM are

concerned with the counterfactual outcome under the observed treatment history up till time $(j-1)$ and specified a future treatment regimen from time j until the outcome is measured: $Y(j+m)_{\bar{A}(j-1)\underline{a}(j)}$. HA-MSM model the mean (or some other parameter) of these counterfactuals, conditional on $V(j)$, a subset of the observed past up till time j ($V(j) \subset (\bar{L}(j), \bar{A}(j-1))$):

$$E[Y(j+m)_{\bar{A}(j-1)\underline{a}(j)} | V(j)] \quad (4)$$

Typically, $V(j) = (\bar{A}(j-1), \bar{S}(j))$, where $\bar{S}(j) \subset \bar{L}(j)$ denotes the covariates conditioned on, or effect modifiers of interest.

We applied this method among individuals experiencing loss of virologic suppression to estimate the effect of future time until switching therapy on CD4 T cell count 8 months later among individuals who remained on their original therapy. These counterfactual outcomes were denoted $Y(j+m)_{\bar{A}(j-1)c(j)}$, where $c(j)$ denotes the counterfactual future time (after time j) until a subject either switches treatment or the outcome is measured at time $j+m$. Note that $c(j)$ is a summary of the counterfactual treatment regimen beginning at time j (i.e. $c(j)$ is a summary of $\underline{a}(j)$). HA-MSM were used to estimate the following history-adjusted mean:

$$E[Y(j+m)_{\bar{A}(j-1)c(j)} | \bar{A}(j-1), \bar{S}(j)] \quad (5)$$

Models were fit only among those individuals who had not already switched therapy ($\bar{A}(j-1) = 1$), and only among those individuals who had not experienced re-suppression of the virus while on the same therapy ($Sup(j) = 0$). In two sets of analyses, the additional effect modifiers of interest were CD4 count at time j ($CD4(j)$), and the presence of an

opportunistic disease at time j ($OD(j)$). Thus $\bar{S}(j) = (CD4(j), Sup(j))$, and $\bar{S}(j) = (OD(j), Sup(j))$, in turn. In other words, among individuals who had not yet switched therapy or been re-suppressed by time point j , HA-MSM were used to model the mean counterfactual CD4 T cell counts that would have been observed if this entire subpopulation (or a random sample) had switched therapy $c(j)$ months after time j , and how these expected counterfactual outcomes differed depending on time elapsed since loss of suppression occurred (j), and on either a patient's CD4 T cell count at time j ($CD4(j)$) or the presence of opportunistic disease ($OD(j)$) at time j .

The HA-MSM parameter identifies an interesting dynamic treatment regimen, or rule for assigning treatment over time in response to individual changes in covariates. In the HIV analyses, the history-adjusted mean estimates, at each time point j , the effect of additional time until switching therapy on future CD4 T cell count among individuals who have not yet switched, conditional on covariates of interest. E.g., the effect of an additional month until switching is estimated as

$$E(Y(j+m)_{\bar{A}(j-1)c(j)+1} \mid \bar{A}(j-1) = 1, \bar{S}(j)) - E(Y(j+m)_{\bar{A}(j-1)c(j)} \mid \bar{A}(j-1) = 1, \bar{S}(j)).$$

This suggests the following treatment rule: when additional time waiting to switch decreases expected future CD4 T cell count, switch immediately; when additional time waiting increases expected future CD4 T cell count, wait until the next time point to switch, and then reevaluate the estimated effect of additional waiting time. More generally, the treatment rule identified by HA-MSM consists of following, at each time point, the first action of the future static treatment plan that optimizes expected outcome. This

treatment plan is updated at each time point in response to changes in patient covariates.

Thus the HA-MSM dynamic treatment rule can be defined as follows:

$$\underline{a}^*(j | V(j)) \equiv \arg \max_{\underline{a}(j)} E(Y(j+m)_{\bar{A}(j-1)\underline{a}(j)} | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j))$$

$$d(j | V(j)) \equiv \underline{a}^*(j | V(j))(1),$$

where $d(j | V(j))$ is the treatment decision at time j ,

and $\underline{a}^*(j | V(j))(1)$ refers to the first action of $\underline{a}^*(j | V(j))$.

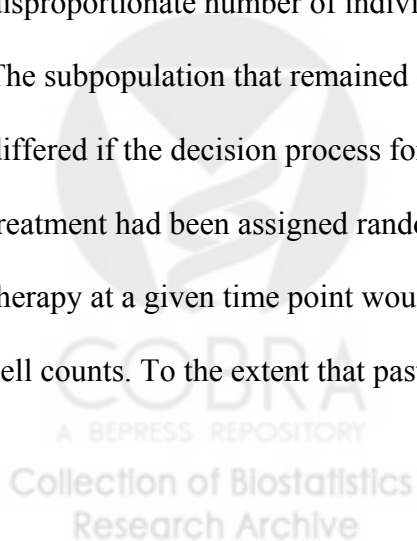
4. MOTIVATION FOR STATICALLY OPTIMAL TREATMENT RULES

The history-adjusted mean (4) is a parameter of both the full data and the treatment mechanism in the observed data [3]. To see this, consider that assignment of treatment up to time j affects $V(j)$ values at time j (i.e. membership in the subpopulations, or strata, of interest at time j). Thus, the observed treatment up till time j can affect the counterfactual mean within the strata of interest.

In other words, if treatment up till time j had been assigned differently (for example, if treatment up till time j had been assigned randomly), members of a given subpopulation of interest at time j (individuals for whom $V(j)=v(j)$) would not necessarily be exchangeable with the corresponding subpopulation in the observed data. Specifically, the two groups would differ as a result of differences in the covariates used to assign treatment. To the extent that these covariates also affect outcome, estimates of the effect of future treatment would also differ between the two groups. As a result, the optimal future treatment plan identified for a given stratum of $V(j)$ in the observed data can fail to optimize outcome if applied to the corresponding stratum in an experiment where treatment was assigned differently.

Thus, the dynamic treatment regime estimated by HA-MSM depends on the observed treatment mechanism; a treatment rule estimated from an observational cohort might not remain statically optimal if applied to an identical cohort with a different treatment mechanism. This is particularly troubling, in the sense that the rule itself is a treatment mechanism. Thus if a HA-MSM-derived treatment rule were applied to an identical population beginning at time 0 (the start of follow up), at later time points the rule might no longer continue to select the optimal treatment plan.

For example, in the HIV analyses, HA-MSM were used to identify the best future treatment plan (switch immediately or wait to switch) among the subpopulation who had not yet switched therapy. Membership in the subpopulation of individuals who remained on their non-suppressive therapy changed over time, as a result of the process for deciding when to switch therapy (the treatment mechanism). Specifically, the treatment mechanism revealed that individuals were more likely to switch therapy if they had lower CD4 T cell counts. As a result, the population that remained on non-suppressive therapy at a given time point, among whom the best future treatment plan was estimated, contained a disproportionate number of individuals who had maintained high CD4 T cell counts. The subpopulation that remained on non-suppressive therapy over time would have differed if the decision process for switching treatment had differed. For example, if treatment had been assigned randomly, the subpopulation remaining on non-suppressive therapy at a given time point would have included more individuals with low prior CD4 T cell counts. To the extent that past CD4 T cell count affects future CD4 T cell count, the



expected future CD4 T cell count would differ would differ between the two subpopulations, and thus the optimal treatment decision could differ.

In other words, since membership in a given stratum of interest can depend on the way that treatment was assigned, the effect of future treatment decisions and thus the optimal treatment plan for the stratum can depend on the observed treatment mechanism. Our goal, however, is to identify stratum-specific future treatment plans that are expected to optimize future outcome *regardless* of the way in which past treatment was assigned. The resulting statically optimal individualized treatment rules will then continue to identify the optimal future plan at each time point if applied to a comparable population beginning at any time point during follow up. For example, a statically optimal rule for deciding when to switch therapy would choose the best treatment plan (switch or not) among people remaining on their non-suppressive regimen at a given time point, regardless of how the decision to switch was made up till that time point. Thus a statically optimal rule will identify the optimal treatment plan at each time point if applied in the context of a clinical trial (where the rule of interest has been applied since time 0), or alternatively, if applied in the context of clinical practice (where some other decision process has been applied since time 0).

5. STATICALLY OPTIMAL TREATMENT RULES

HA-MSM identify statically optimal individualized treatment rules when the history-adjusted mean no longer depends on the observed treatment history, or in other

words, when the observed history-adjusted mean equals the counterfactual history-adjusted mean:

$$E(Y(j+m)_{\bar{A}(j-1)\underline{a}(j)} | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)) = E(Y(j+m)_{\bar{a}(j-1)\underline{a}(j)} | \bar{S}(j)_{\bar{a}(j-1)}). \quad (6)$$

When equality (6) holds, then the optimal future treatment plan at each time point j estimated by HA-MSM will be the same as the optimal future treatment plan at each time point j under any fixed treatment history up till that time point:

$$\begin{aligned} \underline{a}^*(j | V(j)) &\equiv \arg \max_{\underline{a}(j)} E(Y(j+m)_{\bar{A}(j-1)\underline{a}(j)} | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)) \\ &= \arg \max_{\underline{a}(j)} E(Y(j+m)_{\bar{a}(j-1)\underline{a}(j)} | \bar{S}(j)_{\bar{a}(j-1)}) \end{aligned}$$

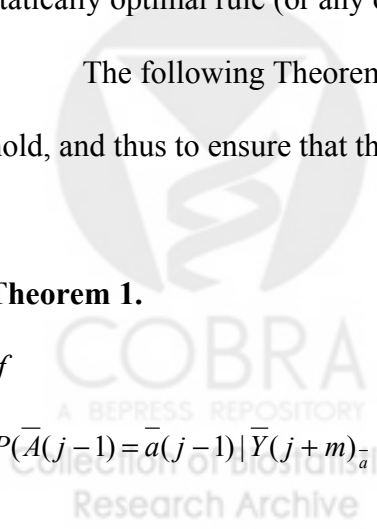
Note that, given the covariates on which the statically optimal rule is based ($\bar{S}(j)$), the treatment history that corresponds to following the statically optimal rule itself is simply one of these possible fixed treatment histories (given $\bar{S}(j)$, the statically optimal rule deterministically assigns a treatment $a(t)$, $t=0, \dots, j-1$). Thus, when equality (6) holds, then the optimal future treatment plan at a given time point estimated by HA-MSM will indeed identify the optimal future treatment plan if the study population had been following the statically optimal rule (or any other treatment mechanism) up till that time point.

The following Theorem presents an equality that is sufficient for identity (6) to hold, and thus to ensure that the rules estimated by HA-MSM are statically optimal.

Theorem 1.

If

$$P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{Y}(j+m)_{\bar{a}} = y(j+m), \bar{S}(j)_{\bar{a}} = \bar{s}(j)) = P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{S}(j)_{\bar{a}} = \bar{s}(j)) \quad (7)$$

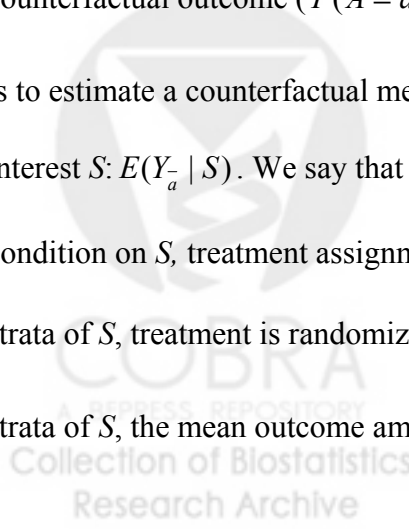


then $E(Y(j+m)_{\bar{A}(j-1)\underline{a}(j)} | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)) = E(Y(j+m)_{\bar{a}(j-1)\underline{a}(j)} | \bar{S}(j)_{\bar{a}(j-1)})$ and the HA-MSM derived rule is statically optimal.

The proof of Theorem 1 is presented in Appendix 1. Here, we present the intuition behind the theorem, which further provides insight into the conditions, discussed below, for the identity (7) to hold.

Non-identity between the observed history-adjusted mean and the counterfactual history-adjusted mean can be seen as a problem of confounding. Consider the case where the goal is to estimate a simple counterfactual mean under treatment history $\bar{a} : E(Y_{\bar{a}})$.

Clearly, unless the entire study population, or a random sample, received treatment history \bar{a} , $E(Y_{\bar{a}})$ is not necessarily equal to the mean outcome among people who received treatment history \bar{a} in the observed data ($E(Y_{\bar{a}}) \neq E(Y_{\bar{A}} | \bar{A} = \bar{a})$). This is a classic example of confounding; the two parameters will differ to the extent there are covariates that affect both treatment assignment and outcome, making treatment assignment dependent on counterfactual outcome ($P(\bar{A} = \bar{a} | Y_{\bar{a}}) \neq P(\bar{A} = \bar{a})$). Next, consider the case where the goal is to estimate a counterfactual mean under treatment history \bar{a} , given baseline covariates of interest $S : E(Y_{\bar{a}} | S)$. We say that S is sufficient to control for confounding if, once we condition on S , treatment assignment no longer depends on counterfactual outcome (within strata of S , treatment is randomized): $P(\bar{A} = \bar{a} | S, Y_{\bar{a}}) = P(\bar{A} = \bar{a} | S)$. In this case, within strata of S , the mean outcome among people who received treatment history \bar{a} in the



observed data is equivalent to the mean counterfactual outcome under treatment \bar{a} :

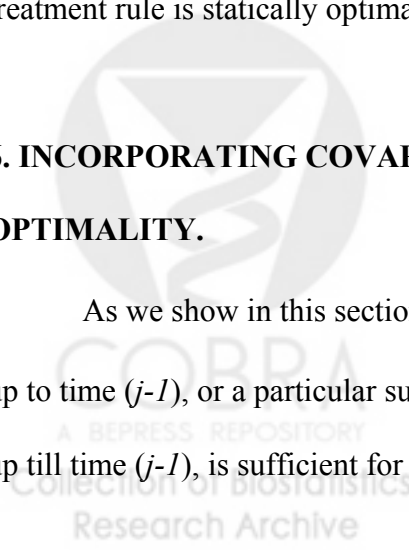
$E(Y_{\bar{a}} | \bar{A} = \bar{a}, S) = E(Y_{\bar{a}} | S)$. This concept is the motivation for the use of multivariable regression for causal inference.

The same concept can be extended to the current situation, only now only considering the equivalence of mean outcomes indexed by observed vs. counterfactual treatment *through time* $j-1$. In order for equality (6) to hold, the covariates of interest $\bar{S}(j)$ must be sufficient to control for confounding of treatment assignment through time $j-1$ ($\bar{A}(j-1)$) on counterfactual outcome $Y_{\bar{A}(j-1)\bar{a}(j)}$. In other words, if treatment assignment through time $j-1$ is independent of counterfactual outcome, given covariates of interest (equality (7)), then the counterfactual history-adjusted mean will equal the observed history-adjusted mean (equality (6)).

If equality (7) holds, then the HA-MSM parameter is no longer dependent on the treatment mechanism. In the following two sections, we discuss choices of $\bar{S}(j)$ sufficient for equality (7) to hold, and thus sufficient to ensure that the HA-MSM individualized treatment rule is statically optimal.

6. INCORPORATING COVARIATE HISTORY TO ENSURE STATIC OPTIMALITY.

As we show in this section, inclusion in $\bar{S}(j)$ of either the entire covariate history up to time $(j-1)$, or a particular summary of the covariates that affect treatment assignment up till time $(j-1)$, is sufficient for equality (7) to hold. The intuitive motivation for this



approach again relies on confounding. We know from Section 5 that if $\bar{S}(j)$ is sufficient to control for confounding of $\bar{A}(j-1)$ on counterfactual outcome $Y_{\bar{A}(j-1)\bar{a}(j)}$, then equality (7) will hold. We further know that confounding arises as a result of covariates that affect both treatment assignment and outcome. Thus, by including in $\bar{S}(j)$ all covariates that affect treatment assignment, we ensure that the HA-MSM-derived treatment rules are statically optimal. We present this result as a Lemma, below (the proof is provided in Appendix 2).

Lemma 1.

If $P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x)$ is only a function of $\bar{S}(j-1)_a \subset \bar{S}(j)_a$, then

$$P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) = P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{S}(j)_a = \bar{s}(j))$$

and (by Theorem 1) the HA-MSM parameter is statically optimal.

Several choices of $\bar{S}(j)$ are sufficient to ensure that

$P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x)$ is only a function of $\bar{S}(j-1)_a \subset \bar{S}(j)_a$, and thus that the HA-MSM individualized treatment rules are statically optimal. For example, this is true if the HA-MSM parameter conditions on the entire covariate history in addition to any effect modifiers of interest. Under the SRA (2), $V(j) = (\bar{S}(j), \bar{A}(j-1))$ then includes all covariates which affect treatment assignment up till time j .

Sufficient Condition 1: $(\bar{S}(j) = (\bar{L}(j-1), \bar{S}^*(j)))$,

where $\bar{S}^*(j)$ denotes the effect modifiers of interest

Including the entire covariate history in the HA-MSM model removes the dependency of the resulting parameter on the treatment mechanism. However, when several covariates are measured over multiple time points, the covariate history $\bar{L}(j)$ can become very high dimensional. Specifying a reasonable HA-MSM model for $E[Y(j+m)_{\bar{A}(j-1)a(j)} | V(j)]$ can thus become challenging or unfeasible.

An alternative to conditioning on the entire covariate history is to condition only those covariates that affect treatment assignment up till time $(j-1)$. The covariates affecting treatment assignment can be summarized by including the treatment mechanism itself (through time $j-1$) as a covariate in $\bar{S}(j)$:

Sufficient Condition 2: $\bar{S}(j) = (g(\bar{A}(j-1) | X^{Full}), \bar{S}^*(j))$,

where $\bar{S}^*(j)$ denotes the effect modifiers of interest

When the treatment mechanism interacts with the exposure of interest, the optimal future treatment plan at time j will depend on the treatment mechanism up till time j in addition to other covariates of interest. In this setting, implementing the resulting rule in practice will require measuring not only the effect modifiers of interest, but also all other covariates that contribute to the treatment mechanism.

7. CAUSAL STRUCTURES SUFFICIENT TO ENSURE STATIC OPTIMALITY

The previous section showed how conditioning the HA-MSM parameter on covariate history in addition to effect modifiers of interest can remove dependency on the observed treatment assignment and permit estimation of a statically optimal treatment rule.

Both approaches presented rely, however, on adjusting for *all* covariates that affect treatment assignment. In contrast, Pearl's causal graph theory [12] provides a tool for identifying a minimal set of covariates sufficient to control for confounding. In this section, we apply *d-separation* from DAG methodology to introduce an alternative condition for $\bar{S}(j)$ sufficient to ensure static optimality of the HA-MSM parameter. Specifically, we show that, under certain causal structures, $\bar{S}(j)$ need not include covariate history (either in its entirety or summarized based on the treatment mechanism) to ensure the static optimality of the treatment rule estimated. These results introduce greater flexibility in the choice of sufficient $\bar{S}(j)$, and provide a graphical tool to identify sufficient alternatives.

Pearl defines the graphical criteria for *d-separation* as follows (Definition 1.2.3):

Definition of d-separation:

“A path p is said to be d-separated (or blocked) by a set of nodes Z if and only if

1. p contains a chain $i \rightarrow m \rightarrow j$ or a fork $i \leftarrow m \rightarrow j$ such that the middle node m is in Z , or

2. p contains an inverted fork (or collider) $i \rightarrow m \leftarrow j$ such that the middle node m is not in Z and such that no descendent of m is in Z

A set Z is said to d-separate X from Y if and only if Z blocks every path from a node in X to a node in Y .”

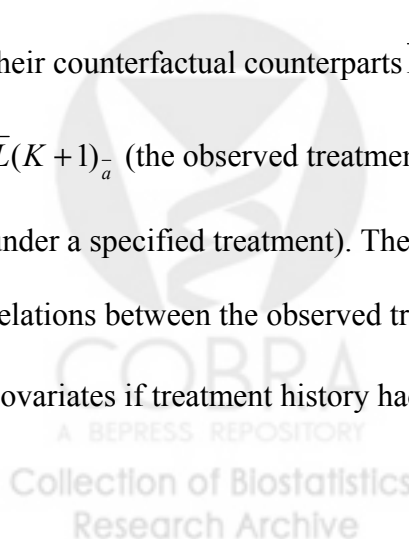
Pearl shows (Theorem 1.2.4) that “If sets X and Y are *d-separated* by Z in a DAG G , then X is independent of Y conditional on Z in every distribution compatible with G ”.

Thus, d -separation of X and Y by Z (written $(X \perp\!\!\!\perp Y \mid Z)_G$) implies that X is conditionally independent of Y given Z (written $(X \perp\!\!\!\perp Y \mid Z)_P$), and thus that $P(X = x \mid Y = y, Z = z) = P(X = x \mid Z = z)$. This theorem provides us with a graphical criterion (presented as a Lemma) for evaluating the sufficiency of $\bar{S}(j)$ for equality (7) to hold.

Lemma 2. *If $\bar{S}_a(j)$ d -separates $\bar{A}(j-1)$ from $Y(j+m)_a$, then $\bar{S}(j)$ is sufficient to ensure that*

$P(\bar{A}(j-1) = \bar{a}(j-1) \mid Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) = P(\bar{A}(j-1) = \bar{a}(j-1) \mid \bar{S}(j)_a = \bar{s}(j))$
and thus (by Theorem 1) that the HA-MSM parameter is statically optimal.

In applying the graphical criterion of d -separation, a DAG must first be converted from a graph showing causal relationships in the observed data, to a graph showing causal relations between the observed treatment and the counterfactual covariate processes. This process involves two steps (Figure 1). First, replace the observed covariates $\bar{L}(K+1)$ by their counterfactual counterparts $\bar{L}(K+1)_a$. Second, erase all arrows from $\bar{A}(K)$ to $\bar{L}(K+1)_a$ (the observed treatment does not affect the counterfactual values of covariates under a specified treatment). The resulting causal graph (Figure 1b) now shows the causal relations between the observed treatment history and the counterfactual values of covariates if treatment history had been set at \bar{a} .



The *d-separation* criterion of Lemma 2 suggests an alternative graph-based proof to Lemma 1. In the causal graph corresponding to the observed treatment history and counterfactual covariate values (Figure 1b), all paths from treatment $\bar{A}(K)$ to counterfactual covariates $\bar{L}(K+1)_a$ are deleted. Thus, the only possible paths connecting $\bar{A}(j-1)$ to $Y(j+m)_a$ are “backdoor paths” (i.e. paths from $\bar{A}(j-1)$ to $Y(j+m)_a$ via arrows into $\bar{A}(j-1)$). Graphically, Lemma 1 states that $\bar{S}(j)$ is sufficient for equality (7) and hence static optimality to hold if $\bar{S}(j-1)_a \subset \bar{S}(j)_a$ includes all covariates with causal arrows into $\bar{A}(j-1)$. But if $\bar{S}_a(j)$ includes all covariates with arrows into $\bar{A}(j-1)$, then $\bar{S}_a(j)$ must block all backdoor paths, and thus all paths, from $\bar{A}(j-1)$ to $Y(j+m)_a$, and thus ensure their *d-separation*. As result, by Lemma 2, equation (7) must hold and $\bar{S}(j)$ is sufficient to ensure static optimality.

The graphical criterion of *d-separation* provides a minimal condition for $\bar{S}(j)$ to be sufficient. As shown, if $\bar{S}(j)_a$ includes all covariates with causal arrows into $\bar{A}(j-1)$ (as required by Lemma 1), then *d-separation* holds. In addition, *d-separation* holds under many causal structures where $\bar{S}(j)_a$ does not include all covariates with causal arrows into $\bar{A}(j-1)$, and thus where neither Sufficient Conditions 1 nor 2 are met (Figure 2).

Given a DAG, Lemma 2 provides a graphical tool for suggesting alternative choices of $\bar{S}(j)$, and for evaluating their sufficiency. Figures 2 and 3 present several sample causal graphs, illustrating examples of $\bar{S}(j)$ sufficient and insufficient to ensure

static optimality. Note, however, that the sufficiency of a particular choice of $\bar{S}(j)$, other than a choice fulfilling Lemma 1, will depend on the specific causal structure of the data. Specifying this causal structure becomes more challenging as the number of covariates and time points increases.

8. RESULTS: STATICALLY OPTIMAL TREATMENT RULES FOR DECIDING WHEN TO SWITCH THERAPY

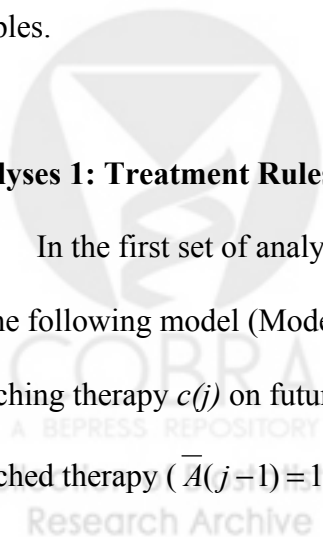
Two sets of analyses were performed, aimed at estimating rules among individuals who had not been re-suppressed based on two covariates of interest, CD4 T cell count at time j and diagnosis with an opportunistic disease at time j

($\bar{S}^*(j) = (Sup(j) = 0, CD4(j))$, and $\bar{S}^*(j) = (Sup(j) = 0, OD(j))$, respectively). All HA-MSM models were fit using the inverse probability of treatment weighted estimator [3, 4]. The treatment mechanism (probability of switching given the past) was fit data-adaptively using the Deletion/Substitution/Addition algorithm and 5-fold cross validation [13].

Confidence intervals for model coefficients are based on 100 non-parametric bootstrap samples.

Analyses 1: Treatment Rules Based on Opportunistic Disease

In the first set of analyses, we first estimated an individualized treatment rule based on the following model (Model 1), which estimates the effect of additional time until switching therapy $c(j)$ on future CD4 T cell count among individuals who have not yet switched therapy ($\bar{A}(j-1) = 1$) or been re-suppressed ($Sup(j) = 0$), given elapsed time since



virologic failure occurred (j) and current diagnosis with an AIDS-defining opportunistic disease ($OD(j)$):

$$E(Y(j+m)_{\bar{A}(j-1)c(j)} | \bar{A}(j-1) = 1, Sup(j) = 0, OD(j)) = \beta_0 + \beta_1 c(j) + \beta_2 j + \beta_3 j \times c(j) + \beta_4 OD(j) + \beta_5 OD(j) \times c(j)$$

The corresponding estimate of the effect of waiting to switch (Table 1) yielded the following treatment rule: In individuals with a current diagnosis of an opportunistic disease, stay on the same therapy and re-evaluate the following month. In individuals without such a diagnosis, switch therapy immediately if less than 8 months have elapsed since loss of suppression occurred, otherwise wait to switch.

In the absence of an explicit causal structure, there is no guarantee that this treatment rule would continue to optimize outcome if it were applied to an identical population, beginning when loss of suppression occurred. In order to estimate a statically optimal treatment rule, the treatment mechanism through time $j-1$ was incorporated into the HA-MSM model (Model 2):

$$E(Y(j+m)_{\bar{A}(j-1)c(j)} | \bar{A}(j-1) = 1, Sup(j) = 0, OD(j)) = \beta_0 + \beta_1 c(j) + \beta_2 j + \beta_3 j \times c(j) + \beta_4 OD(j) + \beta_5 OD(j) \times c(j) + \beta_6 \prod_{l=0}^{j-1} g(A(l) | \bar{A}(l-1), \bar{L}(l))$$

(Alternatively, the treatment mechanism could have been incorporated as an interaction with the exposure of interest ($c(j)$) and/or the other covariates. However, we found no evidence of interactions between the treatment mechanism and time until switching therapy.) The estimated statically optimal rule for deciding when to switch given current opportunistic disease and elapsed time (Table 2), states that, regardless of diagnosis of with opportunistic disease and time elapsed since loss of suppression occurred, treatment should

be switched immediately. The latter rule, rather than the former, would be expected to optimize patient outcome if applied to a comparable population in a clinical trial. In fact, such a rule makes more clinical sense than the former, non-statically optimal, rule in that waiting to switch therapy is unlikely to be more beneficial among those who have a current opportunistic disease.

Analyses 2: Treatment Rules Based on CD4 T Cell Count

In the second set of analyses, we estimated an individualized treatment rule based on the estimated effect of future time until switching ($c(j)$) on future CD4 T cell count among individuals who had not yet switched therapy or been re-suppressed, given current CD4 T cell count ($CD4(j)$) and elapsed time (j) (Table 3). This effect was estimated using the following HA-MSM (Model 3):

$$E(Y(j+m)_{\bar{A}(j-1)c(j)} | \bar{A}(j-1)=1, Sup(j)=0, CD4(j)) = \beta_0 + \beta_1 c(j) + \beta_2 CD4(j) + \beta_3 j + \beta_4 c(j) \times CD4(j) + \beta_5 c(j) \times j + \beta_6 CD4(j) \times j + \beta_7 c(j) \times CD4(j) \times j$$

To estimate a statically optimal individualized treatment rule, we then fit an additional HA-MSM, now including the treatment mechanism as a covariate. Again, after considering potential multi-way interactions, our final HA-MSM model included the treatment mechanism as a main term only (Model 4):

$$E(Y(j+m)_{\bar{A}(j-1)c(j)} | \bar{A}(j-1)=1, Sup(j)=0, CD4(j)) = \beta_0 + \beta_1 c(j) + \beta_2 CD4(j) + \beta_3 j + \beta_4 c(j) \times CD4(j) + \beta_5 c(j) \times j + \beta_6 CD4(j) \times j + \beta_7 c(j) \times CD4(j) \times j + \beta_8 \prod_{l=0}^{j-1} g(A(l) | \bar{A}(l-1), \bar{L}(l))$$

The resulting effect estimates yielded the following statically optimal treatment rule (Table 4):

$$d_j = I[(-9.0 + 0.05 \times CD4(j) + 1.2 \times j - 0.009 \times CD4(j) \times j) < 0]$$

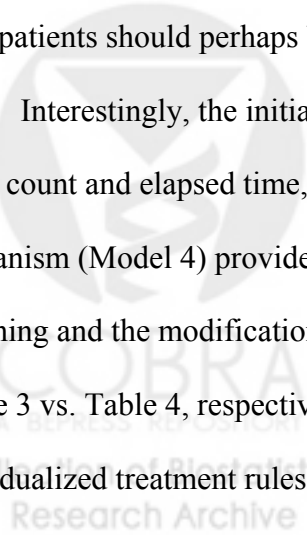
$d_j = 1$: Switch at time j

$d_j = 0$: Wait at time j

In other words, when the expected effect of waiting to switch is negative, switch the patient immediately; when the expected effect is positive, wait to switch and re-evaluate at the subsequent visit.

95% confidence intervals for each of the coefficients in this rule, based on 100 non-parametric bootstrap samples, are presented in Table 4. Alternatively, bootstrap re-sampling can be used to estimate variability in the treatment decision itself by plotting the proportion of bootstrap samples in which the statically optimal treatment rule indicates a switch for a range of CD4 T cell counts and elapsed times since loss of suppression. Such an analysis (Figure 4) suggests that, early after loss of suppression occurs, there is strong evidence that patients with low CD4 T cell count should be switched to a new regimen, while patients with high CD4 T-cell counts can afford to wait. In contrast, among patients who have already spent 5 months on non-suppressive therapy, variability in the decision to switch depends less on current CD4 T cell count, suggesting that the decision to switch for these patients should perhaps be based on other factors.

Interestingly, the initial HA-MSM fit (Model 3), which included only current CD4 T cell count and elapsed time, and the HA-MSM also adjusting for the treatment mechanism (Model 4) provided very similar estimates of the effect of future time until switching and the modification of this effect by current CD4 T cell count and elapsed time (Table 3 vs. Table 4, respectively). As a result, the two approaches yielded very similar individualized treatment rules. The small change in coefficients in the CD4 T cell –based



treatment rule after adjusting for the treatment mechanism can be explained by the postulated causal structure of the data (Figure 5). This DAG also illustrates why the opportunistic disease-based treatment rule was altered by incorporation of the treatment mechanism. In this DAG, CD4 T cell count at time j , but not opportunistic disease diagnosis at time j , is sufficient to *d-separate* treatment history prior to time j from CD4 T cell count m months in the future.

9. DISCUSSION

As mentioned in the introduction, the history-adjusted statically optimal treatment rules discussed in this paper must be differentiated from the optimal dynamic treatment regimes, as estimated by Murphy, *et. al.* [1-3]. A statically optimal rule assigns a patient the first action of the future *static* treatment plan that will optimize his or her expected outcome, given the patient's covariates. The plan can then be updated at subsequent time points in response to changes in patient covariates. In contrast, an optimal dynamic regime selects, at a single point in time, the future *rule* that will optimize expected outcome at a fixed point in the future. Thus the statically optimal rule selects between different static plans, and then updates this selection at subsequent time points, while the optimal dynamic regime selects among candidate *rules* at a single time point.

The two approaches estimate dynamic rules under very different models. There are clearly scenarios where the statically optimal rule results in an inferior expected outcome than the optimal dynamic regime; a simulation illustrating such a scenario is given in [3]. However, in essentially all cases, the models on which both approaches are based will be

incorrectly specified. Future research is needed on the relative performance of the two types of estimated rules in such settings.

The statically optimal rule does offer several practical advantages. Unlike the truly optimal dynamic rule, the statically optimal rule can be estimated using standard software. In addition, the statically optimal rule provides increased flexibility in the outcome that is optimized. The optimal dynamic rule optimizes outcome at a fixed time point, while the statically optimal rule can optimize outcome either at a fixed or moving time point. In many clinical settings, the latter type of outcome is of greater clinical interest. For example, the goal of the clinician in managing resistant HIV is to continually maintain future CD4 T cell count (and patient health); thus, as the patient is followed over time, the interval over which the clinician aims to optimize outcome continually shifts forward.

In summary, HA-MSM estimate modification of causal effects by time-varying variables, and thus can be used to identify individualized treatment rules for modifying treatment in response to a patient's changing covariates. However, the dependence of the HA-MSM parameter on the observed treatment mechanism, or the way in which the assignment of treatment at each time point depended on a subject's past, means that such rules would not necessarily remain statically optimal if applied in a clinical trial. As we have illustrated, treatment rules which are truly statically optimal, or in other words, which no longer depend on the treatment mechanism, can be estimated using HA-MSM given that the covariates on which the rule depends are sufficient to control for confounding of the effect of past treatment on outcome. As a result, by selecting the correct covariates,

HA-MSM can be applied to observational data to identify individualized treatment rules appropriate for evaluation in a clinical trial, or for use in clinical practice.

Causal graphs can be used to identify whether a candidate set of covariates is sufficient to control confounding and thus ensure static optimality, as well as to suggest alternative sets of sufficient covariates. This approach does not rely on consistent estimation of the treatment mechanism, thus the Double Robust HA-MSM estimator [3] will remain consistent if either the treatment mechanism or data-generating distribution are correctly estimated. This approach does, however, rely on use of a DAG that accurately represents the causal relations in the data, and thus relies heavily of the background knowledge of the researcher. As Figure 4 illustrates, the causal relationships in longitudinal data are often complex, and their correct specification may be unpractical.

Inclusion of the entire covariate history, or alternatively, of the treatment mechanism in the HA-MSM model also ensures static optimality, regardless of the specific causal context. Under this approach, the static optimality of the dynamic treatment rule depends on consistent estimation of the treatment mechanism. As a result, the Double Robust HA-MSM estimator [3] will not protect against misspecification of the treatment mechanism when deriving candidate treatment rules, although it may provide gains in efficiency. However, the treatment mechanism-based approach has the advantage of generally requiring a much lower dimensional HA-MSM model than the approach which adjusts for the entire covariate history. The reliance of the former approach on consistent estimation of the treatment mechanism also implies that in settings where the treatment mechanism is known (such as in the context of sequentially randomized trials) HA-MSM

can estimate statically optimal treatment rules without any assumptions beyond the HA-MSM model itself.



REFERENCES:


1. Murphy SA. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B* 2003,65:331-584.
2. Robins JM. Discussion of optimal treatment regimens by Susan A. Murphy. *Journal of the Royal Statistical Society: Series B* 2003,65:355-366.
3. van der Laan MJ, Petersen ML, Joffe MM. History-Adjusted Marginal Structural Models and Statically-Optimal Dynamic Treatment Regimens. *The International Journal of Biostatistics* 2005,1:Article 4. <http://www.bepress.com/ijb/vol1/iss1/4/>
4. Petersen ML, Deeks SG, Martin JM, van der Laan MJ. History-Adjusted Marginal Structural Models to Estimate Time-Varying Effect Modification. *U.C. Berkeley Division of Biostatistics Working Paper Series* 2005:Working Paper 199. <http://www.bepress.com/ucbbiostat/paper199/>
5. Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. *Proceedings of the American Statistical Association* 2000:6-10.
6. Robins JM. Marginal structural models versus structural nested models as tools for causal inference. In: *Statistical models in epidemiology, the environment, and clinical trials*. Edited by Halloran ME, Berry D. New York: Springer-Verlag; 1999:95-134.
7. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000,11:550-560.

8. Hernan MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* 2000,11:561-570.
9. Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. *Epidemiology* 1999,10:37-81.
10. Detels R, Munoz A, McFarlane G, *et al.* Effectiveness of potent antiretroviral therapy on time to AIDS and death in men with known HIV infection duration. Multicenter AIDS Cohort Study Investigators. *JAMA* 1998,280:1497-1503.
11. Deeks SG. Treatment of antiretroviral-drug-resistant HIV-1 infection. *Lancet* 2003,362:2002-2011.
12. Pearl J. *Causality: Models, Reasoning, and Inference*. Cambridge, UK: Cambridge University Press, 2000; pp 11-18.
13. Sinisi S, van der Laan M. Deletion/Substitution/Addition Algorithm in Learning with Applications in Genomics. *Statistical Applications in Genetics and Molecular Biology* 2004,3:18.




FIGURE 1. Illustration of DAG manipulation required prior to evaluation of d -separation to test whether a given $\bar{S}(j)$ is sufficient to ensure static optimality¹.

A. Causal relationships in the observed data.



B. Causal relationships between observed treatment and counterfactual covariates.



¹ Figure 1 presents one possible basic causal structure for longitudinal data. To improve the simplicity of the graphs involved, we present the situation where $j=2$, $K=3$, and $Y(j+m)=Y(j+2)$. Similar reasoning can be used for larger j , K and m .

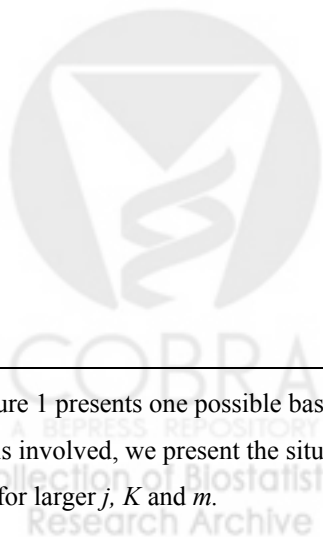




FIGURE 2. DAG-based examples of $\bar{S}(j)$ sufficient to d -separate $\bar{A}(j-1)$ from $Y(j+m)_a$ in specific causal settings.¹


A. Example 1. $S_a(j) = L(j)_a$



B. Example 2. $L(j)_a = (S(j)_a, C(j)_a)$



C. Example 3. $L(j)_a = (S(j)_a, C(j)_a)$



¹To simplify presentation, we use $j=2$, $K=3$, and $Y(j+m)=Y(j+1)$, and covariates and treatment occurring after time j are not shown..

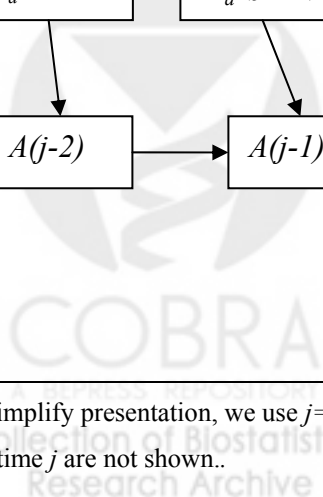
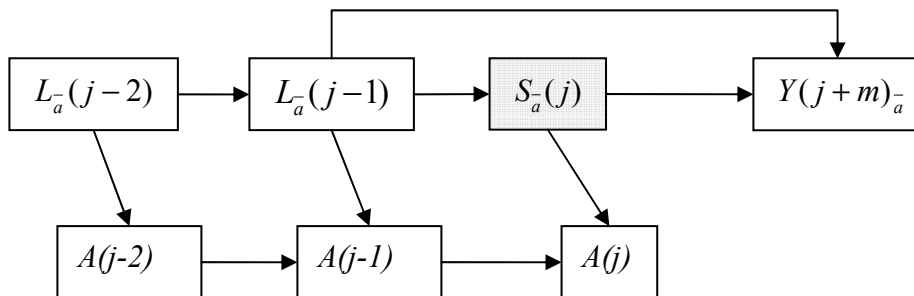



FIGURE 3. DAG-based examples of $\bar{S}(j)$ not sufficient to d -separate $\bar{A}(j-1)$ from $Y(j+m)_a$ in specific causal settings.¹

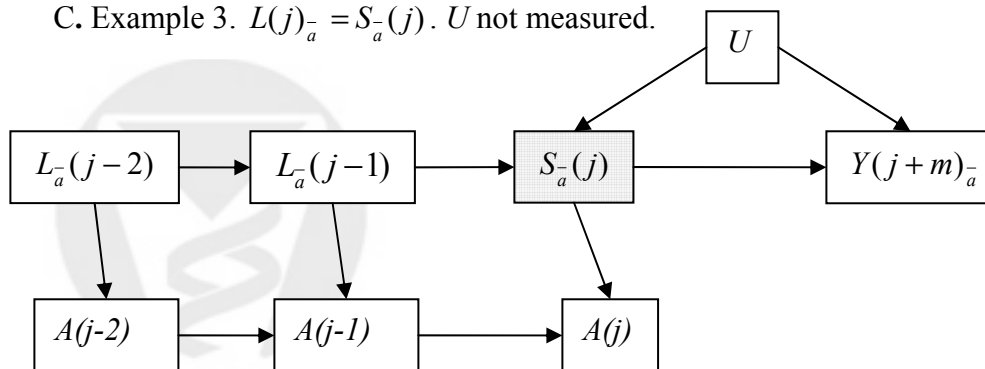
A. Example 1. $L(j)_a = S_a(j)$



B. Example 2. $L(j)_a = (S(j)_a, C(j)_a)$



C. Example 3. $L(j)_a = S_a(j)$. U not measured.



¹ To simplify presentation, we use $j=2$, $K=3$, and $Y(j+m)=Y(j+1)$, and covariates and treatment occurring after time j are not shown.

FIGURE 4. Variability in treatment decision indicated by statically optimal treatment rule, according to CD4 T cell count and elapsed time since loss of suppression occurred.

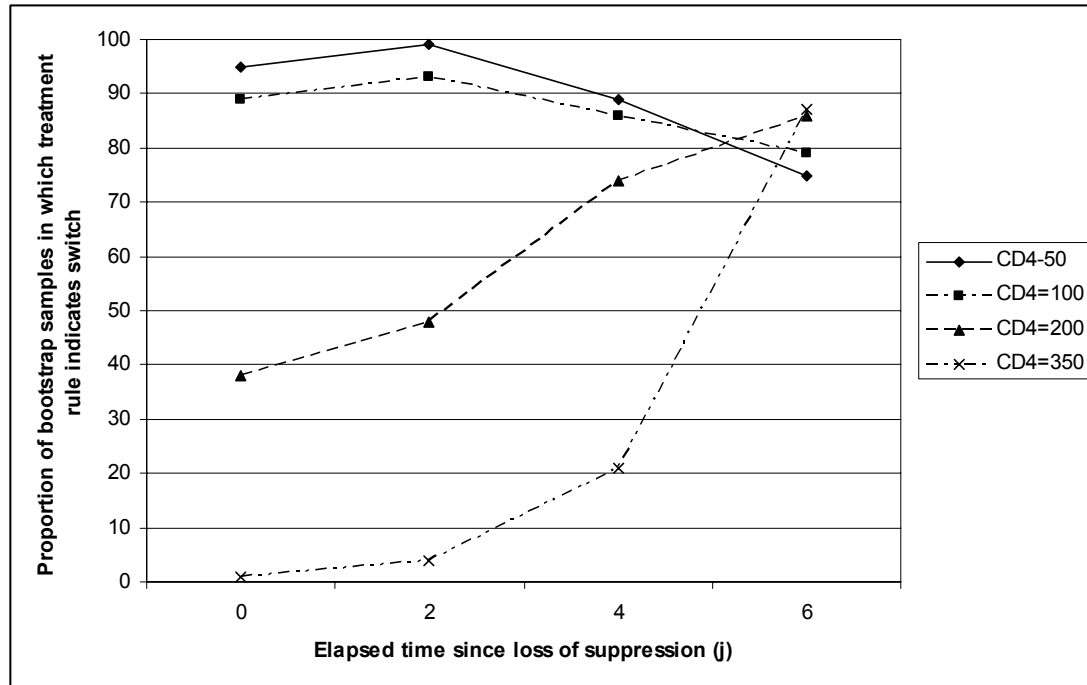
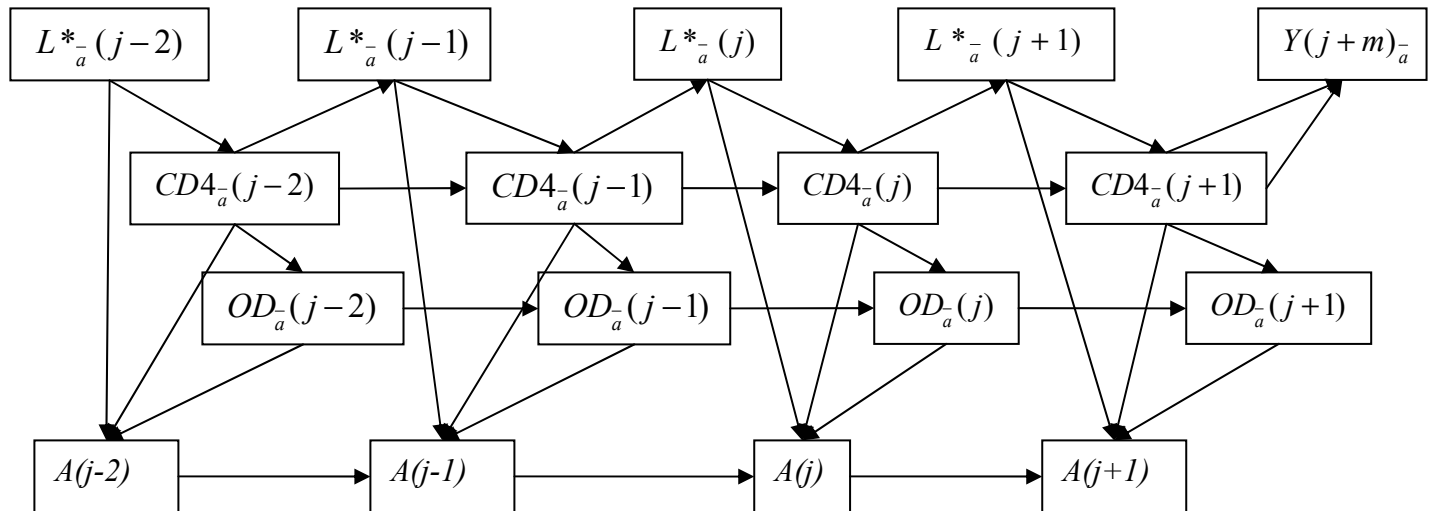


FIGURE 5. Proposed DAG for HIV data example. $\bar{S}(j) = CD4(j)$, but not $\bar{S}(j) = OD(j)$, sufficient for static optimality. Similarly, $CD4(t)$, but not $OD(t)$ ($t=0, \dots, K$), sufficient to control for confounding¹.



¹ To simplify presentation, we use $j=2$, $K=3$, and $Y(j+m)=Y(j+2)$.

$L^*_a(j)$ = additional components of the treatment mechanism (e.g. adherence)

TABLE 1. Individualized treatment rule for analysis 1, based on Model 1 (without treatment mechanism as covariate)

Coefficient	95% CI ^{1,2}
$\beta_1 = -4.0$	-15.1, 7.1
$\beta_3 = 0.56 \times j$	-1.9, 3.1
$\beta_5 = 6.4 \times OD(j)$	-12.4, 25.2

$$d_j = I[(\beta_1 + \beta_3 \times j + \beta_5 \times OD(j)) < 0]$$

Treatment Rule: $d_j = 1$: Switch at time j

$d_j = 0$: Wait at time j



¹ Based on 100 Bootstrap samples

² CI = Confidence Interval

TABLE 2. Individualized treatment rule for analysis 1, based on Model 2 (incorporating treatment mechanism as covariate)

Coefficient	95% CI ^{1 2}
$\beta_1 = -2.6$	-13.1, 7.9
$\beta_3 = -0.28 \times j$	-2.3, 1.8
$\beta_5 = 0.83 \times OD(j)$	-15.7, 17.2

$$d_j = I[(\beta_1 + \beta_3 \times j + \beta_5 \times OD(j)) < 0]$$

Treatment Rule: $d_j = 1$: Switch at time j

$d_j = 0$: Wait at time j



¹ Based on 100 Bootstrap samples

² CI = Confidence Interval

TABLE 3. Individualized treatment rule for analysis 2, based on Model 3 (without treatment mechanism as covariate)

Coefficient	95% CI ^{1 2}
$\beta_1 = -9.2$	-17.6, -7.6
$\beta_4 = 0.05 \times CD4(j)$	0.02, 0.08
$\beta_5 = 1.5 \times j$	-0.4, 3.4
$\beta_7 = -0.009 \times CD4(j) \times j$	-0.02, -0.004

$$d_j = I[(\beta_1 + \beta_4 \times CD4(j) + \beta_5 \times j + \beta_7 \times CD4(j) \times j) < 0]$$

Treatment Rule: $d_j = 1$: Switch at time j

$d_j = 0$: Wait at time j



¹ Based on 100 Bootstrap samples

² CI = Confidence Interval

TABLE 4. Individualized treatment rule based on Model 4 (incorporating treatment mechanism as covariate)

Coefficient	95% CI ^{1 2}
$\beta_1 = -9.0$	-17.5, -0.5
$\beta_4 = 0.05 \times CD4(j)$	0.02, 0.08
$\beta_5 = 1.2 \times j$	-0.7, 3.1
$\beta_7 = -0.009 \times CD4(j) \times j$	-0.02, -0.003

$$d_j = I[(\beta_1 + \beta_4 \times CD4(j) + \beta_5 \times j + \beta_7 \times CD4(j) \times j) < 0]$$

Treatment Rule: $d_j = 1$: Switch at time j

$d_j = 0$: Wait at time j



¹ Based on 100 Bootstrap samples

² CI = Confidence Interval

APPENDIX 1: Proof of Theorem 1- equality sufficient to ensure that HA-MSM estimate statically optimal treatment rules.

Theorem 1.

If

$$P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{Y}(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) = P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{S}(j)_a = \bar{s}(j)) \quad (7)$$

then $E(Y(j+m)_{\bar{A}(j-1)\bar{a}(j)} | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)) = E(Y(j+m)_{\bar{A}(j-1)\bar{a}(j)} | \bar{S}(j)_{\bar{a}(j-1)})$.

Proof:

Note that

$$\begin{aligned} & E(Y(j+m)_{\bar{A}(j-1)\bar{a}(j)} | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j) = \bar{s}(j)) \\ &= E(Y(j+m)_a | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)_a = \bar{s}(j)) \\ &= \sum_y y(j+m) P(Y(j+m)_a = y(j+m) | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)_a = \bar{s}(j)) \end{aligned}$$

and

$$\begin{aligned} & P(Y(j+m)_a = y(j+m) | \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)_a = \bar{s}(j)) \\ &= \frac{P(Y(j+m)_a = y(j+m), \bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)_a = \bar{s}(j))}{P(\bar{A}(j-1) = \bar{a}(j-1), \bar{S}(j)_a = \bar{s}(j))} \\ &= \frac{P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) P(Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j))}{P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{S}(j)_a = \bar{s}(j)) P(\bar{S}(j)_a = \bar{s}(j))} \end{aligned}$$

Now note that $\frac{P(Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j))}{P(\bar{S}(j)_a = \bar{s}(j))}$ is only an F_X parameter (i.e. a

parameter of the Full Data), so it remains to show that the ratio

$$\frac{P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j))}{P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{S}(j)_a = \bar{s}(j))} \text{ does not depend on } g.$$

Specifically, if this ratio =1, or

$$P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) = P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{S}(j)_a = \bar{s}(j))$$

then this is true.



APPENDIX 2: Proof of Lemma 1- sufficient condition for equality (7) to hold, and thus for HA-MSM to estimate statically optimal treatment rules.

Lemma 1.

If $P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x)$ is only a function of $\bar{S}(j-1)_a \subset \bar{S}(j)_a$, then

$$P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) = P(\bar{A}(j-1) = \bar{a}(j-1) | \bar{S}(j)_a = \bar{s}(j))$$

and (by Theorem 1) the HA-MSM parameter is statically optimal.

Proof:

Note that $P(\cdot | Y, S) = E(P(\cdot | X, Y, S) | Y, S)$, so

$$\begin{aligned} & P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) \\ &= E(P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j), X^{Full} = x) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) \end{aligned}$$

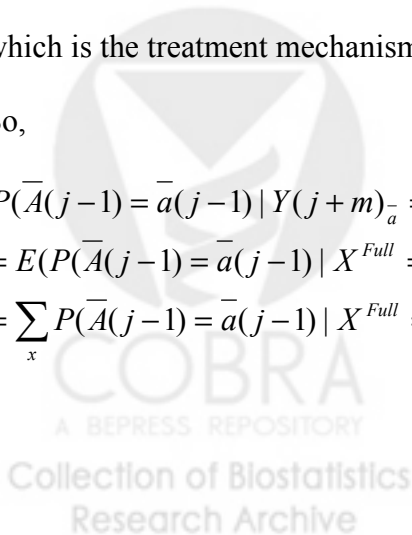
Since $Y(j+m)_a$ and $\bar{S}(j)_a$ are included in X^{Full} ,

$$P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}, X^{Full} = x) = P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x)$$

which is the treatment mechanism up till time $(j-1)$.

So,

$$\begin{aligned} & P(\bar{A}(j-1) = \bar{a}(j-1) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) \\ &= E(P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x) | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) \\ &= \sum_x P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x) P(X^{Full} = x | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) \end{aligned}$$



If now $\bar{S}(j)_a$ is such that the treatment probability $P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x)$ is

only a function of $\bar{S}(j-1)_a \subset \bar{S}(j)_a$ then it follows that, given

$Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j), P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x)$ is a constant, and thus

$$\begin{aligned} & \sum_x P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x) P(X^{Full} = x | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) \\ &= P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x) \sum_x P(X^{Full} = x | Y(j+m)_a = y(j+m), \bar{S}(j)_a = \bar{s}(j)) \\ &= P(\bar{A}(j-1) = \bar{a}(j-1) | X^{Full} = x) \end{aligned}$$

The same approach to the right hand side shows the two are equal and completes the proof.

