



Published in final edited form as:
Int J Biostat. 2007 ; 3(1): Article3.

Causal Effect Models for Realistic Individualized Treatment and Intention to Treat Rules*

Mark J. van der Laan* and Maya L. Petersen†

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

†University of California, Berkeley, mayaliv@gmail.com

Abstract

Marginal structural models (MSM) are an important class of models in causal inference. Given a longitudinal data structure observed on a sample of n independent and identically distributed experimental units, MSM model the counterfactual outcome distribution corresponding with a static treatment intervention, conditional on user-supplied baseline covariates. Identification of a static treatment regimen-specific outcome distribution based on observational data requires, beyond the standard sequential randomization assumption, the assumption that each experimental unit has positive probability of following the static treatment regimen. The latter assumption is called the experimental treatment assignment (ETA) assumption, and is parameter-specific. In many studies the ETA is violated because some of the static treatment interventions to be compared cannot be followed by all experimental units, due either to baseline characteristics or to the occurrence of certain events over time. For example, the development of adverse effects or contraindications can force a subject to stop an assigned treatment regimen.

In this article we propose causal effect models for a user-supplied set of realistic individualized treatment rules. Realistic individualized treatment rules are defined as treatment rules which always map into the set of possible treatment options. Thus, causal effect models for realistic treatment rules do not rely on the ETA assumption and are fully identifiable from the data. Further, these models can be chosen to generalize marginal structural models for static treatment interventions. The estimating function methodology of Robins and Rotnitzky (1992) (analogue to its application in Murphy, et. al. (2001) for a single treatment rule) provides us with the corresponding locally efficient double robust inverse probability of treatment weighted estimator.

In addition, we define causal effect models for “intention-to-treat” regimens. The proposed intention-to-treat interventions enforce a static intervention until the time point at which the next treatment does not belong to the set of possible treatment options, at which point the intervention is stopped. We provide locally efficient estimators of such intention-to-treat causal effects.

Keywords

counterfactual; causal effect; causal inference; double robust estimating function; dynamic treatment regimen; estimating function; individualized stopped treatment regimen; individualized treatment rule; inverse probability of treatment weighted estimating functions; locally efficient estimation; static treatment intervention

*We thank James Robins for helpful discussions and suggestions.

1 Introduction

A wide range of statistical tools are available to estimate the causal effects of static treatment interventions. The identifiability of such effects relies on the assumption that treatment assignment in the observed data is not based deterministically on a subject's past (the assumption of experimental treatment assignment, or ETA). However, in practical applications the full set of treatment options is often not available to all individuals. Moreover, when the treatment of interest is assigned longitudinally over time, subjects for whom a given treatment regimen was initially possible may develop conditions that reduce their set of future treatment options. Common occurrences such as these result in violation of the ETA assumption, potentially causing considerable bias in estimators of the causal treatment effect.

In this article we introduce two new classes of causal models that address this pressing challenge to practical data analysis. These classes of causal models are indexed by the following two types of intervention: 1) realistic treatment rules; and 2) intention-to-treat interventions. By definition, both types of intervention assign treatment at each time point only from among those treatments which are possible given a subject's past. We further introduce two types of causal models indexed by realistic individualized treatment rules. The first type estimates the causal effects of a user-supplied set of individualized treatment rules (or dynamic treatment regimes), and can thus be used to identify the optimal rule from among this user-supplied set. The second type estimates the causal effects of realistic rules indexed by static treatment regimens; the parameter estimated is the causal effect of remaining on a static treatment regimen for only as long as such a static regimen is possible. Similarly, causal models for intention-to-treat interventions estimate the effect of remaining on a static treatment regimen only to the extent possible given a subject's covariates. As we explain below, the difference between intention-to-treat and realistic rules lies in the type of intervention that is assigned after the initial static regimen becomes impossible.

By defining a specific type of treatment intervention, the causal parameters indexed by both realistic treatment rules and intention-to-treat interventions avoid violation of the ETA assumption. As a result, the causal effects discussed in this article are fully identifiable based on the data. The article provides a detailed interpretation of the causal effects indexed these different types of interventions, using both the formal counterfactual framework and an example drawn from the treatment of HIV. The article further develops inverse probability weighted, likelihood-based, and double robust estimators of these causal effects. The estimation of causal effects indexed by a user-supplied set of realistic treatment rules is illustrated using a data analysis based on the HIV example.

This introduction begins by defining the data structure that underlies the causal parameters and corresponding estimators developed in the article. The causal effect of a static treatment and its relation to the ETA assumption is then reviewed. Next, we introduce our first class of causal models, indexed by realistic treatment rules. We then introduce the intention-to-treat intervention and corresponding causal effect. These three types of causal effect (ie. indexed by a static intervention, realistic rule, and intention-to-treat intervention) are compared using the HIV example. We provide a general road map to the rest of the article, and review the relevant literature.

1.1 Data structure

Consider a data generating experiment in which the experimental unit results in the following time-ordered sequential data structure

$$O = (L(0), A(0), L(1), A(1), \dots, L(T), A(T), L(T+1)),$$

where $A(j)$ denotes a treatment assignment at time j , $L(j)$ denotes all variables measured on the experimental unit after $A(j-1)$ and before $A(j)$, and $T+1$ is a fixed or random end-point such as a survival time. We assume that $T+1 \leq \tau+1$ with probability 1 for a fixed τ . Suppose we observe n independently and identically distributed copies O_1, \dots, O_n of O . For simplicity, throughout this article, we will treat all random variables as discrete, but all formulas have natural continuous analogues.

We let $R(t) \equiv I(T \leq t)$ be a component of $L(t)$, and we truncate the A and L process at T so that $A(t) = A(\min(t, T))$, $L(t) = L(\min(t, T+1))$. In this manner, we can now also represent the observed longitudinal data structure O on the experimental unit as a vector of fixed length,

$$O = (L(0), A(0), L(1), A(1), \dots, L(\tau), A(\tau), L(\tau+1)),$$

where we remind the reader that after time $T+1$ the data structure becomes degenerate in the sense that $A(T+j) = A(T)$, and $L(T+1+j) = L(T+1)$ for $j = 1, 2, \dots$

Let Y be a real-valued function of L , which will denote the outcome of interest. For example, $Y = T+1$ might be the survival time $T+1$, or it might be an outcome $Y(\tau+1)$ of a time-dependent process $Y(\cdot)$ measured at a fixed time $\tau+1$. We use the notation $\bar{L}(t) \equiv (L(0), \dots, L(t))$, but the complete covariate/outcome and treatment process are also denoted with $L = \bar{L}(\tau+1)$ and $A = \bar{A}(\tau)$.

The time-dependent treatment options process—Let $\mathcal{A}(t)$ be the support of the marginal random variable $\bar{A}(t) \equiv (A(0), \dots, A(t))$, $t = 0, \dots, \tau$. Let $\mathcal{D}(t)$ represent a set of possible treatment options for $A(t)$, given an experimental unit with history $\bar{A}(t-1)$, $\bar{L}(t)$, in the sense that

$$g_0(a(t) | \bar{L}(t), \bar{A}(t-1)) \equiv Pr(A(t) = a(t) | \bar{L}(t), \bar{A}(t-1)) > 0 \text{ for } a(t) \in \mathcal{D}(t).$$

It is assumed that $\mathcal{D}(t)$ is a function of $L(t)$: e.g., $\mathcal{D}(t)$ could be one of the components of $L(t)$.

1.2 The causal effect of a static treatment intervention

The current literature on causal inference provides models and corresponding methods for the estimation of static treatment effects. Typically, the data of interest are based on sampling subjects from a particular population and recording the treatment assignments, co-variables, and outcomes of interest for these subjects over time. Marginal structural models (MSM), introduced by Robins (e.g., Robins (1999), Robins (2000)), provide a powerful tool for causal inference in the context of such longitudinal data structures, and address many of the limitations of the traditional regression approach. MSM model the dependence of the distribution of treatment regimen-specific counterfactual outcomes (or outcome processes) on the treatment regimen. In other words, MSM model the population distribution of the outcome process that would be observed if all members of the population were to follow a particular static treatment regimen. The causal effect of a change in treatment is estimated as the difference in the population distribution of the outcome under the two treatment regimens being compared. For example, marginal structural models model as a function of $\bar{a}(t)$ the mean outcome under an

intervention setting $\bar{A}(t) = \bar{a}(t)$ with probability 1, possibly conditional on user-supplied baseline covariates. Inverse Probability of Treatment Weighted (IPTW) estimators, locally efficient double robust IPTW (DR-IPTW) estimators, and likelihood-based estimators have been proposed by Robins and co-authors for the unknown causal parameters in the marginal structural model. van der Laan and Robins (2003) provide a comprehensive overview of the development of these methods, together with a review of the relevant literature.

Identification of the causal effects of static treatment interventions based on observational data relies on the standard sequential randomization assumption (SRA) on the treatment mechanism, otherwise known as the assumption that there are no unmeasured confounders. In addition, because these methods aim to produce the results of a trial in which static treatment interventions are randomly assigned *and* each subject is forced to fully comply with the assigned intervention, they further require that treatment in the observed data not be deterministically assigned at any time point based on a subject's observed past; this latter assumption is called the experimental treatment assignment (ETA) assumption. In most studies all static treatment interventions cannot be followed by all sampled subjects, due either to baseline characteristics or to the occurrence of certain events over time. For example, clinical progression may force a clinician to initiate treatment before an assigned time, or the development of serious adverse effects or a contraindication, such as pregnancy, may force a clinician to stop an assigned treatment regimen. In practice, the ETA assumption can be somewhat weakened, to say that treatment cannot be deterministically assigned in response to that subset of a subject's observed past which is sufficient for the SRA to hold; thus the development of a condition that forces a subject to stop an assigned therapy need not lead to an ETA violation if it is causally unrelated to the outcome of interest.

Violation of the ETA assumption, whether theoretical (i.e. due to the occurrence of events which make a given treatment impossible) or practical (i.e. due to finite sample size), is known to result in potentially extreme bias in the IPTW estimators of marginal structural model parameters. Further, when ETA is violated, both likelihood-based estimators and DR-IPTW estimators rely fully on model assumptions (beyond the SRA) which cannot be tested from the data (Neugebauer and van der Laan (2005b)). The use of stabilizing weights (Robins (1999)) can somewhat mitigate the problem of ETA violations, by weakening the assumption from the requirement that

$$\inf_{\bar{a} \in \mathcal{A}} g_0(a(t) | \bar{A}(t-1), \bar{L}(t)) > 0, t=0, \dots, \tau$$

to the requirement that

$$\sup_{\bar{a} \in \mathcal{A}} \frac{g(a(t) | V)}{g(a(t) | \bar{A}(t-1), \bar{L}(t))} < \infty, t=0, \dots, \tau,$$

where $g_0(a(t) | \bar{A}(t-1), \bar{L}) = P(A(t) = a(t) | \bar{A}(t-1), \bar{L}(t))$ is the true treatment mechanism, and V denotes any baseline covariates of interest included in the marginal structural model. In many applications, however, violation of the ETA assumption remains a significant problem.

In addition to the potentially serious bias in effect estimates caused by ETA violations, the fact that models for static treatment interventions and their corresponding estimates aim to reproduce the results of typically unrealistic randomized trials has been a source of

philosophical criticism. Finally, even in settings where ETA violations of static treatment interventions do not occur, it is frequently of interest to estimate the causal effect of a dynamic rather than static treatment regimen. In other words, the researcher may be interested in the difference in mean outcome that would be observed between individuals randomized to receive different treatment *strategies* (or individualized treatment rules, in which treatment is changed in response to patient evolution), rather than different static regimens. In this article we propose a set of causal models and corresponding estimators that address these challenges.

1.3 Realistic individualized treatment rules

Lack of identifiability of the counterfactual distribution of the data under a static treatment intervention results from a non-zero probability of sampling an experimental unit for which the static intervention cannot occur. In this article we define realistic individualized treatment rules as rules which always assign treatments (in response to observed history) that fall in the set of possible treatment options. Causal effects comparing realistic individualized treatment rules are now identifiable from the observed data distribution.

Such a class of realistic individualized treatment rules can be user-defined. This approach is comparable to designing a randomized trial to compare realistic individualized treatment rules of interest. For example, the researcher may specify *a priori* a set of interesting rules which assign treatment over time by responding to patient covariates in a clinically realistic way, thus avoiding violations of the ETA assumption. Modeling the mean outcome for such a user-supplied class of realistic individualized treatment rules also yields a model for the optimal rule among the user-supplied class (i.e. by selecting the rule which maximizes or minimizes the mean outcome). The causal models for realistic treatment rules presented in this paper thus provide an alternative method for modelling and estimation of optimal dynamic treatment regimes from among a user-supplied set, based on a generalization of structural nested models (Robins (1989), Robins (1998), Robins (1999), Robins (1994)), as developed in (Murphy (2003), Robins (2003)).

It may not always be straightforward to propose such an interesting set of realistic individualized treatment rules. For example, the set of possible treatment options, given a subject's covariates ($\mathcal{D}(t)$) may not be known or collected as part of the study. In such a case, then we propose to define the set of possible treatment rules by employing a stabilized version of the treatment mechanism:

$$\mathcal{D}(t) \equiv \{a(t) : \frac{g_0(a(t) | V)}{g_0(a(t) | \bar{A}(t-1), \bar{L}(t))} < \alpha < \inf_{t=0, \dots, \tau} \}$$

for some α , and if the treatment mechanism g_0 is unknown, then one estimates this set by substitution of an estimator g_n of g_0 .

In addition, it may be the case that the researcher is in fact interested in the effect of a static (rather than dynamic) treatment regimen, *to the extent* that subjects are realistically able to follow it. One option is to map a static treatment intervention into a corresponding realistic individualized treatment rule, in which the subject follows the assigned static treatment intervention for only as long as the assigned intervention remains possible. When the assigned regimen no longer falls within the set of possible treatment options for the subject (as a result of the subject's covariate history), then the subject is assigned a particular alternative treatment in the set of remaining treatment options (e.g., the one "closest" to the treatment assigned by the static intervention). Again, this new static regimen is applied until the subject is forced to

switch again, and so on. In this manner, these individualized treatment rules are indexed by static treatment regimens, and provide natural realistic approximations of the intended static treatment intervention. Such rules correspond to randomized controlled trials that incorporate explicit rules for altering patients' randomized static treatment assignment in order, for example, to protect patient health.

Note that if the ETA assumption holds, so that the set of possible treatment options at time t can be chosen to be equal to the set of all marginally possible treatments at time t , then the realistic individualized treatment rules indexed by static interventions reduce to the class of static treatment interventions, and thus the causal model for such realistic rules reduces to a standard marginal structural model for static treatments. As a consequence, the realistic causal models introduced in this article generalize causal effect models for static treatment interventions (MSM's) that rely on the ETA assumption to causal effect models for individualized treatment rules, indexed by static interventions, that also apply when the ETA assumption is violated.

1.4 Intention-to-treat interventions

The previous section introduced a new class of causal models for realistic individualized treatment rules. These rules assign a treatment at each time point, based on a subject's current covariates, that always falls within the set of possible treatment options for that subject; the rules may be, but are not necessarily, indexed by static treatment regimens. In this section we introduce a related class of causal models that estimate the effects of “intention-to-treat” interventions indexed by static treatment regimens. As with realistic treatment rules indexed by static regimens, intention-to-treat interventions enforce the assigned static intervention up till the time point t at which the next prescribed treatment does not fall in the set of possible treatment options $\mathcal{D}(t + 1)$. If a realistic treatment rule were being applied, at this point a pre-specified alternative regimen would be assigned. In contrast, in an intention-to-treat intervention, when the assigned regimen is no longer possible then all intervention on that subject is stopped.

The name “intention-to-treat” is based on the loose association of this type of intervention with the practice of analyzing the results of randomized controlled trials based on treatment assignment, rather than on the treatment that subjects were observed to follow. In other words, even though in practice some individuals may find it impossible to follow the treatment arm to which they are randomly assigned, conventional clinical trial analysis treats these individuals as if they had followed their assigned treatment. A key difference between this classic definition of “intention-to-treat” and the intention-to-treat parameter described in this paper is that the latter assumes that the set of possible treatment options for a subject is fully determined given a subject's past, whereas in the context of a controlled trial, failure to adhere to assigned treatment may occur even in the absence of conditions which make the assigned treatment impossible.

The causal effects of these intention-to-treat interventions indexed by static treatment regimens are now fully identifiable from the data. As a consequence, as with realistic treatment rules, we can develop locally efficient estimators of these causal effects without the need to assume the often unrealistic ETA assumption. As with realistic treatment rules, this model for intention-to-treat interventions indexed by static treatments generalizes marginal structural models for static interventions, which depend on the ETA assumption, to causal models for corresponding intention-to-treat interventions that reduce to marginal structural models if the ETA assumption holds.

1.5 Example

Suppose that we sample subjects from an HIV-infected population receiving antiretroviral therapy, who at time 0 experience a persistent rebound in plasma HIV RNA level as a result of viral resistance to their prescribed drug regimens. Suppose that $Y(8)$ is the CD4 T cell count measured 8 months after rebound and that the measurements $L(t)$, $t = 0, \dots, 8$ include plasma HIV RNA level (viral load), CD4 T cell count, and other time-dependent characteristics of interest. Let $(A(t), t = 0, \dots, 7)$ be the indicator process which equals 1 up till the time point at which a subject modifies his or her antiretroviral regimen, and then jumps to 0. One might now be interested in estimation of the causal effect of time until treatment modification on CD4 T cell count at 8 months, based on this sample of patients experiencing a rebound of the virus at time 0. Specifically, we refer to Petersen et al. (2005) for a description of the SCOPE cohort of HIV-infected patients, and of the interest and relevance of the “when to switch?” question in the HIV-AIDS research community, with relevant references. In particular, it has been observed that a drug can still have a significant beneficial effect on a resistant virus by making it less lethal and/or fit, so that an increase in viral load does not necessarily imply a decrease in CD4 T cell count.

Consideration of hypothetical randomized trials can be used to illustrate how the three types of causal parameters described in this article compare with the effect of a static treatment regimen as typically estimated using marginal structural models.

Estimating the effect of a static modification time—First, consider a trial in which each subject is randomly assigned a time at which to modify treatment. Note that this is a randomized trial of a static intervention, in that subjects are assigned a modification time at baseline and this time is not altered in response to the condition of the patient or virus. In order to estimate the mean outcome in the trial arm assigned to modify treatment at time t , one would need that every subject in this group (or at least a representative subgroup of these subjects) is able to fully comply with the assigned modification time t . However, suppose that some subjects in the population develop an opportunistic infection or adverse drug effects before time t , and as a result are unable to remain on their baseline regimen. Such patients cannot comply with the assigned modification time. Assuming that the measurement of such deterministic events is necessary to ensure sequential randomization, or put another way, assuming that individuals who are unable to comply with their assigned modification time t are not representative of all subjects assigned to modify at time t with respect to their counterfactual outcome, the causal effect of time until treatment modification is not identifiable from the data. As a consequence, any of the candidate estimators of a marginal structural model modeling the causal effect of time until treatment modification on mean CD4 T cell count at 8 months suffer from potentially serious bias. Further, one might argue that such a causal effect is not in fact of primary interest, as it corresponds with an intervention that would never be employed in practice.

Estimating the effect of a realistic rule/intention-to-treat intervention indexed by a static modification time—Alternatively, consider a trial in which each subject is assigned a realistic treatment rule indexed by a random modification time t . A subject in such a trial remains on his or her original therapy until the minimum of time t and the first time point at which an event occurs which forces him or her to modify therapy. Under ETA (i.e. the absence of deterministic events that force treatment modification), these realistic treatment rules are identical to the static modification regimens described in the prior example. However, unlike the effects of the static modification regimens, the causal effects of the realistic rules are still interpretable and identifiable if ETA is violated. In this case, the realistic treatment rule indexed by the static modification time t is identical to the “intention-to-modify at time t ” intervention, due to the fact that the treatment process $A(t)$ is binary.

Estimating the effect of a user-supplied set of realistic rules for modifying therapy—The prior example describes estimation of the effect of realistic treatment rules indexed by *static* modification times. In clinical practice, however, the effects of a set of dynamic strategies, which assign a patient a rule for when to modify therapy based on (e.g.) the evolution of CD4 T cell count over time, are likely to be of greater clinical interest than the effects of set modification times assigned at baseline (even if such static modification times are allowed to be changed if they become impossible). With this motivation, consider a third clinical trial, in which subjects are assigned to modify treatment only when their CD4 T cell count falls below a certain randomly assigned threshold θ . Again, such a user-supplied set of treatment rules should be realistic. This implies that either any patient characteristics that make the assigned threshold impossible to comply with must not be required for the SRA to hold, or such covariates should be incorporated into the treatment rules considered. For example, we could now define the following user-supplied set of realistic individualized treatment rules indexed by a threshold θ : “modify therapy if CD4 T cell count drops below θ or if an event occurs that forces modification”.

1.6 Organization

This article introduces causal models for realistic individualized treatment rules and intention-to-treat interventions, and develops corresponding estimators. In Section 2 the causal model for realistic (and thereby identifiable) individualized treatment rules is presented, and the corresponding locally efficient double robust inverse probability of treatment weighted estimator is derived. This model and methodology for the special (but common) case in which the treatment is assigned at a single point in time in response to baseline covariates is provided in Appendix A. In Section 3 we present a data analysis, based on the example presented above, that illustrates this model and methodology. In the analysis we estimate the causal effect of a user-supplied set of realistic rules for modifying antiretroviral therapy, based on a subject's current and baseline CD4 T cell counts.

The remainder of the article focuses on causal effect models for intention-to-treat regimens. Specifically, in Section 4 we define a causal inference framework which allows us to define the causal effects of a range of interventions on the data generating distribution of the data structure O , and, in particular, allows us to define our wished non-parametric identifiable intention-to-treat causal effect parameter. This framework represents a set of assumptions which do not put any restrictions on the data generating distribution, but are essential for the definition and identification of the wished causal effect of an intention-to-treat regimen from the data generating distribution. Given the causal inference framework, we define the intention-to-treat counterfactual processes, and corresponding models for the conditional mean of the intention-to-treat counterfactual outcome.

In Section 5 we present the intention-to-treat causal effect model for the point treatment data structure ($W = L(0), A, Y$). We further present the corresponding efficient Double Robust Inverse Probability of Treatment Weighted (DR-IPTW) estimating function and the locally efficient double robust estimator. The latter estimator (derived in Appendix B) is locally efficient in the sense that its consistency (and asymptotic linearity) relies on either correct specification of the treatment mechanism $P(A = a | W)$ or the regression $E(Y | A, W)$, and it is efficient if both are correctly specified. We also present the likelihood-based estimator and the simpler IPTW estimator, which is a special case of the DR-IPTW estimator. The presentation of the estimating function-based estimators of the intention-to-treat causal parameter for the general longitudinal data structure is deferred to Appendices C-D, since this work happens to be quite involved.

Section 6 is devoted to a discussion.

1.7 Some immediately relevant literature

Based on personal communication we became aware of completely independent unpublished work on individualized treatment rules by Andrea Rotnitzky, who presented similar/overlapping ideas in a November 2005 NIH grant proposal titled “Methods for Analysis with Missing and/or Censored Data and for Causal Inference”. Specifically, in the aforementioned grant, Rotnitzky proposed the consideration of parametric and semiparametric models for the marginal means of counterfactual variables under a class of dynamic treatment regimes as a device for estimating the optimal treatment regime in the class. She also indicated how doubly-robust locally efficient Inverse Probability Weighted estimators of the model parameters can be constructed. These estimators essentially agree with those derived at the end of Section 2 of the present paper. Finally, she also noted that an important application of the methodology was to estimate the optimal CD4 T cell count level at which to either start or switch antiretroviral therapies, much along the lines of the example presented in Section 3 here. We were not aware of this work, but certainly wish to acknowledge her independent and overlapping original work in this area of research.

The method of inverse probability of treatment weighting in order to compare dynamic treatment regimens and their corresponding double robust estimators was presented in Murphy et al. (2001), and Hernan et al. (2006), and, as a method, goes back to the general Inverse Probability of Censoring Weighting as presented in Robins and Rotnitzky (1992) and Robins (1993). Murphy et al. (2001) proposes a model for a single dynamic treatment regimen conditional on baseline covariates, while the current article generalizes that to a model for a user-supplied class of dynamic treatment regimens. The introduction of dynamic treatment regimens, and the notion and idea of (what we call) realistic individualized treatment rules goes back to Robins (1986) who discusses such regimens in the context of a study aiming to estimate the effects of exposure of chemicals on employees, noting that static regimens cannot be identified since the subjects can only be exposed if at work. Finally, we also wish to point out the relation between the statistical framework/assumptions as we presented for defining and identifying the causal effects of the intention-to-treat regimens and the work in Robins (1986) (specifically, pages 1422-1423 in Robins (1986)).

2 Causal effect models for realistic individualized treatment rules

In this section we present a causal effect model for realistic individualized treatment interventions, and present the corresponding locally efficient double robust inverse probability of treatment weighted estimator following the general estimating function methodology of Robins and Rotnitzky (1992) and van der Laan and Robins (2003).

2.1 The counterfactual framework for realistic individualized treatment rules

We use the statistical framework of counterfactuals on which marginal structural models are based. This framework was introduced in Neyman (1990), extended to causal effects of time-independent treatments by Rubin (1978), and further extended to a formal theory of causal inference for direct and indirect effects of time-varying treatments from experimental and observational longitudinal studies by Robins (1986, 1987). This causal framework for treatment interventions \bar{a} assumes the existence of counterfactuals indexed by static treatment interventions \bar{a} , the corresponding link between the observed data and these counterfactuals (i.e., consistency assumption), and the sequential randomization assumption (SRA). By applying the result in (Gill and Robins (2001), Yu and van der Laan (2002)), it follows that, by construction, assuming the consistency and randomization assumptions puts no restriction on the data generating distribution. These assumptions do, however, allow us to define the causal parameter of interest as a parameter of the data generating distribution.

As defined in Robins (1986), an individualized treatment rule d is a function $(d(0), \dots, d(\tau))$, where the j -th function, $(\bar{A}(j-1), \bar{L}(j)) \rightarrow d(j)(\bar{A}(j-1), \bar{L}(j))$, maps the history at time j into a treatment choice for $A(j)$, $j = 0, \dots, \tau$.

Consistency assumption—We define the full data as the collection $X = (L_{\bar{a}} : \bar{a} \in \mathcal{A})$ of counterfactual processes $L_{\bar{a}}$ indexed by static treatment interventions varying over the support of the marginal distribution of $\bar{A} = (A(0), \dots, A(\tau))$. We also assume the temporal ordering assumption, $L_{\bar{a}}(j) = L_{\bar{a}(j-1)}(j)$, and the consistency assumption stating that $O = (\bar{A}, L_{\bar{A}})$ or, as a chronological data structure:

$$O = (L(0), A(0), L_{A(0)}(1), A(1), \dots, L_{\bar{A}(\tau-1)}(\tau), A(\tau), L_{\bar{A}(\tau)}(\tau+1)).$$

Dynamic treatment counterfactuals—Given this standard consistency assumption, for any rule d , the counterfactual L_d indexed by a dynamic treatment d can be defined as $L_{\bar{a}}$ with $\bar{a} = \bar{a}(X, d)$ defined as the following function of X and the rule d : $a(0) = d(0)(L(0))$, $a(1) = d(1)(a(0), L_{a(0)}(1))$, and, in general, $a(j) = d(j)(\bar{a}(j-1), L_{\bar{a}(j-1)}(j))$, $j = 0, \dots, \tau$. Thus, given the existence of the random variable X defined as the collection of static treatment-specific counterfactuals, one can also define the dynamic treatment regimen-specific counterfactuals $L_d \equiv L_{\bar{a}(X,d)}$ as a measurable function of X and the rule d . We recall that the treatment options process $t \rightarrow \mathcal{D}(t)$ is included in the $t \rightarrow L(t)$ process, so that \mathcal{D}_d denotes the d -specific treatment options process for the experimental unit.

It is also of interest to note that, for each experimental unit, the rule d maps into a unique treatment regimen $\bar{a}(d, X)$. However, a static treatment intervention \bar{a} can correspond with various individualized treatment rules d : e.g. $L_{\bar{a}} = L_{d_1} = L_{d_2}$ for two different rules d_1 and d_2 in a set of dynamic regimens \mathcal{A}^* . If an experimental unit follows rule d starting at time 0, then it follows that d_j is, in fact, only a function of $\bar{L}_d(j)$. For the sake of notational convenience, in that case we will use the notation $\bar{L}_d(j) \rightarrow d(j) \bar{L}_d(j)$.

Sequential randomization assumption—We will assume the (strong) sequential randomization assumption: i.e., for each $j = 0, \dots, \tau$, $A(j)$ is independent of X , given $\bar{L}(j)$, $\bar{A}(j-1)$. The data generating distribution of O will be denoted with $P_0 = P_{F_{X0}, g_0}$, and is indexed by the distribution F_{X0} of X and the conditional probability distribution, $g_0(\cdot | X)$, of \bar{A} , given X .

Realistic dynamic treatment assumption—Let \mathcal{A}^* be a set of dynamic treatment regimens so that for any $d \in \mathcal{A}^*$ we have

$$P(d(j) \bar{L}_d(j) \in \mathcal{D}_d(j), j=0, \dots, \tau) = 1. \tag{1}$$

That is, for each possible history at time j under a dynamic treatment regimen $d \in \mathcal{A}^*$, the next treatment assigned by this individualized treatment rule d at time $j+1$ is an element of the set $\mathcal{D}_d(j+1)$ of possible treatment options. This condition on the rule d guarantees that the distribution of L_d is identifiable by the G -computation formula (Robins (1999), Gill and Robins (2001), Yu and van der Laan (2002)):

$$P(L_d=l) = \prod_{j=0}^{\tau+1} P(L(j)=l(j) | \bar{A}(j-1)=\bar{d}(j)(l), \bar{L}(j-1)=\bar{l}(j-1)).$$

where we defined $\bar{d}(j)(l) \equiv (d(1)(l(0)), \dots, d(j-1)(l(j-1)))$.

Realistic individualized treatment rules indexed by static treatment regimens—

Let $C_{\bar{a}}$ be the counterfactual stopping time defined as

$$C_{\bar{a}} \equiv \min \{t \in \{-1, 0, \dots, \tau\} : a(t+1) \notin \mathcal{D}_{\bar{a}}(t+1) \text{ or } t=\tau\}.$$

Given a static treatment regimen \bar{a} , one can define a dynamic treatment regimen as one which follows the static treatment regimen \bar{a} until time point $t = C_{\bar{a}}$ at which $a(t+1) \notin \mathcal{D}(t+1)$ or $t = \tau$, and subsequently one proceeds assigning treatments in the set of treatment options according to a particular user-supplied rule.

For example, the following construction describes such a set of dynamic treatment regimens indexed by static treatment interventions \bar{a} . Suppose that the maximal set of treatment options is \mathcal{S} in the sense that $\mathcal{D}(j) \subset \mathcal{S}$ for all $j = 0, \dots, \tau$, with probability 1. In addition, define a dissimilarity measure between any pair of elements in \mathcal{S} so that for each $s \in \mathcal{S}$, we can identify the element in $\mathcal{D}(j)$ closest to s . We could now define the following individualized treatment rule indexed by a static treatment regimen \bar{a} : 1) follow static treatment regimen \bar{a} until time point $t = C_{\bar{a}}$ at which $a(t+1) \notin \mathcal{D}(t+1)$ or $t = \tau$; 2) if $t < \tau$ (that is, it was not possible to fully comply with the static regimen \bar{a}), then set the next treatment equal to the element in $\mathcal{D}(t+1)$ closest to $a(t+1)$; 3) keep this treatment constant until the time point at which the treatment is not an element of the set of treatment options so that a switch of treatment is required, or until the endpoint τ ; 4) if the treatment needs to be switched before τ , then switch again to the element in the set of treatment options closest to the current treatment; 4) continue in this manner until one reaches the end point τ . Notice that this defines an individualized treatment rule as a deterministic function of a static intervention \bar{a} . Therefore, we can denote this set of dynamic treatment rules with $d_{\bar{a}}, \bar{a} \in \mathcal{A}$.

2.2 Causal effects of realistic individualized treatment rules

The above standard causal inference assumptions put no restrictions on the data generating distribution and thereby cannot be tested based on the data. In particular, the model for the distribution of the data implied by the above assumptions is still unspecified/nonparametric.

We define the parameter of interest on this nonparametric model as the conditional mean of Y_d , given a subset V of the baseline covariates $L(0)$, for all $d \in \mathcal{A}^*$. In order to deal with the curse of dimensionality, one can choose between two approaches. Firstly, one can assume a model

$$E_0(Y_d|V) = m(d, V|\beta_0) \tag{2}$$

for some parametrization $(d, V) \rightarrow m(d, V|\beta)$ indexed by a finite dimensional Euclidean parameter β . In this model $\beta(F_X)$ is the parameter of interest, and $\beta_0 = \beta(F_{X0})$ is the true value of this parameter. For example, if $d = d_{\bar{a}}$ is a deterministic function of a static treatment intervention, as in our example above, then we would have

$$E_0(Y_{d_{\bar{a}}}|V) = m(\bar{a}, V|\beta_0).$$

Alternatively, if one believes that the model assumed on $E(Y_d | V)$ is not realistic, then it might be sensible to define the parameter of interest as

$$\beta_h(F_X) \equiv \arg \min_{\beta} \sum_{d,V} (E_{F_X}(Y_d | V) - m(d,V|\beta))^2 h(d,V).$$

If model (2) holds, then $\beta_{0h} = \beta_0$ for all h . One can map $\beta(F_{X0})$ into a corresponding optimal individualized treatment rule within each strata V :

$$d(F_{X0})(V) = \arg \max_{d \in \mathcal{A}^*} m(d,V|\beta_0).$$

Note that the parameters $\beta(F_X)$ and $\beta_h(F_X)$ are parameters of F_X . As a consequence, we can apply the general estimating function methodology of Robins and Rotnitzky (1992) and van der Laan and Robins (2003), to obtain the class of all estimating functions, including the optimal DR-IPTW estimating function, which equals the efficient influence curve when evaluated at the true parameter values. This general estimating function methodology involves three steps: 1) identify the class of all full data estimating functions (formally, the space spanned by the gradients of the path-wise derivative of the parameter of interest, also called the orthogonal complement of the nuisance tangent space); 2) construct an inverse probability of treatment weighted class of estimating functions which are such that the conditional expectation, given X , maps into the class of full data estimating functions; and, 3) map this class of IPTW estimating functions into the double robust IPTW estimating functions by subtracting the projection on the tangent space spanned by all scores of the treatment mechanism under the sole model assumption SRA. For details, we refer to the original paper Robins and Rotnitzky (1992) which laid out this general approach for censored data models and to Chapters 1 and 2 of van der Laan and Robins (2003).

Firstly, we need to determine the class of full data estimating functions one would obtain in the full data model for X . It follows that this class of full data estimating functions is given by:

$$\left\{ \sum_{d \in \mathcal{A}^*} h(d,V) \frac{d}{d\beta_0} m(d,V|\beta_0) (Y_d - m(d,V|\beta_0)) : h \right\}.$$

In the case that one defines the parameter of interest as

$$\beta_h(F_X) \equiv \arg \min_{\beta} \sum_{d,V} (E_{F_X}(Y_d|V) - m(d,V|\beta))^2 h(d,V),$$

then the only full data estimating function is

$$\sum_{d \in \mathcal{A}^*} h(d,V) \frac{d}{d\beta_0} m(d,V|\beta_0) (Y_d - m(d,V|\beta_0)).$$

We now need to find an IPTW estimating function which has the property that its conditional expectation, given X , maps into the class of full data estimating functions. As established in the following result, we can use

$$D_{h,IPTW}(O|g_0, \beta_0) = \sum_{d \in \mathcal{A}^*} \frac{I(\bar{a}=d(\bar{L}))}{g(\bar{a}|X)} h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Y - m(d, V|\beta_0)),$$

where $\bar{a} = d(\bar{L})$ is defined as

$$(a(0) = d_0(L(0)), a(1) = d(1)(a(0), \bar{L}(1)), \dots, a(\tau) = (\bar{a}(\tau - 1), \bar{L}(\tau))).$$

Result 1 Assume that for all individualized treatment rules $d \in \mathcal{A}^*$, we have

$$Pr(g(\bar{a}(X, d) | X) > 0) = 1,$$

where $\bar{a}(X, d)$ is the treatment regimen followed by the experimental unit with full data counterfactuals X if the experimental unit follows rule d : $a(0) = d(0)(L(0))$, $a(1) = d(1)(\bar{L}_{a(0)}(1))$, and, in general, $a(j) = d(j)(\bar{L}_{\bar{a}(j-1)}(j))$, $j = 0, \dots, \tau$.

We have for all h

$$E(D_{h,IPTW}(O|g_0, \beta_0) | X) = \sum_{d \in \mathcal{A}^*} h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Y_d - m(d, V|\beta_0)).$$

As a consequence, if $E(Y_d | V) = m(d, V | \beta_0)$, then

$$E_0 D_{h,IPTW}(O|g_0, \beta_0) = 0 \text{ for all } h,$$

and, we always have for all h

$$E_0 D_{h,IPTW}(O|g_0, \beta_{h0}) = 0.$$

Proof. Because $g(\bar{a}(d, X) | X) > 0$, the conditional expectation $E(D_{h,IPTW}(g_0, \beta_0) | X)$ equals

$$\begin{aligned} & \sum_{\bar{a} \in \mathcal{A}} \sum_{d \in \mathcal{A}^*} I(\bar{a} = d(\bar{L}_{\bar{a}})) h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Y_{\bar{a}} - m(d, V|\beta_0)) \\ & = \sum_{d \in \mathcal{A}^*} \sum_{\bar{a} \in \mathcal{A}} I(\bar{a} = d(\bar{L}_{\bar{a}})) h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Y_d - m(d, V|\beta_0)). \end{aligned}$$

Now, we note that $\bar{a} = d(\bar{L}_{\bar{a}})$ is equivalent with the unique solution $a(0) = d(L(0))$, $a(j) = d(\bar{L}_{\bar{a}(j-1)}(j))(j)$, $j = 1, \dots, \tau$. Thus, the inner $\sum_{\bar{a} \in \mathcal{A}}$ reduces to the single term $h(d, V)d/d\beta_0 m(d, V | \beta_0)(Y_d - m(d, V | \beta_0))$, so that the conditional expectation reduces to

$$\sum_{d \in \mathcal{A}^*} h(d, V) \frac{d}{d\beta_0} m(d, V | \beta_0)(Y_d - m(d, V | \beta_0)),$$

which completes the proof.

Finally, we map this IPTW estimating function for β_h into the efficient estimating function by subtracting its projection on the tangent space of the treatment mechanism under SRA. The following result describes this double robust IPTW estimating function, and thereby the efficient influence curve. The proof of this result is a direct consequence of Theorem 1.3 and Theorem 1.6 in van der Laan and Robins (2003).

Result 2 The efficient influence curve of β_h in the (nonparametric) model for the data generating distribution P_0 at P_0 is given by $-c(\beta_{h0})^{-1} D_{h,DR}(O | g_0, Q_0, \beta_{h0})$, where

$$D_{h,DR}(O | g_0, Q_0, \beta_{h0}) = D_{h,IPTW}(O | g_0, \beta_0) - \sum_{t=0}^{\tau} \{E_{g_0, Q_0}(D_{h,IPTW}(g_0, \beta_0) | \bar{A}(t), \bar{L}(t)) - E_{g_0, Q_0}(D_{h,IPTW}(g_0, \beta_0) | \bar{A}(t-1), \bar{L}(t))\}$$

and $c(\beta) \equiv \frac{d}{d\beta} E_0 D_{h,DR}(O | g_0, Q_0, \beta)$. If $E_0(Y_d | V) = m(d, V | \beta_0)$, then for all h

$$E_0 D_{h,DR}(O | g, Q, \beta_0) = 0 \text{ if } g = g_0 \text{ or } Q = Q_0.$$

In general, for all h ,

$$E_0 D_{h,DR}(O | g, Q, \beta_{h0}) = 0 \text{ if } g = g_0 \text{ or } Q = Q_0.$$

Inverse probability of treatment weighted and double robust locally efficient estimators—Given an estimator g_n of the treatment mechanism g_0 , and a possibly data dependent index h_n , we define the IPTW estimator as the solution $\beta_{h_n, IPTW}$ of

$$0 = \sum_{i=1}^n D_{h_n, IPTW}(O_i | g_n, \beta).$$

Similarly, given an estimator (g_n, Q_n) of the nuisance parameter (g_0, Q_0) , and a possibly data dependent index h_n , we define the double robust locally efficient estimator as the solution $\beta_{h_n, DR}$ of

$$0 = \sum_{i=1}^n D_{h_n, DR}(O_i | g_n, Q_n, \beta).$$

Under regularity conditions, the estimator $\beta_{h_n, DR}$ is consistent and asymptotically linear if either g_n converges to g_0 or Q_n converges to Q_0 , and, if both nuisance parameters are consistently estimated, then $\beta_{h_n, DR}$ is an asymptotically efficient estimator of β_{h_0} . Therefore we call such an estimator $\beta_{h_n, DR}$ locally efficient. In contrast, the consistency of $\beta_{h_n, IPTW}$ requires that g_n converges to g_0 . For the formal statement for the asymptotics of the double robust estimator with the required regularity conditions, we refer to Theorem 2.4 and 2.5 in van der Laan and Robins (2003).

In order to avoid technicalities, for statistical inference we propose the bootstrap method which is known to be asymptotically valid under the same conditions required to establish the asymptotic linearity of the estimators $\beta_{h_n, IPTW}$ and $\beta_{h_n, DR}$.

3 Data example: Realistic individualized rules for treatment modification in HIV-infected patients experiencing viral rebound

In this section we present the results of a data analysis based on the example described in subsection (1.5). Specifically, data drawn from the Study of the Consequences of the Protease Inhibitor Era (SCOPE) were used to estimate counterfactual mean CD4 T cell count 8 months after confirmed virologic rebound on an antiretroviral treatment regimen, under a set of user-supplied realistic treatment rules. Specifically, we considered rules indexed by a range of CD4 T cell count thresholds θ . Thus, we aimed to replicate the results of a clinical trial in which subjects were assigned to modify their initial failing regimen only when CD4 T cell count reached a randomly assigned threshold.

Recall that \bar{A} in this example is an indicator process that jumps only once, when a subject modifies therapy for the first time; thus, treatment decisions are made only for those individuals who have not yet modified their original antiretroviral regimen. We defined the following set of realistic individualized treatment rules d_θ , indexed by threshold CD4 T cell count θ :

$$\begin{aligned} d_\theta(t) &\equiv I(CD4(t) \leq \theta) \\ d_\theta(t) &= 1: \text{Modify therapy at time } t \\ d_\theta(t) &= 0: \text{Do not modify therapy at time } t. \end{aligned}$$

This set of treatment rules uses a subject's current CD4 T cell count ($CD4(t)$) to assign a treatment decision $d_\theta(t)$ (modify treatment or not) at each time point t from baseline until a subject modifies therapy. Note that for a given subject, d_θ applied from baseline onwards deterministically implies a static treatment regimen \bar{a} .

By calling this set of rules d_θ “realistic”, we imply that, given the subset of a subject's covariates necessary for the SRA to hold, the rule d_θ does not assign, at any time point, a treatment action that is impossible for that subject. Below, we discuss whether this realistic dynamic treatment assumption is reasonable for the set of user-supplied rules considered in the current data example.

We estimated the following parameter by assuming a model on the counterfactual mean future CD4 T cell count under a realistic treatment rule indexed by d_θ :

$$E(Y_{d_\theta})=m(\theta|\beta).$$

Thus β was our parameter of interest, providing a summary of how the counterfactual outcome varied depending on the threshold CD4 T cell count at which a subject modified his or her virologically failing therapy. In particular, estimation of β implied estimation of an optimal threshold θ at which to modify the failing antiretroviral regimen (i.e. the optimal threshold was defined as the θ which maximized the expected counterfactual CD4 T cell count 8 months later).

We further modelled this counterfactual dynamic mean conditional on baseline CD4 T cell count. In other words, we estimated the following causal parameter:

$$E(Y_{d_\theta}|V)=m(\theta,V|\beta).$$

where $V \equiv CD4(0)$ denotes CD4 T cell count at time of confirmed virologic failure (baseline). Estimation of this parameter allowed us to address the question of whether the optimal threshold at which to modify therapy varied depending on a subject's baseline CD4 T cell count.

In defining θ and V , CD4 T cell count was categorized using 23 discrete levels (corresponding to cell counts of 0-50 cells, 51-100 cells, etc...).

3.1 IPTW estimation

We assumed the following models on the causal parameter of interest.

$$m(\theta|\beta)=\beta_0+\beta_1\theta+\beta_2\theta^2,$$

and

$$m(\theta,V|\beta)=\beta_0+\beta_1\theta+\beta_2\theta^2+\beta_3\theta \times V+\beta_4\theta^2 \times V+\beta_5V.$$

As noted above, if one is uncomfortable with assuming these models, the causal parameter of interest could alternatively be viewed as the projection of the true causal parameters onto working models $m(\theta|\beta)$ and $m(\theta, V|\beta)$.

The parameter β was estimated using the IPTW estimating function

$$D_{h,IPTW}(O|g_0,\beta_0) = \sum_{d_\theta} \frac{I(A=d_\theta(CD4))}{g_0(A|X)} h(d_\theta, V) \frac{d}{d\beta_0} m(\theta, V|\beta_0)(Y - m(\theta, V|\beta_0)),$$

where we used $h(d_\theta, V) \equiv 1$. Under this estimating function, β was estimated using weighted least squares regression, with each subject contributing one line of data for each threshold θ consistent with the subject's observed treatment history, and using the following weights:

$$\text{weight} = \frac{I(\bar{A} = d_\theta(C\bar{D}4))}{g(\bar{A}|X)}$$

Estimation of β thus required, for each possible threshold θ , determination of whether each subject complied with the treatment rule implied by that threshold ($I(\bar{A} = d_\theta(C\bar{D}4))$). This determination was made according to the following algorithm:

- Among subjects who modified therapy at some time M before 8 months, we evaluated whether $CD4(t) > CD4(M)$ for $t = 0, \dots, M - 1$. If yes, the subject was considered to have followed a treatment rule corresponding to $\theta = CD4(M)$. If no, the subject failed to follow a rule for treatment modification based on any CD4 T cell count threshold, and was assigned no value for θ .
- Subjects who did not modify therapy prior to measurement of the outcome were considered to have followed multiple treatment rules, corresponding to each threshold CD4 T cell category below the minimum CD4 T cell count category observed over the course of follow-up.

In implementing the IPTW estimator, the treatment mechanism $g(\bar{A}|X)$ was estimated using logistic regression of the probability of switching therapy at each time point given a subject's observed past. The regression fits were estimated data-adaptively using the Deletion/Substitution/Addition algorithm of Sinisi and van der Laan (2004), and 5-fold cross validation. Potentially informative censoring was addressed by modeling the censoring mechanism and employing inverse probability of censoring weights (as outlined in van der Laan and Robins (2003)). Standard errors were estimated based on 100 bootstrap samples.

3.2 Results

One hundred thirty three subjects experienced 167 episodes of confirmed virologic failure during SCOPE follow-up; of these, 33 subjects were censored before the outcome at 8 months was obtained. For a full description of this sample, and the covariates measured on each subject over time, see (Petersen et al. (2005)). Of the 100 subjects who failed an antiretroviral therapy regimen and were not censored prior to measurement of the outcome, 56 subjects (57 episodes of failure) had observed treatment histories that corresponding with following the rule d_θ (as defined above) for at least one threshold θ .

The estimated treatment mechanism is reported in Table (1). The resulting weights ranged from 1.01 to 53; in order to reduce variability, weights were truncated at 10, which resulted in truncation for 8 of the subjects. Under the assumption that the treatment mechanism was consistently estimated (as required for the consistency of the IPTW estimator), we further considered the assumption that the rules d_θ considered are in fact realistic. As a means of partially examining the assumption that modification is not deterministically assigned given time-dependent covariates, we assessed the stabilizing weights that would have been employed in a standard marginal structural model estimating the effect of a static switch time; we

estimated $\sup_{\bar{a} \in \bar{A}} \frac{g(\bar{a}|V)}{g(\bar{a}|X)} = 1.2$ and $\sup_{\bar{a} \in \bar{A}} \frac{g(\bar{a}|V)}{g(\bar{a}|X)} = 7.1$. Given that a subject can only switch

regimens once, $\sup_{\bar{a} \in \bar{A}} \frac{g(\bar{a}|V)}{g(\bar{a}|X)}$ was estimated as less than or equal to $1.2^7 \times 7.1 = 25.4 \ll \infty$, suggesting the absence of extreme ETA violations for static modification rules (i.e. those assigning modification at a fixed time). We note, however, that this does not necessarily imply that the dynamic rules d_θ are realistic. Further investigation of possible ETA violations under

these rules (for example, by applying a bootstrap simulation as described in Wang et al. (2006) for the point treatment setting) seems sensible and necessary, but is not performed here.

Tables 2 and 3 show estimates of β for the models $m(\theta|\beta)$ and $m(\theta, V|\beta)$, respectively. Based on the results in Table (2), Figure (1) illustrates how estimated mean counterfactual CD4 T cell count varies as a function of modification threshold θ . The estimate β_n of $m(\theta|\beta)$ suggests that the optimal threshold for modification is a CD4 T cell count category higher than the maximum CD4 T cell count category observed at baseline ($\theta_{opt} = 31.7$, corresponding to > 1500 cells), suggesting that, on average, all subjects would benefit from modifying therapy immediately following loss of suppression.

Figure (2) uses the estimated β_n of $m(\theta, V|\beta_0)$, reported in Table 3, to plot how the mean counterfactual outcome varies as a function of modification threshold θ and baseline CD4 T cell count. The results of Table 3, and Figure 2 suggest that, regardless of a subject's baseline CD4 T cell count, the highest expected counterfactual CD4 T cell count 8 months later is achieved by switching therapy immediately (i.e. at a threshold corresponding to the subject's baseline CD4 T cell count).

While suggestive, it should be emphasized that these findings are preliminary, and are intended as an example of the proposed methodology, rather than as a meaningful guide to clinical practice. Significant limitations of the current data example include small sample size, and the definition of treatment modification employed. Given small sample size, modification here was defined broadly as interruption or addition of at least 1 drug to the patient's failing antiretroviral regimen. Thus treatment simplification and interruption were included in the definition of modification, while the real clinical question of interest focuses on when the patient should be switched to a new combination regimen with the goal of re-suppressing the virus. Analyses using several large HIV cohorts based on a refined definition of treatment modification and consideration of alternative outcomes are currently underway, with the goal of providing more clinically relevant results.

4 Causal effect models for intention-to-treat interventions

4.1 The counterfactual framework for intention-to-treat causal models

The counterfactual causal inference framework for treatment interventions $\bar{a}(t)$ up till time t assumes the existence of counterfactuals indexed by static treatment interventions $\bar{a}(t)$, the corresponding link between the observed data and these counterfactuals (i.e., consistency assumption), and the sequential randomization assumption (SRA). Our framework below simply assumes the consistency and sequential randomization assumptions for all t . By applying the result of Gill and Robins (2001) and Yu and van der Laan (2002) for all t , it follows that, by construction, assuming these consistency and randomization assumptions for all t puts no restrictions on the data generating distribution. However, these assumptions do allow us to define the intention-to-treat causal parameter of interest as a parameter of the data generating distribution.

Existence of t -specific static treatment counterfactuals—For each t and each possible $\bar{a}(t) \in \mathcal{A}(t)$, we define

$$O_{\bar{a}(t)} \equiv (L_{\bar{a}(t)}, A_{\bar{a}(t)})$$

as the data one would have observed on the experimental unit if it would have been assigned $\bar{A}(t) = \bar{a}(t)$. Thus the first $t + 1$ components of $A_{\bar{a}(t)}$ are set at $\bar{a}(t)$, but the subsequent treatment

actions are random: $A_{\bar{a}(t)}(0) = a(0), \dots, A_{\bar{a}(t)}(t) = a(t)$. It is assumed that for all t and $\bar{a}(t) \in \mathcal{A}(t)$, we have

$$L_{\bar{a}(t)} = L_{A_{\bar{a}(t)}}.$$

We define $X(t) \equiv (L_{\bar{a}(t)}, A_{\bar{a}(t)}) : \bar{a}(t) \in \mathcal{A}(t)$ as the collection of treatment-specific processes corresponding with setting the first $t + 1$ treatment actions, $t = 0, \dots, \tau$. Thus, $X(\tau) = (L_{\bar{a}} : \bar{a})$ denotes the collection of counterfactual processes $L_{\bar{a}}$ indexed by fully set static treatment regimens $\bar{a} = (a(0), \dots, a(\tau))$.

t-specific temporal ordering assumption—For each time point t , we assume the usual temporal ordering assumption:

$$O_{\bar{a}(t)}(j) = O_{\bar{a}(\min(j-1, t))}(j).$$

This states that the counterfactual data at time j is only affected by past interventions.

t-specific consistency assumption—It is assumed that for all $t = 0, \dots, \tau$

$$(A, L) = (\bar{A}(t), O_{\bar{A}(t)}) = (A_{\bar{A}(t)}, L_{\bar{A}(t)}).$$

That is, we can represent O as a missing data structure on the full data structure $X(t) = \{O_{\bar{a}(t)} : \bar{a}(t) \in \mathcal{A}(t)\}$, where the missingness variable is $\bar{A}(t)$, $t = 0, \dots, \tau$. In particular, for $t = \tau$, this presents our observed longitudinal data structure as a missing data structure on a collection of treatment regimen-specific processes $X(\tau)$:

$$O = (A = \bar{A}(\tau), L = L_A).$$

t-specific sequential randomization assumption—For each t , we assume the sequential randomization assumption: for all $j = 0, \dots, t$

$$A(j) \perp X(t) \mid \bar{A}(j-1), \bar{L}(j). \tag{3}$$

We will refer to this as the strong sequential randomization assumption (SSRA). This implies, in particular, the typical sequential randomization assumption (SRA): for all $j = 0, \dots, \tau$

$$A(j) \perp X(\tau) \mid \bar{A}(j-1), \bar{L}(j). \tag{4}$$

That is, at each time-point, conditional on the observed past, the treatment at this time-point is conditionally independent of the full data $X(\tau)$. The latter sequential randomization assumption implies (and is, in essence, equivalent with) the coarsening at random (CAR) assumption on $G_{\bar{A}|X(\tau)}$ for the observed data O w.r.t. full data structure $X(\tau)$. In censored data structures, one

frequently assumes coarsening at random (CAR) (Heitjan and Rubin (1991), Jacobsen and Keiding (1995), Gill et al. (1997), in increasing generality).

Taking the τ -specific missing data representation of the observed data structure, it follows that the data generating distribution $P_{F_{X(\tau)0}, G_0}$ of O is indexed by a distribution of $X(\tau) = (L_{\bar{a}} : \bar{a})$, and the conditional probability distribution $G_0(\cdot | X(\tau))$ of \bar{A} , given $X(\tau)$. We will refer to the latter as the treatment mechanism, and we denote its probability density with $g_0(\cdot | X(\tau))$. By the chronological ordering, and our conventions above, the τ -specific missing data structure assumption is equivalent with

$$O=(L(0),A(0),L_{A(0)}(1),A(1),\dots,L_{\bar{A}(T_A-1)}(T_A),A(T_A),L_{\bar{A}(T_A)}(T_A+1)).$$

By our missing data representations for all t , we have $A = A_{\bar{A}}$, but also $A = A_{\bar{a}(t)}$ for any $\bar{a}(t) = \bar{A}(t)$, and, as a consequence, $L_{\bar{A}} = L_{\bar{a}(t)}$ for any $\bar{a}(t) = \bar{A}(t)$.

Identifiability results for static treatment interventions under the experimental treatment assignment (ETA) assumption—Under the SRA and the experimental treatment assignment assumption (ETA), it is possible to identify the treatment-specific counterfactual distributions from the observed data partial likelihood, through the G -computation formula (Robins (1999), Gill and Robins (2001), Yu and van der Laan (2002)). That is, under the assumption that $g(\bar{a} | X(\tau)) > 0$, the SRA allows us to identify the marginal distribution of $L_{\bar{a}}$, while the SSRA allows us to also identify the marginal distribution of $O_{\bar{a}(t)} = (A_{\bar{a}(t)}, L_{\bar{a}(t)})$ for any $t = 0, \dots, \tau$. Specifically, for each t , we have the following t -specific factorization of the likelihood of O :

$$dP_{F_{X(t)g_{\bar{A}(t)}|X(t)}}(O) = Q_{X(t),g_{\bar{A}(t)}|X(t)}(O) g_{\bar{A}(t)}(\bar{A}(t) | X(t)),$$

where

$$Q_{X(t),g_{\bar{A}(t)}|X(t)}(\bar{L}, \bar{A}(t), \underline{A}(t+1)) = \prod_{j=0}^{t+1} P(L(j) | \bar{L}(j-1), \bar{A}(j-1)) P(\underline{A}(t+1), \underline{L}(t+2) | \bar{L}(t+1), \bar{A}(t)),$$

and

$$g_{\bar{A}(t)}(\bar{A}(t) | X(t)) = \prod_{j=0}^t g(A(t) | \bar{A}(t-1), \bar{L}(t)).$$

For $t < \tau$, we define $\underline{A}(t) = (A(t), \dots, A(\tau))$ and $\underline{L}(t) = (L(t), \dots, L(\tau + 1))$. In addition,

$$P(\underline{A}(t+1), \underline{L}(t+2) | \bar{L}(t+1), \bar{A}(t)) \equiv \prod_{t+1}^{\tau} g(A(t) | \bar{A}(t-1), \bar{L}(t)) \prod_{t+2}^{\tau+1} P(L(t) | \bar{L}(t-1), \bar{A}(t-1)).$$

If we assume SSRA, and the ETA assumption $g_0(\bar{a}(t) | X(t)) > 0$ a.e., then we have that the probability distribution of $O_{\bar{a}(t)}$ is given by the following likelihood-based formula (G -computation formula)

$$P_{O_{\bar{a}(t)}}(\bar{L}, \underline{a}(t+1)) = Q_{0X(t),t}(\bar{L}, \bar{a}(t), \underline{a}(t+1)).$$

In other words, by setting $\bar{A}(t) = \bar{a}(t)$ in the likelihood factor $Q_{0X(t),t}$, one obtains the density of $O_{\bar{a}(t)}$. In many applications, as discussed in the introduction, this $\bar{a}(t)$ -specific experimental treatment assignment assumption $Pr(g_0(\bar{a}(t) | X(t)) > 0) = 1$ does not hold for lots of static treatment regimens $\bar{a}(t)$. The intention-to-treat parameter developed here is identifiable without the need to assume these typically unrealistic ETA-assumptions.

The observed data model implied by the causal inference assumptions—The model for the observed data structure implied by the above consistency assumptions and the strong SRA is nonparametric. As a consequence, the strong SRA and the consistency assumptions cannot be tested, but these assumptions provide us with a set of assumptions which provide the wished causal interpretation of our target parameters, defined below, of the data generating distribution. Possible data generating distributions are the elements of the nonparametric structural equation model corresponding with the causal graph implied by the time-ordering: i.e., let $L(j) = g_j(\bar{L}(j-1), \bar{A}(j-1), U)$, $A(j) = f_j(\bar{A}(j-1), \bar{L}(j), e(j))$ for arbitrary deterministic functions f_j, g_j , an arbitrary random variable U , and an exogenous random vector e . This nonparametric structural equation model is indeed a saturated model, and, for all $t \in \{0, 1, \dots, \tau\}$, it satisfies the consistency assumption and the SRA w.r.t. to the counterfactuals $X(t)$ implied by this structural equation model (see Pearl (2001), Gill and Robins (2001), Yu and van der Laan (2002)).

4.2 The intention-to-treat causal parameter

This section formally defines “intention-to-treat” counterfactuals, indexed by static treatment regimens. Specifically, for every $\bar{a} \in \mathcal{A}$, we define the individualized stopped treatment-specific process

$$X_{d(\bar{a})} = (L_{d(\bar{a})}, A_{d(\bar{a})}) \equiv (L_{\bar{a}(C_{\bar{a}})}, A_{\bar{a}(C_{\bar{a}})}),$$

where $C_{\bar{a}}$ is a counterfactual stopping time defined as

$$C_{\bar{a}} \equiv \min\{t \in \{-1, 0, \dots, \tau\} : a(t+1) \notin \mathcal{D}_{\bar{a}}(t+1) \text{ or } t=\tau\}.$$

That is, $X_{d(\bar{a})}$ is the process we would have observed on the subject if the subject had followed the static treatment \bar{a} till the end τ , or till time $C_{\bar{a}}$ at which \bar{a} for the next time point corresponds with a treatment outside the set of options $\mathcal{D}_{\bar{a}}(C_{\bar{a}} + 1)$. After the stopping time $C_{\bar{a}}$, the experimental unit is subjected to the data generating process applicable in the counterfactual

world in which one has followed \bar{a} up till time $C_{\bar{a}}$; that is, it follows its counterfactual treatment process $A_{\bar{a}(t)}$ with $t = C_{\bar{a}}$. In particular, $Y_{d(\bar{a})}$ denotes the treatment-specific outcome of interest. For example, $Y_{d(\bar{a})} = T_{d(\bar{a})} + 1$ might be the survival time under treatment regimen $d(\bar{a})$, or it might be the counterfactual outcome $Y_{d(\bar{a})}(\tau + 1)$ of a time-dependent process $Y_{d(\bar{a})}(\cdot)$ measured at a fixed time $\tau + 1$.

Missing data structure on intention-to-treat treatment-specific counterfactuals

—It is of interest to understand the information the observed data provide about these intention-to-treat counterfactuals. For any \bar{a} , we define the observed

$$C(\bar{a}) \equiv \min \{t: A(t+1) \neq a(t+1) \text{ or } a(t+1) \notin \mathcal{D}(t+1) \text{ or } t=\tau\}.$$

Thus $C(\bar{a})$ is the amount of time the experimental unit has followed $d(\bar{a})$ (if it did not even follow $a(0)$, then it equals -1), where $C(\bar{a}) \in \{-1, 0, 1, \dots, \tau\}$. Consider the indicator

$$\Delta(\bar{a}) = I(A(C(\bar{a})+1) \notin \mathcal{D}(C(\bar{a})+1) \text{ or } C(\bar{a}) = \tau). \tag{5}$$

We note that, if $\Delta(\bar{a}) = 1$, then the experimental unit has followed the intention-to-treat treatment regimen $d(\bar{a})$. Formally, we have the following link between the observed data structure and the intention-to-treat treatment-specific counterfactuals:

$$(A, L) = (A_{d(\bar{a})}, L_{d(\bar{a})}) \text{ if } \Delta(\bar{a}) = 1.$$

Thus, one could represent the observed data structure O also as

$$O = (\Delta(\bar{a}), \Delta(\bar{a})(A_{d(\bar{a})}, L_{d(\bar{a})}); \bar{a} \in \mathcal{A}).$$

That is, for each static treatment regimen \bar{a} , we observe if the experimental unit followed the individualized stopped treatment regimen $d(\bar{a})$, and if it did, then we observe its corresponding intention-to-treat counterfactual process.

Intention-to-treat causal effect parameter—Let $V \subset L(0)$ be a user-supplied set of baseline co-variables. Consider the model

$$E_0(Y_{d(\bar{a})} | V) = m(\bar{a}, V | \beta_0), \tag{6}$$

for some parametrization $\beta \rightarrow m(\cdot | \beta)$ and parameter value β_0 . Let $\beta(P_{F_{X(\tau)}, G})$ be the parameter of interest defined on the model for the observed data structure O defined by the assumptions above and the model (6), so that $\beta_0 = \beta(P_{F_{X(\tau)0}, G_0})$ denotes the true parameter value corresponding with the true data generating distribution P_0 .

We prefer not to assume the model $m(\cdot | \beta)$, but rather to use it as a working model to define a smooth version of $E_0(Y_{d(\bar{a})} | V)$ (see Neugebauer and van der Laan (2005a)). Specifically, following Neugebauer and van der Laan (2005a), we define our parameter of interest nonparametrically as

$$\beta_h(P) \equiv \arg \min_{\beta} \sum_{\bar{a}, V} (m(\bar{a}, V | \beta) - E_p(Y_{d(\bar{a})} | V))^2 h(\bar{a}, V),$$

where the weight function h is user-supplied. Thus in this case, our model is still nonparametric, but our parameter is defined by a working model $m(\cdot | \beta)$ and a weight function h . Note that, if (6) holds at P , then $\beta_h(P) = \beta(P)$ for all h . It is also of interest to note that β_h is a parameter of both the full data distribution of $X(\tau) = (L_{\bar{a}} : \bar{a} \in \mathcal{A})$ and the treatment mechanism $G_{\bar{a}|X}$.

Identifiability of intention-to-treat-specific distribution—We have the following identifiability result providing the mapping from the likelihood of O to the distribution of the intention-to-treat counterfactual data structure $O_{d(\bar{a})} = (A_{d(\bar{a})}, L_{d(\bar{a})})$.

Result 3 We have the following identifiability result:

$$P_{O_{d(\bar{a})}}(a^*, l) = I(\bar{a}^*(c_a(l)) = \bar{a}(c_a(l))) \times \prod_{t=0}^{\tau-1} P_{L(t) | \bar{L}(t-1), \bar{A}(t-1)}(l(t) | \bar{l}(t-1), \bar{a}^*(t-1)) \prod_{t=c(l)+1}^{\tau} g_0(a^*(t) | \bar{a}^*(t-1), \bar{l}(t)) \tag{7}$$

where $c_a(l) \equiv \min \{t \in \{-1, \dots, \tau\} : a(t+1) \notin \mathcal{D}(l)(t+1) \text{ or } t = \tau\}$ is the realization of the stopping time for treatment \bar{a} as identified by $L = l$ and \bar{a} .

The above identifiability result can be used to define a likelihood-based estimator. The consistency of this estimator will rely on correct estimation of the complete data generating mechanism: i.e., both the treatment mechanism g_0 and the Q_0 -factor of the density of O need to be consistently estimated. Alternatively, estimating function-based estimators can be derived that only rely on correct estimation of the treatment mechanism g_0 , or that are possibly double robust w.r.t misspecification of g_0, Q_0 . The inverse probability of treatment weighted and (possibly) double robust estimating functions, and corresponding estimators, are presented for the longitudinal data setting in Appendices C-D. Specifically, the efficient influence curve of β_h at P_0 is derived for the general longitudinal data structure, and the corresponding locally efficient estimating function and estimator are presented.

For pedagogical purposes, in the section that follows we provide a comprehensive analysis of our intention-to-treat causal effect model for the much simpler point treatment data structure, and present the corresponding likelihood-based, IPTW and DR-IPTW estimators of β_{h0} for a given h . The corresponding class of IPTW and locally efficient estimators of β_0 under the assumption that $m(\cdot | \beta)$ is a correctly specified model is obtained by letting h be arbitrary.

5 Intention-to-treat effects for point treatment

We observe the chronological data structure $O = (W, A, Y)$, where W are baseline co-variables, A is treatment, and Y is a final outcome. We assume the usual consistency assumption which states that $X = (W, (Y_a : a \in \mathcal{A}))$, and $O = (W, A, Y_A)$ is a missing data structure on X . In addition, we assume the randomization assumption which states that A is independent of X , given W : $g_0(a | X) \equiv Pr(A = a | X) = g_0(a | W) = Pr(A = a | W)$. Let $\mathcal{D} \subset W$ be a set of possible treatment options in the sense that $g_0(a | W) > 0$ for $a \in \mathcal{D}$.

Intention to Treat Causal Effect—Let $V \subset W$ be a user-supplied set of baseline co-variables. Let $Y_{d(a)} \equiv YI(a \notin \mathcal{D}) + Y_a I(a \in \mathcal{D})$ and $A_{d(a)} \equiv aI(a \in \mathcal{D}) + AI(a \notin \mathcal{D})$. Let $(W, A_{d(a)}, Y_{d(a)})$ denote the data we would observe on the experimental unit if it followed the intention-to-treat treatment $d(a)$. The parameter of interest is $\psi_0(a, V) = \Psi(P_0)(a, V) \equiv E_{P_0}(Y_{d(a)} | V)$. Note that this parameter corresponds with the mean outcome one would observe if one only intervenes (by setting $A = a$) on the experimental units for which a is a possible treatment option in the sense that $a \in \mathcal{D}$. In order to deal with the curse of dimensionality, we consider a working model $\{m(a, V | \beta) : \beta\}$ for $\psi_0(a, V)$, indexed by a Euclidean parameter β . For a user-supplied function h , let

$$\beta_h(P) \equiv \arg \min_{\beta} E_p \sum_a (\Psi(P)(a, V) - m(a, V | \beta))^2 h(a, V) \tag{8}$$

Let $\beta_{h0} = \beta_h(P_{F_{X_0}, G_0})$ be the true parameter value corresponding with the true data generating distribution $P_0 = P_{F_{X_0}, G_0}$. Note that β_h is a parameter of both the full data distribution of $X = (W, (Y_a : a \in \mathcal{A}))$ and the treatment mechanism $G_{A|X}$. We note that, if one is willing to assume that the model $m(\cdot | \beta)$ is correctly specified, then $\beta_h(P) = \beta(P)$ does not depend on h , and each estimator we present for β_h in this section is a valid estimator for β .

For any $a \in \mathcal{A}$, consider the indicator

$$\Delta(a) = I(A=a \text{ or } a \notin \mathcal{D}). \tag{9}$$

We note that, if $\Delta(a) = 1$, then the experimental unit has followed treatment $d(a)$. It is also possible that $A = a$ and $a \notin \mathcal{D}$, except if $\mathcal{D} = \{a : g_0(a | W) > 0\}$. Formally, we have the following representation of the observed data in terms of the intention-to-treat counterfactuals $(W, A_{d(a)}, Y_{d(a)})$:

$$O = (W, (\Delta(a), \Delta(a)A_{d(a)}, Y_{d(a)}) : a \in \mathcal{A}).$$

Thus, the observation $O = (W, A, Y)$ is equivalent with 1) observing the baseline co-variables W , and 2) for each a , observing if the experimental unit followed $d(a)$, and if it did, then observing $(A_{d(a)}, Y_{d(a)})$.

The model for the distribution of O is still nonparametric under the above assumptions. As a consequence, in this model all regular asymptotically linear estimators of β_{h0} at P_0 are efficient. In the next three subsections we present three estimators of β_h : likelihood-based estimator, inverse probability of treatment weighted estimator, and the estimator based on the efficient influence curve which we refer to as the double robust IPTW estimator, which is also locally efficient.

5.1 Likelihood-based estimation

The parameter $E(Y_{d(a)} | V)$ is identifiable from the observed data distribution under the above stated consistency assumption and randomization assumption. This is shown by the following result.

Result 4 Consider a joint random variable (X, A) with $X = (W, (Y_a : a \in \mathcal{A}))$, and assume that $g_0(A | X) = g_0(a | W)$. Let $\mathcal{D} \subset W$ be such that $P(\min_{a \in \mathcal{D}} g_0(a | W) > 0) = 1$. Let $(W, A, Y) =$

(W, A, Y_A) . Define the random variable $Y_{d(a)} \equiv Y_A I(a \notin \mathcal{D}) + Y(a) I(a \in \mathcal{D})$. For any $V \subset W$, we have

$$E_0(Y_{d(a)}|V) = E_0(E_0(Y|A=a, W) I(a \in \mathcal{D}) + E(Y|A, W) I(a \notin \mathcal{D}) | V)$$

In general, we have that the probability distribution of $(W, A_{d(a)}, Y_{d(a)})$ at w, a^*, y is given by

$$P_{d(a)}(w, a^*, y) = P_w(w) \left\{ I(a=a^*) P_{Y|A,W}(y|a, W) \right\}^{I(a \in \mathcal{D}(w))} \times \left\{ g_0(a^*|w) P_{Y|A,W}(y|a^*, w) \right\}^{I(a \notin \mathcal{D}(w))} .$$

One can generate the intention-to-treat counterfactuals $(W, A_{d(a)}, Y_{d(a)})$ straightforwardly. Given the marginal distribution of W , the conditional distribution of A given W , and the conditional distribution of Y given (A, W) , one generates $W, A_{d(a)}, Y_{d(a)}$ as follows: 1) generate W from P_W ; 2) if $a \notin \mathcal{D}$, then generate A from $P_{A|W}$ and set $A_{d(a)} = A$, else set $A = A_{d(a)} = a$; 3) generate Y from $P_{Y|W, A}(\cdot | W, A)$ and set $Y_{d(a)} = Y$.

By applying this data generating experiment to an estimate of the data generating distribution, one obtains a large sample $(\hat{W}_b, \hat{A}_{d(a),b}, \hat{Y}_{d(a),b}), b = 1, \dots, B$ for all $a \in \mathcal{A}$, which yields a simulation-based estimate of the distribution of $(W, A_{d(a)}, Y_{d(a)})$. Such an estimate could now also be mapped into an estimate of β_{h0} by regressing the simulated $\hat{Y}_{d(a),b}$ on a, \hat{V}_b according to the regression model $\{m(\cdot | \beta) : \beta\}$ using weights $h(a, \hat{V}_b), a \in \mathcal{A}, b = 1, \dots, B$.

If one is only concerned with estimation of the conditional mean $E(Y_{d(a)} | V)$, then it suffices to directly estimate $Q_0(a, W) = E_0(Y | A = a, W)$ with an estimator Q_n , and regress

$$Q_{n,d(a)}(a, W) \equiv Q_n(a, W) I(a \in \mathcal{D}) + Q_n(A, W) I(a \notin \mathcal{D})$$

on a, V according to the model $m(\cdot | \beta)$. That is, the likelihood-based estimator of β_{h0} can be defined as

$$\beta_n(Q_n) = \arg \min_{\beta} \sum_{i=1}^n \sum_a (Q_{n,d(a)}(a, W_i) - m(a, V_i | \beta))^2 h(a, V_i).$$

5.2 Inverse probability of treatment weighted estimation

The proposed inverse probability of treatment weighted estimator of β_{h0} is based on the following result.

Result 5 Let $\Delta(a) = I(A = a \text{ or } a \notin \mathcal{A})$. We have

$$\begin{aligned} E_0(\Delta(a)|X) &= I(a \notin \mathcal{D}) + I(a \in \mathcal{D}) g_0(a|X) \\ &= g_0(a|X)^{I(a \in \mathcal{D})}. \end{aligned}$$

We also have for any set of baseline co-variables $V \subset W$

$$E_0(Y_{d(a)}|V) = E_0\left(\frac{Y\Delta(a)}{g_0(a|X)^{I(a \in \mathcal{D})}}|V\right).$$

Proof: The first statement is trivial. Regarding the second statement we note that $\frac{Y\Delta(a)}{g_0(a|X)^{I(a \in \mathcal{D})}}$ equals

$$I(a \notin \mathcal{D})Y_A + I(a \in \mathcal{D})\frac{I(A=a)}{g_0(a|X)}Y_a.$$

The conditional expectation of the second term, given X , equals $I(a \in \mathcal{D})Y_a$. Thus, the conditional expectation, given W , of $Y\Delta(a)/g_0(a|X)^{I(a \in \mathcal{D})}$ equals the conditional expectation of $I(a \notin \mathcal{D})Y_A + I(a \in \mathcal{D})Y_{d(a)}$, given W , which proves the second statement of the result.

IPTW loss-based learning of intention-to-treat causal effect—We note that $\psi_0(a, V) \equiv E_0(Y_{d(a)}|V)$ can be estimated nonparametrically by using available machine learning/data adaptive regression algorithms. The above result shows

$$\begin{aligned} E_0(Y_{d(a)}|V) &= E_0\left(\frac{Y\Delta(a)}{g_0(a|X)^{I(a \in \mathcal{D})}}|V\right) \\ &\equiv E_0(Y_g(a)|V). \end{aligned}$$

Thus, for any user-supplied function h , we have

$$\psi_0 = \arg \min_{\psi} E_0 L_h(O, \psi | g_0),$$

where the loss function is defined as

$$L_h(O, \psi | g) \equiv \sum_{a \in \mathcal{A}} (Y_g(a) - \psi(a, V))^2 h(a, V).$$

As a consequence, we can estimate ψ_0 with the unified loss-based estimation methodology of van der Laan and Dudoit (2003) with the loss function given by $L_h(O, \psi | g)$ for any choice h . For example, given an estimator g_n of g_0 , one can estimate ψ_0 by data-adaptively regressing

$Y_{g_n,i}(a)$ on a, V_i , with weights $h(a, V_i)$, $a \in \mathcal{A}$, $i = 1, \dots, n$, using a machine learning algorithm such as the cross-validated deletion/substitution/addition (CV-DSA) algorithm of Sinisi and van der Laan (2004).

Similarly, we can apply the unified loss function-based learning approach to the inverse probability of treatment weighted loss function

$$L_h(O, \psi | g) \equiv \sum_{a \in \mathcal{A}} \frac{\Delta(a)}{g(a|X)} (Y - \psi(a, V))^2 h(a, V).$$

For example, given an estimator g_n of g_0 , one can estimate ψ_0 by data adaptively regressing Y_i on a, V_i , with weights $h(a, V_i) \Delta_i(a) / g_n(a | X_i)^{I(a \in \mathcal{D}_i)}$, $a \in \mathcal{A}$, $i = 1, \dots, n$, using a machine learning algorithm.

IPTW estimation of the intention-to-treat causal effect—The above first loss function implies the following estimator of β_{h0} :

$$\beta_n = \arg \min_{\beta} \sum_{i=1}^n \sum_a (Y_{g_n,i}(a) - m(a, V_i | \beta))^2 h(a, V_i),$$

which is a standard weighted least squares regression of $(Y_{g_n,i}(a) : a)$ on V_i for a repeated (over a) measures type data set, where the weights are given by $(h(a, V_i) : a)$. The second loss function implies the following estimator of β_{0h} :

$$\beta_n = \arg \min_{\beta} \sum_{i=1}^n \sum_a (Y_i - m(a, V_i | \beta))^2 h(a, V_i) \frac{\Delta_i(a)}{g_n(a | X_i)^{I(a \in \mathcal{D}_i)}}.$$

This is now a standard weighted least squares regression of Y_i on a, V_i for a repeated (across $a \in \mathcal{A}$) measures type data set, where the weights are given by $h(a, V_i) \Delta_i(a) / g_n(a | X_i)^{I(a \in \mathcal{D}_i)}$.

The latter weighted least squares regression estimator corresponds with the following IPTW estimating function

$$\begin{aligned} D_h(O | \beta, g) &\equiv \sum_{a \in \mathcal{A}} h(a, V) \frac{d}{d\beta} m(a, V | \beta) (Y - m(a, V | \beta)) \frac{\Delta(a)}{g(a | X)^{I(a \in \mathcal{D})}} \\ &= I(A \in \mathcal{D}) \frac{h(A, V)}{g(A | X)} \frac{d}{d\beta} m(A, V | \beta) (Y - m(A, V | \beta)) \\ &\quad + \sum_{a \notin \mathcal{D}} h(a, V) \frac{d}{d\beta} m(a, V | \beta) (Y - m(a, V | \beta)) \end{aligned}$$

By Result 5 we have that this IPTW estimating function is unbiased for β_{h0} :

$$E_0 D_h(O | \beta_0, g_0) = 0.$$

Relation to IPTW estimating function for marginal structural model—We note that in the special case that $\mathcal{D} = \mathcal{A}$ with probability 1, we have that

$$D_h(O|\beta, g) = \frac{h(A, V)}{g(A|X)} \frac{d}{d\beta} m(A, V|\beta) (Y - m(A, V|\beta))$$

reduces to the standard IPTW estimating function for a marginal structural model $E(Y_a | V) = m(a, V | \beta)$, which is known to be unbiased if indeed the ETA assumption, $\inf_{a \in \mathcal{A}} g(a | W) > 0$,

or its stabilized version, $\sup_{a \in \mathcal{A}} \frac{h(a, V)}{g(a|W)} < \infty$, holds.

5.3 Locally efficient double robust estimation

The following result provides the optimal estimating function based on the efficient influence curve of β_h at P_0 . The proof and derivation of the corresponding influence curve is provided in Appendix B.

Result 6 Consider the following estimating function:

$$\begin{aligned} D_{h,DR}(\beta_0, g_0, Q_0) &= \sum_a \frac{\Delta(a)}{g_0(a|X)^{I(a \in \mathcal{D})}} h(a, V) \frac{d}{d\beta} m(a, V|\beta_0) (Y - m(a, V|\beta_0)) \\ &\quad - \sum_{a \in \mathcal{D}} (I(A=a) - g_0(a|W)) \frac{h(a, V) \frac{d}{d\beta_0} m(a, V|\beta_0)}{g_0(a|W)} (Q_0(a, W) - m(a, V|\beta_0)). \end{aligned}$$

If $E(Y_{d(a)} | V) = m(a, V | \beta_0)$, then for all functions h

$$E_0 D_{h,DR}(\beta_0, g, Q) = 0 \text{ if } g = g_0 \text{ or } Q = Q_0.$$

If $\beta_{h0} = \arg \min_{\beta} E_0 \Sigma_a (E_0(Y_{d(a)} | V) - m(a, V | \beta))^2 h(a, V)$, then

$$E_0 D_{h,DR}(\beta_{h0}, g, Q) = 0 \text{ if } g = g_0 \text{ or } Q = Q_0.$$

The efficient influence curve of β_h at P_0 is given by $-c(\beta_{h0})^{-1} D_{h,DR}(\beta_{h0}, g_0, Q_0)$.

If P_0 is such that $E_0(Y_{d(a)} | V) = m(a, V | \beta_0)$, then β_0 does not depend on h so that $D_{h,DR}$ yields an estimating function for all functions h .

Locally efficient double robust IPTW estimator—Given an estimator g_n, Q_n of g_0, Q_0 , we can define the estimator $\beta_{h,n,DR}$ as the solution of the estimating equation

$$0 = \sum_{i=1}^n D_{h,n,DR}(O_i | \beta, g_n, Q_n).$$

If $m(\cdot | \beta)$ is linear in β , then this estimating equation is linear in β so that its solution exists in closed form. This estimator is locally efficient under regularity conditions, in the sense that it is consistent, asymptotically linear and efficient if both g_n and Q_n are consistent, and it remains consistent and asymptotically linear if only one of these two nuisance parameters is incorrectly estimated. In order to avoid technicalities, we propose the bootstrap method to obtain an estimate of the sampling distribution of $\beta_{h_n, DR}$ and to construct corresponding confidence intervals.

6 Discussion

Violations of the ETA assumption have the potential to severely bias IPTW estimates of static treatment effects. In particular, both data applications and simulation studies, such as Neugebauer and van der Laan (2005b), have exposed the importance of “practical” violations in the ETA assumption, which arise due to finite sample size. In recognition of the importance of this issue, we developed a diagnostic tool, based on bootstrap simulation, that can be applied to provide an estimate of ETA bias, in essence quantifying the lack of finite sample identifiability for the causal effect of interest (Wang et al. (2006)).

Unfortunately, in many data sets the ETA bias of the IPTW estimator is a serious concern, and having diagnosed the impact of ETA bias, one is left with the question of how to address it. In the case that the parameter of interest is a causal effect of a treatment at a single point in time, then the experimental units causing the ETA bias can be identified by their baseline covariates. Therefore, it might seem a reasonable approach to only estimate the causal effect conditional on the experimental unit having baseline covariates for which all treatments have positive probability (e.g., larger than a user-supplied $\delta > 0$). However, this seemingly sensible and natural approach forces one to restrict to a sub-distribution which may not be the sub-distribution of interest. In addition, it will require throwing away the observations not drawn from this sub-distribution. Due to the resulting forced reduction in sample size, it does not necessarily follow that the finite sample ETA bias shrinks. So, even in the point-treatment case, there does not seem to be a simple manner to deal with the ETA bias.

If the treatment is time-dependent then such a sub-sampling approach fails to be valid because the experimental units causing ETA bias are not known at baseline $t = 0$. Instead, the experimental units causing the ETA bias make themselves known during the course of the study by developing time-dependent covariates which change their set of treatment options. As a consequence, if the parameter of interest is the causal effect of a static longitudinal treatment intervention, then deleting the experimental units causing ETA bias corresponds with adjusting for variables on the pathway between the treatment and outcome of interest, an approach that is known to result in non-interpretable parameters.

To summarize, static treatment interventions are typically not realistic, and, as a consequence, are typically either non-identifiable or extremely hard to estimate based on finite samples. It is this issue which motivated the current article proposing two classes of causal effect models which do not rely on the ETA assumption, but restrict attention to interventions for which the data carries information.

In this article, we have introduced causal effect models for intention-to-treat interventions and realistic treatment rules indexed by static treatment interventions. By choosing the realistic individualized treatment rules appropriately, the proposed causal effects of realistic individualized treatment rules generalize causal effects of static treatment interventions, are always identifiable from the data while remaining interpretable, and are easier to learn based on finite samples. Intention-to-treat interventions have similar advantages to realistic treatment rules indexed by static treatments; however, in contrast to causal parameters indexed by

realistic treatment rules, the intention-to-treat causal parameter is a function of the treatment mechanism. As a result, a change in the way that treatment is assigned can change the causal effect being estimated. Thus, in applications for which the treatment mechanism followed in the observed cohort is considered a nuisance, causal effect models for realistic individualized treatment rules are the preferred approach.

In addition, we have introduced models for realistic individualized treatment rules that allow the user to supply his or her own set of realistic individualized treatment rules to be compared. As illustrated in the data example, such models for realistic individualized treatment rules identify the optimal individualized treatment rule among the user-supplied set of realistic individualized treatment rules.

Both of our proposed causal effect models force the user to identify for each experimental unit at each point in time a set of possible treatment options. We believe that this is actually a nice feature since it forces the researcher to ask the very questions which are needed to be able to obtain a collection of *identifiable and realistic* treatment regimens from data. Consultation with subject matter experts must clearly play a central role in answering these questions. For example, the researcher might need to determine which events correspond with a reduction of treatment options for the patient. If such knowledge is not available, then we propose to learn the treatment mechanism from the data and map the fitted treatment mechanism in a time-dependent set of possible treatment options for each experimental unit.

Appendix A: Causal effect models for realistic point treatment rules

In order to illustrate causal effect models for realistic treatment rules we walk through the special and simple case that treatment is assigned at a single point in time.

The observed data structure is $O = (W, A, Y)$, where W is a vector of baseline covariates, A is a subsequent treatment, and Y is a final outcome of interest.

Consistency assumption

We define the full data as the collection $X = (W, (Y(a), a \in \mathcal{A}))$ of counterfactual outcomes $Y(a)$ indexed by static treatment interventions varying over the support of the marginal distribution of A . The consistency assumption states that $O = (W, A, Y = Y(A))$.

Dynamic treatment counterfactuals

Given this standard consistency assumption, for any rule d , Y_d can be defined as Y_a with $a = d(W)$. Thus, given the existence of the random variable X defined as the collection of static treatment-specific counterfactuals, one can also define the dynamic treatment regimen-specific counterfactuals $Y_d \equiv Y_{d(W)}$ as a measurable function of X and the rule d .

Randomization assumption (RA)

We will assume the randomization assumption stating that A is independent of X , given W . The data generating distribution of O will be denoted with $P_0 = P_{F_{X|0}, g_0}$, and is indexed by the distribution $F_{X|0}$ of X and the conditional probability distribution, $g_0(\cdot | X)$, of A , given X , where $g_0(A | X) = g_0(A | W)$.

Realistic dynamic treatment assumption

Let \mathcal{A}^* be a set of dynamic treatment regimens so that for any $d \in \mathcal{A}^*$ we have

$$P(d(W) \in \mathcal{D}) = 1. \quad (10)$$

That is, for each possible baseline covariate W , the treatment assigned by this treatment rule d is an element of the set \mathcal{D} of possible treatment options. This condition on the rule d guarantees that the distribution of (W, Y_d) is identifiable by the G -computation formula (Robins (1999), Gill and Robins (2001), Yu and van der Laan (2002)):

$$P(W=w, Y_d=y) = P(W=w) P(Y=y | A=d(w), W=w).$$

Realistic individualized treatment rules indexed by static treatment regimens

Given a static treatment a , one can define a dynamic treatment regimen as one which follows the static treatment a if $a \in \mathcal{D}$, and if $a \notin \mathcal{D}$, then one assigns treatment in the set of treatment options according to a particular user-supplied rule. For example, the following construction describes such a set of dynamic treatment regimens indexed by static treatment interventions a . Suppose that the maximal set of treatment options is \mathcal{S} in the sense that $\mathcal{D} \subset \mathcal{S}$ for all subjects with probability 1. In addition, define a dissimilarity measure between any pair of elements in \mathcal{S} so that for each $s \in \mathcal{S}$ we can identify the element in \mathcal{D} closest to s . We could now define the following individualized treatment rule indexed by a static treatment a : If $a \in \mathcal{D}$, then set $A = a$, else set A equal to the element in \mathcal{D} closest to a . Notice that this defines an individualized treatment rule as a deterministic function of a static intervention $A = a$. Therefore, we can denote this set of treatment options with $d_a, a \in \mathcal{A}$.

Causal effect model for realistic individualized treatment rules

The above standard causal inference assumptions put no restrictions on the data generating distribution and thereby cannot be tested based on the data. In particular, the model for the distribution of the data implied by the above assumptions is still unspecified/nonparametric.

We define the parameter of interest on this nonparametric model as the conditional mean of Y_d , given a subset V of the baseline covariates W , for all $d \in \mathcal{A}^*$. In order to deal with the curse of dimensionality, one can follow two types of approaches. Firstly, one could assume a model

$$E_0(Y_d|V) = m(d, V|\beta_0) \quad (11)$$

for some parametrization $(d, V) \rightarrow m(d, V|\beta)$ indexed by a finite dimensional Euclidean parameter β . In this model $\beta(F_X)$ is the parameter of interest, and $\beta_0 = \beta(F_{X_0})$ is the true value of this parameter. For example, if $d = d_a$ is a deterministic function of a static treatment intervention, as in our example above, then we would have

$$E_0(Y_{d(a)}|V) = m(a, V|\beta_0).$$

Alternatively, if one believes such a model is not realistic, then it might be sensible to define the parameter of interest as

$$\beta_h(F_X) \equiv \arg \min_{\beta} \sum_{d,V} (E_{F_X}(Y_d|V) - m(d,V|\beta))^2 h(d,V).$$

If model (11) holds, then $\beta_{0h} = \beta_0$ for all h . One can map $\beta(F_{X0})$ into a corresponding optimal individualized treatment rule within each strata V :

$$d(F_{X0})(V) = \arg \max_{d \in \mathcal{A}^*} m(d,V|\beta_0).$$

Note that the parameters $\beta(F_X)$ and $\beta_h(F_X)$ are parameters of F_X . As a consequence, we can apply the general estimating function methodology as presented in van der Laan and Robins (2003) to obtain the class of all estimating functions, including the optimal double robust inverse probability of treatment weighted estimating function, which equals the efficient influence curve when evaluated at the true parameter values. The general methodology involves three steps: 1) identify the class of all full data estimating functions (formally, the space spanned by the gradients of the path-wise derivative of the parameter of interest, also called the orthogonal complement of the nuisance tangent space), 2) construct an inverse probability of treatment weighted class of estimating functions which are such that the conditional expectation, given X , maps into the class of full data estimating functions, 3) map this class of IPTW estimating functions in the so-called double robust IPTW estimating functions by subtracting the projection on tangent space spanned by all scores of the treatment mechanism under the sole model assumption RA. For details, we refer to the original paper Robins and Rotnitzky (1992) which laid out this general approach for censored data models or to Chapter 1 and 2 of van der Laan and Robins (2003).

Firstly, we need to determine the class of full data estimating functions one would obtain in the full data model for X . It follows that this class of full data estimating functions is given by:

$$\left\{ \sum_{d \in \mathcal{A}^*} h(d,V) \frac{d}{d\beta_0} m(d,V|\beta_0)(Y_d - m(d,V|\beta_0)) : h \right\}.$$

In the case that one defines the parameter of interest as

$$\beta_h(F_X) \equiv \arg \min_{\beta} \sum_{d,V} (E_{F_X}(Y_d|V) - m(d,V|\beta))^2 h(d,V),$$

then the only full data estimating function is

$$\sum_{d \in \mathcal{A}^*} h(d,V) \frac{d}{d\beta_0} m(d,V|\beta_0)(Y_d - m(d,V|\beta_0)).$$

We now need to find an IPTW-estimating function which has the property that its conditional expectation, given X , maps into the class of full data estimating functions. We can use

$$D_{h, IPTW}(O|g_0, \beta_0) = \sum_{d \in \mathcal{A}^*} \frac{I(A=d(W))}{g(A|X)} h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Y - m(d, V|\beta_0)).$$

The following result establishes the wished result.

Result 7 Assume that for all individualized treatment rules $d \in \mathcal{A}^*$ we have

$$Pr(g(d(W)|X) > 0) = 1.$$

We have for all h

$$E(D_{h, IPTW}(O|g_0, \beta_0)|X) = \sum_{d \in \mathcal{A}^*} h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Y_d - m(d, V|\beta_0)).$$

As a consequence, if $E(Y_d | V) = m(d, V | \beta_0)$, then

$$E_0 D_{h, IPTW}(O|g_0, \beta_0) = 0 \text{ for all } h,$$

and, we always have for all h

$$E_0 D_{h, IPTW}(O|g_0, \beta_{h0}) = 0.$$

Finally, we map this IPTW estimating function for β_h into the efficient estimating function by subtracting its projection on the tangent space of the treatment mechanism under RA. The following result describes this double robust IPTW estimating function, and thereby the efficient influence curve. The proof of this result is a direct consequence of Theorem 1.3 and Theorem 1.6 in van der Laan and Robins (2003).

Result 8 Let

$$D_{h, RA}(O|g_0, Q_0) = \sum_{d \in \mathcal{A}^*} \frac{I(A=d(W))}{g(A|X)} h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Q_0(A, W) - m(d, V|\beta_0)) - \sum_{d \in \mathcal{A}^*} h(d, V) \frac{d}{d\beta_0} m(d, V|\beta_0)(Q_0(d(W), W) - m(d, V|\beta_0)),$$

where $Q_0(A, W) = E_0(Y | A, W)$ and $\beta_0 = \beta(Q_0)$. The efficient influence curve of β_h in the (nonparametric) model for the data generating distribution P_0 at P_0 is given by $-c(\beta_{h0})^{-1} D_{h, DR}(O | g_0, Q_0, \beta_{h0})$, where

$$D_{h,DR}(O|g_0, Q_0, \beta_{h0}) = D_{h,IPW}(O|g_0, \beta_0) - D_{h,RA}(O|g_0, Q_0).$$

and $c(\beta) \equiv \frac{d}{d\beta} E_0 D_{h,DR}(O|g_0, Q_0, \beta)$. If $E_0(Y_d | V) = m(d, V | \beta_0)$, then for all h

$$E_0 D_{h,DR}(O|g, Q, \beta_0) = 0 \text{ if } g = g_0 \text{ or } Q = Q_0.$$

In general, for all h ,

$$E_0 D_{h,DR}(O|g, Q, \beta_{h0}) = 0 \text{ if } g = g_0 \text{ or } Q = Q_0.$$

Double robust locally efficient estimator

Given an estimator (g_n, Q_n) of the nuisance parameter (g_0, Q_0) , and a possibly data-dependent index h_n , we define the double robust locally efficient estimator as the solution β_{hn} of

$$0 = \sum_{i=1}^n D_{h_n,DR}(O_i | g_n, Q_n, \beta).$$

Under regularity conditions, the estimator β_{hn} is consistent and asymptotically linear if either g_n converges to g_0 or Q_n converges to Q_0 , and, if both nuisance parameters are consistently estimated, then β_{hn} is an asymptotically efficient estimator of β_{h0} . Therefore we call such an estimator β_{hn} locally efficient. For the formal statement for the asymptotics of this double robust estimator with the required regularity conditions, we refer to Theorem 2.4 and 2.5 in van der Laan and Robins (2003). In order to avoid technicalities, for statistical inference we propose the bootstrap method which is known to be asymptotically valid under the same conditions required to establish the asymptotic linearity of the double robust estimator $\beta_{h_n,DR}$.

Appendix B: Proofs and derivations for point treatment intention-to-treat causal effects

Proof of Result 6. We will first show the double robustness result for $D_{h,DR}$. Firstly, if $g = g_0$, then the first term has mean zero, and the second term has trivially mean zero. Consider now the case that $Q = Q_0$. Write the first terms as a sum of two terms $\sum_a \Delta(a)/g_0 | W)^{I(a \in \mathcal{D})} S(O) = \sum_{a \in \mathcal{D}} I(A = a)/g_0 S(O) + \sum_{a \notin \mathcal{D}} S(O)$ for some S , and write the second term as a difference of two terms as well. This gives:

$$\begin{aligned} & \sum_a \frac{I(A=a, a \in \mathcal{D})}{g(a|X)} h^*(Y - m) + \sum_a I(a \notin \mathcal{D}) h^*(Y - m) \\ & - \sum_{a \in \mathcal{D}} \frac{I(A=a)}{g} h^*(Q_0 - m) + \sum_{a \in \mathcal{D}} h^*(Q_0(a, W) - m). \end{aligned}$$

The expectation of the sum of the first and the third term equals zero. The second and fourth term can be written as (use that $Q_0(a, W) = E(Y_a | W)$)

$$\sum_a h^*(a, V) (YI(a \notin \mathcal{D}) + Y_a I(a \in \mathcal{D}) - m(a, V | \beta_0)) = \sum_a h^*(a, V) (Y_{d(a)} - m(a, V | \beta_0))$$

which has mean zero. This proves that $E_0 D_{h, DR}(\beta_0, g, Q_0) = 0$.

It remains to derive the efficient influence curve of the nonparametric parameter $\beta_h(P)$ and show that it is indeed given by $-c(\beta_0)^{-1} D_{h, DR}(\beta_0, g_0, Q_0)$. Since our model for the observed data structure O is non-parametric, we can use the following equivalent formulation of the model and parameter of interest in terms of the distribution of the observed data. We observe $(W, A, Y) \sim P_0$. Consider a working model $\{m(a, V | \beta) : \beta\}$ for $\psi_0(a, V) = \Psi(P_0) \equiv E_{P_0}(E_{P_0}(Y | A = a, W)I(a \in \mathcal{D}) + E_{P_0}(Y | A, W)I(a \notin \mathcal{D}) | V)$, indexed by a Euclidean parameter β . Let

$$\beta_h(P) \equiv \arg \min_{\beta} E_p \sum_a (\Psi(P)(a, V) - m(a, V | \beta))^2 h(a, V)$$

be the parameter of interest, and let the model for P_0 be nonparametric. We have that β_h is exactly the same parameter (of the data generating distribution) as defined above in terms of intention-to-treat counterfactuals. Therefore, the efficient influence curve of β_h at P_0 in this nonparametric model is also the efficient influence curve in the model in which we assume the additional non-identifiable non-testable consistency and randomization assumptions. Let $\beta_{h0} = \beta_h(P_0)$ denote the true parameter value.

Consider the estimator

$$\beta_n = \arg \min_{\beta} \sum_{i=1}^n \sum_a (Y_i - m(a, V_i | \beta))^2 h(a, V_i) \frac{\Delta_i(a)}{g_n(a | X_i)^{I(a \in \mathcal{D}_i)}}$$

We will derive the influence curve of this estimator in the case that g_n is a nonparametric estimator. Because the influence curve of a regular asymptotically linear estimator in a saturated model equals the efficient influence curve, this exercise will result in the wished efficient influence curve.

Derivation of influence curve of nonparametric estimator

Firstly, we note that β_n is the solution of

$$0 = P_n D_h(\beta, g_n) \equiv \frac{1}{n} \sum_{i=1}^n D_h(O_i | \beta, g_n),$$

where

$$D_h(O|\beta, g_n) = \sum_a \frac{\Delta(a)}{g_n(a|X)} h(a, V) \frac{d}{d\beta} m(a, V|\beta) (Y - m(a, V|\beta)),$$

where we use the notation $Pf \equiv \int f(o)dP(O)$. A standard M -estimator analysis shows that, in first order, we have

$$\beta_n - \beta_0 \approx -c(\beta_0)^{-1} \{ (P_n - P_0) D_h(\beta_0, g_0) + P_0 \{ D_h(\beta_0, g_n) - D_h(\beta_0, g_0) \} \},$$

where $c(\beta_0) = \frac{d}{d\beta_0} P_0 D_h(\beta_0, g_0)$. So, it remains to determine the influence curve $D_1(P_0)$ of the latter term $P_0 \{ D_h(\beta_0, g_n) - D_h(\beta_0, g_0) \}$. Then, the influence curve of β_n is given by:

$$IC(P_0) = -c(\beta_0)^{-1} \{ D_h(\beta_0, g_0) + D_1(P_0) \}.$$

Derivation of the influence curve $D_1(P_0)$

We note that

$$\begin{aligned} \frac{\Delta(a)}{g_n^{I(a \in \mathcal{D})}} - \frac{\Delta(a)}{g_0^{I(a \in \mathcal{D})}} &= I(a \in \mathcal{D}, A=a) \left(\frac{1}{g_n} - \frac{1}{g_0} \right) \\ &\approx -I(a \in \mathcal{D}, A=a) \frac{g_n - g_0}{g_0^2}. \end{aligned}$$

Thus,

$$\begin{aligned} P_0(D_h(\beta_0, g_n) - D_h(\beta_0, g_0)) &= \\ - \sum_a P_0 I(A=a, a \in \mathcal{D}) \frac{(g_n - g_0)(a|W)}{g_0^2(a|W)} h^*(a, V) (Y - m(a, V|\beta)), \end{aligned}$$

where we denote $h^* = hd/d\beta m$. This can be written as:

$$-P_{w_0} \sum_{a \in \mathcal{D}} \frac{(g_n - g_0)(a|W)}{g_0(a|W)} h^*(a, V) (Q_0(a, W) - m(a, V|\beta_0)).$$

We have

$$\begin{aligned} g_n(a|w) - g_0(a|w) &= \frac{(p_n - p_0)(a, w)}{p_0(w)} - \frac{p_0(a, w)}{p_0^2(w)} (p_n - P_0)(w) \\ &= \frac{p_n(a, w)}{p_0(w)} - \frac{p_0(a, w)}{p_0^2(w)} p_n(w) \\ &= \frac{p_n(a, w)}{p_0(w)} - \frac{g_0(a|w)}{p_0(w)} p_n(w), \end{aligned}$$

where $p_n(w) = \frac{1}{n} \sum_i I(W_i=w)$, $p_0(w) = Pr(W=w)$, $p_n(a,w) = \frac{1}{n} \sum_i I(A_i=a, W_i=w)$, and $p_0(a,w) = Pr(A=a, W=w)$. So $-D_{1i}$ is given by

$$P_{w_0} \sum_{a \in \mathcal{D}(W)} \left(\frac{I(A_i=a, W_i=w)}{p_0(W)} - \frac{g_0(a|W)}{p_0(W)} I(W_i=W) \right) h^*(a, V) \left(Q_0(a, W) - \frac{m_0(a, V)}{g_0(a|W)} \right).$$

Now, note that for a given function f $P_{w_0} I(W_i=W) f(W) / p_0(W) = \sum_w I(W_i=w) f(w) = f(W_i)$. Thus,

$$-D_{1i} = \sum_{a \in \mathcal{D}_i} (I(A_i=a) - g_0(a|W_i)) \frac{h^*(a, V_i)}{g_0(a|W_i)} (Q_0(a, W_i) - m(a, V_i|\beta_0)).$$

We conclude that the efficient influence curve $IC^*(P_0)$ of $\beta(P)$ at P_0 is given by:

$$\begin{aligned} -c(\beta_0) IC^*(P_0) &= D_h(\beta_0, g_0, Q_0) \\ &\equiv D_h(\beta_0, g_0) - D_{1h}(\beta_0, g_0, Q_0) \\ &\equiv \sum_a \frac{\Delta(a)}{g_0(a|X)} h(a, V) \frac{d}{d\beta} m(a, V|\beta_0) (Y - m(a, V|\beta_0)) \\ &\quad - \sum_{a \in \mathcal{D}} (I(A=a) - g_0(a|W)) \frac{h^*(a, V)}{g_0(a|W)} (Q_0(a, W) - m(a, V|\beta_0)). \end{aligned}$$

This completes the proof of Result 6.

Appendix C: Inverse probability of treatment weighted estimation of intention-to-treat causal effects of time-dependent treatments

The IPTW estimation methodology is based on the following identifiability result for the intention-to-treat treatment-specific distributions.

Result 9 For any \bar{a} , we define the observed

$$C(\bar{a}) \equiv \min \{t: A(t+1) \neq a(t+1) \text{ or } a(t+1) \notin \mathcal{D}(t+1) \text{ or } t=\tau\}.$$

Consider the indicator

$$\Delta(\bar{a}) = I(A(C(\bar{a})+1) \notin \mathcal{D}(C(\bar{a})+1) \text{ or } C(\bar{a}) = \tau). \tag{12}$$

We have

$$E_0(\Delta(\bar{a}) | X(\tau)) = g_0(\bar{a}(C(\bar{a})) | X(\tau)) = \prod_{t=0}^{C(\bar{a})} P(A(t) = a(t) | \bar{A}(t-1) = \bar{a}(t-1), \bar{L}(t)),$$

where the latter product is defined as 1 if $C(\bar{a}) = -1$. We also have that, for any set of baseline co-variables $V \subset L(0)$,

$$E_0\left(\frac{Y\Delta(\bar{a})}{g_0(\bar{a}(C(\bar{a})) | X(\tau))} | V\right) = E_0(T_{d(\bar{a})} | V).$$

Proof: Firstly, we note that

$$\Delta(\bar{a}) = \sum_{c=-1}^{\tau} I(\bar{A}(c) = \bar{a}(c), a(c+1) \notin \mathcal{D}_a(c+1) = \mathcal{D}_{\bar{a}}(c+1), c \leq C_{\bar{a}}(\bar{a})),$$

where, for simplicity, we define $I(a(\tau+1) \notin \mathcal{D}(\tau+1)) = 1$. Here we noted that $\mathcal{D}_a(c+1) = \mathcal{D}_{\bar{a}(c)}(c+1)$, and $I(c \leq C_{\bar{a}}(\bar{a})) = I(c \leq C_{\bar{a}(c)}(\bar{a}))$. In addition, we noted that at most one of the indicators in the sum can be equal to 1. Now, take the conditional expectation, given $X(\tau)$, which gives

$$\sum_{c=-1}^{\tau} g_0(\bar{a}(c) | X(\tau)) I(a(c+1) \notin \mathcal{D}_{\bar{a}}(c+1), c \leq C_{\bar{a}}(\bar{a})).$$

We have that for $c < C_{\bar{a}}(\bar{a})$, $a(c+1) \in \mathcal{D}_{\bar{a}}(c+1)$, and for $c > C_{\bar{a}}(\bar{a})$ the indicator is 0. Thus, the latter sum equals

$$g_0(\bar{a}(C_{\bar{a}}(\bar{a})) | X(\tau)).$$

This proves the first statement in the result.

Regarding the second statement, firstly, we note that $\frac{Y\Delta(\bar{a})}{g_0(\bar{a}(C(\bar{a})) | X(\tau))}$ equals

$$\sum_{c=-1}^{\tau} \frac{Y_{\bar{a}(c)}}{g_0(\bar{a}(c) | X(c))} I(\bar{A}(c) = \bar{a}(c), a(c+1) \notin \mathcal{D}_a(c+1), c \leq C_{\bar{a}}(\bar{a})),$$

where $g_0(\bar{a}(c) | X(c))$ is defined as 1 at $c = -1$. We also used that $g_0(\cdot | X(c)) = g_0(\cdot | X(\tau))$. For $c = -1$, the term equals $Y_{d(\bar{a})} I(a(0) \notin \mathcal{D}_a(0), c \leq C_{\bar{a}}(\bar{a}))$, and we will now show that for the terms with $c \geq 0$ the conditional expectation, given $X(c)$, equals $Y_{d(\bar{a})} I(a(c+1) \notin \mathcal{D}_{\bar{a}}(c+1), c \leq C_{\bar{a}}(\bar{a}))$. Consider the c -specific term for $c \geq 0$. We take the conditional expectation, given $X(c)$ (so that $Y_{\bar{a}(c)}$ and $g_0(\bar{a}(c) | X(c))$ is fixed), which yields

$$\begin{aligned} & \frac{g_0(\bar{a}(c)|X(c))}{g_0(\bar{a}(c)|X(c))} Y_{\bar{a}(c)} I(a(c+1) \notin \mathcal{D}_{\bar{a}}(c+1), c \leq C_{\bar{a}}(\bar{a})) \\ & = Y_{\bar{a}(c)} I(a(c+1) \notin \mathcal{D}_{\bar{a}}(c+1), c \leq C_{\bar{a}}(\bar{a})). \end{aligned}$$

We have that for $c < C_{\bar{a}}(\bar{a})$, $a(c+1) \in \mathcal{D}_{\bar{a}}(c+1)$ and for $c > C_{\bar{a}}(\bar{a})$ the indicator is 0. Thus, the sum over $c \in \{-1, \dots, \tau\}$ of the conditional expectations of the c -specific term, given $X(c)$, reduces to a single term corresponding with $c = C_{\bar{a}}(\bar{a})$ given by

$$Y_{\bar{a}(C_{\bar{a}}(\bar{a}))} = Y_{d(\bar{a})}.$$

Finally, note that $V \subset X(c)$ for all $c \geq 0$. This proves the second statement of the result.

IPTW-estimating function for β_h

We can estimate β_0 with the following IPTW estimator

$$\beta_n = \arg \min_{\beta} \sum_{i=1}^n \sum_{\bar{a}} (Y_i - m(\bar{a}, V_i | \beta))^2 h(\bar{a}, V_i) \frac{\Delta_i(\bar{a})}{g_n(\bar{A}_i(C_i(\bar{a})) | X_i(\tau))}.$$

This estimator is now a standard weighted least squares regression of $(Y_i : \bar{a})$ on V_i for a repeated measures type data set, where the weights are given by $h(\bar{a}, V_i) \Delta_i(\bar{a}) / g_n(\bar{A}_i(C_i(\bar{a})) | X_i(\tau))$, $i = 1, \dots, n$.

The latter weighted least squares regression corresponds with the following h -specific IPTW estimating function:

$$D_h(O | \beta, g) \equiv \sum_{\bar{a} \in \mathcal{A}} h(\bar{a}, V) \frac{d}{d\beta} m(\bar{a}, V | \beta) (Y - m(\bar{a}, V | \beta)) \frac{\Delta(\bar{a})}{g(\bar{A}(C(\bar{a})) | X(\tau))}.$$

By Result 9 we have that this estimating function is unbiased for $\beta_0 = \beta_{h0}$:

$$E_0 D_h(O | \beta_0, g_0) = 0.$$

In order to implement the above mentioned IPTW estimators of β_{h0} , or, ψ_0 itself, one needs to know the set $\{\bar{a} : \Delta(\bar{a}) = 1\}$ and the corresponding stopping times $C(\bar{a})$ for each observed O .

Algorithm for generating followed intention to treat treatments

Let \mathcal{A}^1 denote the set of treatment left over during the algorithm, and let \mathcal{E} denote the wished set of treatments with corresponding stopping times. We initiate $\mathcal{A}^1 = \mathcal{A}$, and initiate \mathcal{E} at the empty set.

Given $L(0)$, set $\mathcal{E} = \mathcal{E} \cup \{(\bar{a}, -1) : a(0) \notin \mathcal{D}(0)\}$; (thus, we add all $\bar{a} \in \mathcal{A}^1$ with $a(0) \notin \mathcal{D}(0)$, and we set $C(\bar{a}) = -1$).

$\mathcal{A}^1 = \mathcal{A}^1 / \{\bar{a} \in \mathcal{A}_1 : a(0) \notin \mathcal{D}(0)\}$: that is, we delete the selected treatments from \mathcal{A}_1 .

Given $L(0), A(0), L(1)$, set $\mathcal{E} = \mathcal{E} \cup \{(\bar{a}, 0) : \bar{a} \in \mathcal{A}^1, a(0) = A(0), a(1) \notin \mathcal{D}(1)\}$.

$\mathcal{A}^1 = \mathcal{A}^1 / \{\bar{a} \in \mathcal{A}^1 : a(0) = A(0), a(1) \notin \mathcal{D}(1)\}$.

In general, for $j = 0, \dots$, given $L(0), A(0), \dots, A(j-1), L(j)$, set $\mathcal{E} = \mathcal{E} \cup \{(\bar{a}, j-1) : \bar{a}, j-1 : \bar{a} \in \mathcal{A}^1, \bar{a}(j-1) = \bar{A}(j-1), a(j) \notin \mathcal{D}(j)\}$.

$\mathcal{A}^1 = \mathcal{A}^1 / \{\bar{a} \in \mathcal{A}_1 : \bar{a}(j-1), = \bar{A}(j-1), a(j) \notin \mathcal{D}(j)\}$. Proceed till $j = \tau$ or \mathcal{A}^1 is empty.

Appendix D: The optimal estimating function and corresponding locally efficient estimator for intention-to-treat causal effects of time-dependent treatments

The following result presents the efficient influence curve for β_h at P_0 , and its corresponding optimal estimating function.

Result 10 Given a working model $\{m(\bar{a}, V | \beta) : \beta\}$ for $\psi_0(\bar{a}, V) = \Psi(P_0)(\bar{a}, V) \equiv E_{P_0}(Y_{d(\bar{a})} | V)$ indexed by a Euclidean parameter β , our parameter of interest is defined on the nonparametric model for P_0 as

$$\beta_h(P) \equiv \arg \min_{\beta} E_p \sum_{\bar{a}} (\Psi(P)(\bar{a}, V) - m(\bar{a}, V | \beta))^2 h(\bar{a}, V).$$

Let $\beta_{h0} = \beta_h(P_0)$ denote the true parameter value. Consider the following class of estimating functions:

$$D_{h,DR}(\beta_0, g_0, Q_0) \equiv D_h(\beta_0, g_0) - D_h(\beta_0, g_0, Q_0),$$

where

$$\begin{aligned} D_h(O | \beta_0, g_0) &= \sum_{\bar{a}} \frac{\Delta(\bar{a})h(\bar{a}, V) \frac{d}{d\beta_0} m(\bar{a}, V | \beta_0)}{g_0(\bar{A}(C(\bar{a})) | X(\tau))} (Y - m(\bar{a}, V | \beta_0)) \\ D_{h,t}^*(\beta_0, g_0) &= \sum_{\bar{a}} I(C(\bar{a}) \geq t) \frac{\Delta(\bar{a})h(\bar{a}, V) \frac{d}{d\beta} m(\bar{a}, V | \beta)}{g_0(\bar{A}(C(\bar{a})) | X(\tau))} (Y - m(\bar{a}, V | \beta_0)) \\ &\quad t=0, \dots, \tau \\ D_h(\beta_0, Q_0, g_0) &= \sum_{t=0}^{\tau} E_{g_0, Q_0} (D_{h,t}^*(\beta_0, g_0) | \bar{A}(t), \bar{L}(t)) - E_{g_0, Q_0} (D_{h,t}^*(\beta_0, g_0) | \bar{A}(t-1), \bar{L}(t)). \end{aligned}$$

Here $g_0(\bar{A}(\tau) | X(\tau)) = \prod_{t=0}^{\tau} g_0(A(t) | \bar{A}(t-1), \bar{L}(t))$ and $Q_0(O) = \prod_{t=0}^{\tau+1} Q(L(t) | \bar{L}(t-1), \bar{A}(t-1))$.

We have that the efficient influence curve of β_h at P_0 is given by

$$IC^*(O) = -c(\beta_{h0})^{-1} D_{h,DR}(\beta_{h0}, g_0, Q_0).$$

If $E(Y_{d(\bar{a})} | V) = m(\bar{a}, V | \beta_0)$, then for all functions h

$$E_0 D_{h,DR}(\beta_0, g_0, Q) = 0 \text{ for all } Q.$$

If $\beta_{h0} = \arg \min_{\beta} E_0 \Sigma_{\alpha}(E_0(Y_{d(\bar{a})} | V) - m(\bar{a}, V | \beta))^2 h(\bar{a}, V)$, then

$$E_0 D_{h,DR}(\beta_{h0}, g_0, Q) = 0 \text{ for all } Q.$$

For the point treatment data structure $O = (L(0), A(0), Y)$, we have the following double robustness result:

$$E_0 D_{h,DR}(\beta_0, g, Q) = 0 \text{ if either } g = g_0 \text{ or } Q = Q_0.$$

We have not been able to establish the double robustness of $D_{h,DR}$ for time-dependent treatment processes, and suggest that the double robustness as stated for the point treatment data structure might only hold for point treatment. However, a particular type of generalized double robustness might be achievable, as defined in Robins and Rotnitzky (2001), but this remains to be established in future work.

Given an estimator g_n, Q_n of g_0, Q_0 , we can define the estimator $\beta_{hn,DR}$ as the solution of the estimating equation

$$0 = \sum_{i=1}^n D_{h,DR}(O_i | \beta, g_n, Q_n).$$

If $m(\cdot | \beta)$ is linear in β , then this estimating equation is linear in β so that its solution exists in closed form. This estimator is locally efficient under regularity conditions, in the sense that it is consistent, asymptotically linear and efficient if both g_n and Q_n are consistent, and it remains consistent and asymptotically linear if g_0 is consistently estimated. In order to avoid technicalities, we propose the bootstrap method to obtain an estimate of the sampling distribution of $\beta_{hn,DR}$ and to construct corresponding confidence intervals.

Proof of result 10. We will first show the robustness of the unbiasedness of the estimating function w.r.t. miss-specification of Q : $E_0 D_{h,DR}(\beta_0, g_0, Q) = 0$ for all Q . Firstly, we have $E_0 D_h(\beta_0, g_0) = 0$. In addition, we have $D_h(\beta_0, g_0, Q) = \Sigma_t r_t(\bar{A}(t), \bar{L}(t)) - E_{g_0}(r_t | \bar{A}(t-1), \bar{L}(t))$

for $r_t = E_{Q, g_0}(D_{h,t}^*(Q, g) | \bar{A}(t), \bar{L}(t))$ so that each t -specific term has conditional mean zero, given $\bar{A}(t), L(t)$ (for all functions r). This shows that $E_0 D_{h, DR}(\beta_0, g_0, Q) = 0$ for all Q .

Derivation of influence curve of nonparametric estimator

Consider the estimator

$$\beta_n = \arg \min_{\beta} \sum_{i=1}^n \sum_{\bar{a}} (Y_i - m(\bar{a}, V_i | \beta))^2 h(\bar{a}, V_i) \frac{\Delta_i(\bar{a})}{g_n(\bar{A}_i(C_i(\bar{a})) | X_i(\tau))}$$

We will derive the influence curve of this estimator in the case that g_n is a nonparametric estimator. Because the influence curve of a regular asymptotically linear estimator in a saturated model equals the efficient influence curve, this exercise will result in the wished efficient influence curve. In the sequel, we will use the notation \approx to indicate a first order approximation: since all our random variables are discrete and finite, the claimed asymptotic linearity of the estimator with corresponding influence curve can be fully formalized. Firstly, we note that β_n is the solution of

$$0 = P_n D_h(\beta, g_n) = 0,$$

where

$$D_h(O | \beta, g_n) = \sum_{\bar{a}} \frac{\Delta(\bar{a})}{g_n(C(\bar{a}) | X(\tau))} h(\bar{a}, V) \frac{d}{d\beta} m(\bar{a}, V | \beta) (Y - m(\bar{a}, V | \beta)).$$

In first order we have

$$\beta_n - \beta_0 = -c(\beta_0)^{-1} \{ (P_n - P_0) D_h(\beta_0, g_0) + P_0 \{ D_h(\beta_0, g_n) - D_h(\beta_0, g_0) \} \},$$

where $c(\beta_0) = \frac{d}{d\beta_0} P_0 D_h(\beta_0, g_0)$. So, we need to determine the influence curve $D_1(P_0)$ of the latter term $P_0 \{ D_h(\beta_0, g_n) - D_h(\beta_0, g_0) \}$. Then, the influence curve of β_n is given by:

$$IC(P_0) = -c(\beta_0)^{-1} \{ D_h(\beta_0, g_0) + D_1(P_0) \}.$$

We note that

$$\frac{\Delta(\bar{a})}{g_n(\bar{A}(C(\bar{a})) | X(\tau))} - \frac{\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a})) | X(\tau))} \approx -\Delta(\bar{a}) \frac{(g_n - g_0)(\bar{A}(C(\bar{a})) | X(\tau))}{g_0^2(\bar{A}(C(\bar{a})) | X(\tau))},$$

where we remind the reader that this term equals zero if $C(\bar{a}) = -1$, even when $\Delta(\bar{a}) = 1$, since in that case $\Delta(\bar{a})/g(C(\bar{a}) | X(\tau)) \equiv 1$ for both $g = g_n$ and $g = g_0$.

Thus,

$$P_0(D_h(\beta_0, g_n) - D_h(\beta_0, g_0)) \approx -\sum_{\bar{a}} P_0 \frac{\Delta(\bar{a})}{g_0^2(\bar{A}(C(\bar{a})|X(\tau))} (g_n - g_0) (\bar{A}(C(\bar{a})|X(\tau)) h^*(\bar{a}, V) (Y - m(\bar{a}, V|\beta)),$$

where we denote $h^* = hd/d\beta m$. Let

$$B(\bar{a}, O) \equiv \frac{\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a})|X(\tau))} h^*(\bar{a}, V) (Y - m(\bar{a}, V|\beta)).$$

Then the latter expectation w.r.t. P_0 can be rewritten as follows:

$$-\sum_{\bar{a}} P_0 \frac{B(\bar{a}, O)}{g_0(\bar{A}(C(\bar{a})|X(\tau))} (g_n - g_0) (\bar{A}(C(\bar{a})|X(\tau)).$$

Define $g_n(\underline{a}(l+1, c)|X(\tau)) \equiv \prod_{j=l+1}^c g_0(a(j)|\bar{a}(j-1), X(\tau))$. Now, we note that

$$\begin{aligned} & (g_n - g_0)(\bar{a}(c)|X(\tau)) \\ &= \prod_{j=0}^c g_n(a(j)|\bar{a}(j-1), X(\tau)) - \prod_{j=0}^c g_0(a(j)|\bar{a}(j-1), X(\tau)) \\ &= \sum_{l=0}^c g_n(\bar{a}(l-1)|X(\tau)) (g_n - g_0)(a(l)|\bar{a}(l-1), X(\tau)) g_0(\underline{a}(l+1, c)|X(\tau)) \\ &\approx \sum_{l=0}^c \frac{g_0(\bar{a}(c)|X(\tau))}{g_0(a(l)|\bar{a}(l-1), X(\tau))} (g_n - g_0)(a(l)|\bar{a}(l-1), X(\tau)). \end{aligned}$$

Substitution of this latter expression with $c = C(\bar{a})$ gives us now:

$$-\sum_{\bar{a}} P_0 B(\bar{a}, O) \left(\sum_{l=0}^{C(\bar{a})} \frac{(g_n - g_0)(a(l)|\bar{a}(l-1), X(\tau))}{g_0(a(l)|\bar{a}(l-1), X(\tau))} \right).$$

Let $W(l) = (\bar{A}(l-1), \bar{L}(l))$. We have

$$\begin{aligned} & g_n(a(l)|w(l)) - g_0(a(l)|w(l)) \\ &= \frac{p_n(a(l), w(l))}{p_0(w(l))} - \frac{p_0(a(l), w(l))}{p_0^2(w(l))} (p_n - p_0)(w(l)) \\ &= \frac{p_n(a(l), w(l))}{p_0(w(l))} - \frac{p_0(a(l), w(l))}{p_0^2(w(l))} p_n(w(l)) \\ &= \frac{p_n(a(l), w(l))}{p_0(w(l))} - g_0(a(l)|w(l)) \frac{p_n(w(l))}{p_0(w(l))}, \end{aligned}$$

Where

$$\begin{aligned}
 p_n(w(l)) &= \frac{1}{n} \sum_i I(W_i(l) = w(l)), \\
 p_n(a(l), w(l)) &= \frac{1}{n} \sum_i I(A_i(l) = a(l), W_i(l) = w(l)), \\
 p_0(w(l)) &= P(W(l) = w(l)),
 \end{aligned}$$

and

$$p_0(a(l), w(l)) = P(A(l) = a(l), W(l) = w(l)).$$

So we obtain

$$\begin{aligned}
 & -\sum_{\bar{a}} P_0 B(\bar{a}, O) \sum_{l=0}^{C(\bar{a})} \frac{p_n(A(l), W(l)) - g_0(A(l) | W(l)) p_n(W(l))}{g_0(A(l) | W(l)) p_0(W(l))} \\
 &= -\frac{1}{n} \sum_{i=1}^n \sum_{\bar{a}} P_0 B(\bar{a}, O) \\
 & \left(\sum_{l=0}^{C(\bar{a})} \frac{I(W_i(l) = W(l))}{g_0(A(l) | W(l)) p_0(W(l))} (I(A_i(l) = A(l)) - g_0(A(l) | W(l))) \right) \\
 &= -\frac{1}{n} \sum_{i=1}^n \sum_{\bar{a}} \sum_{l=0}^{\tau} P_0 B(\bar{a}, O) I(l \leq C(\bar{a})) \\
 & \left(\frac{I(W_i(l) = W(l))}{g_0(A(l) | W(l)) p_0(W(l))} (I(A_i(l) = A(l)) - g_0(A(l) | W(l))) \right) \\
 &= -\frac{1}{n} \sum_{i=1}^n \sum_{\bar{a}} \sum_{l=0}^{\tau} P_{0, A(l), W(l)} E_0 \left(I(l \leq C(\bar{a})) B(\bar{a}, O) | A(l), W(l) \right) \\
 & \frac{I(A_i(l) = A(l), W_i(l) = W(l)) - g_0(A(l) | W(l)) I(W_i(l) = W(l))}{g_0(A(l) | W(l)) p_0(W(l))} \\
 &= -\frac{1}{n} \sum_{i=1}^n \sum_{\bar{a}} \sum_{l=0}^{\tau} \sum_{a^*(l)} \\
 & E_0 \left(I(l \leq C(\bar{a})) B(\bar{a}, O) | A(l) = a^*(l), W_i(l) \right) (I(A_i(l) = a^*(l)) - g_0(a^*(l) | W_i(l))) \\
 &= -\frac{1}{n} \sum_{i=1}^n \sum_{\bar{a}} \sum_{l=0}^{\tau} \\
 & E_0 \left(I(l \leq C(\bar{a})) B(\bar{a}, O) | A(l) = A_i(l), W_i(l) \right) - E_0 \left(I(l \leq C(\bar{a})) B(\bar{a}, O) | W_i(l) \right)
 \end{aligned}$$

Thus we can represent $-D_1(O)$ as:

$$\begin{aligned}
 & \sum_{\bar{a}} \sum_{l=0}^{\tau} E_0 \left(I(C(\bar{a}) \geq l) B(\bar{a}, O) | A(l), W(l) \right) - E_0 \left(I(C(\bar{a}) \geq l) B(\bar{a}, O) | W(l) \right) \\
 &= \sum_{l=0}^{\tau} \left(E_0 \left(\sum_{\bar{a}} I(C(\bar{a}) \geq l) \frac{\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a})) | X(\tau))} h^*(\bar{a}, V) (Y - m(\bar{a}, V | \beta)) | A(l), W(l) \right) - \right. \\
 & \left. E_0 \left(\sum_{\bar{a}} I(C(\bar{a}) \geq l) \frac{\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a})) | X(\tau))} h^*(\bar{a}, V) (Y - m(\bar{a}, V | \beta)) | W(l) \right) \right).
 \end{aligned}$$

So,

$$-D_1(O) = \sum_{l=0}^{\tau} E_0(D_{h,l}^*(g_0, \beta) | A(l), W(l)) - E_0(D_{h,l}^*(g_0, \beta_0) | W(l)),$$

Where

$$D_{h,t}^*(O|\beta_0, g_0) \equiv \sum_{\bar{a}} I(C(\bar{a}) \geq t) \frac{\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a}))|X(\tau))} h^*(\bar{a}, V) (Y - m(\bar{a}, V|\beta)).$$

We conclude that the efficient influence curve $IC^*(P_0)$ of $\beta(P)$ at P_0 is given by:

$$\begin{aligned} -c(\beta_0)IC^*(P_0) &= D_h(\beta_0, g_0, Q_0) \\ &= D_h(\beta_0, g_0) - D_h(\beta_0, g_0, Q_0) \\ &\equiv D_h(\beta_0, g_0) - \sum_{t=0}^{\tau} E_0(D_{h,t}^*(\beta_0, g_0)|\bar{A}(t), \bar{L}(t)) \\ &\quad + \sum_{t=0}^{\tau} E_0(D_{h,t}^*(\beta_0, g_0)|\bar{A}(t-1), \bar{L}(t)). \end{aligned}$$

This completes the proof of Result 10.

References

- Gill R, Robins JM. Causal inference in complex longitudinal studies: continuous case. *Ann Stat* 2001;29(6):1785–811.
- Gill, RD.; van der Laan, MJ.; Robins, JM. Coarsening at random: characterizations, conjectures and counter-examples. In: Lin, DY.; Fleming, TR., editors. *Proceedings of the First Seattle Symposium in Biostatistics*; New York: Springer Verlag; 1997. p. 255-94.
- Heitjan DF, Rubin DB. Ignorability and coarse data. *Annals of statistics* December;1991 19(4):2244–2253.
- Hernan MA, Lanoy E, Costagliola D, Robins JM. Comparison of dynamic treatment regimes via inverse probability weighting. *Basic & Clinical Pharmacology & Toxicology* 2006;98:237–242. [PubMed: 16611197]
- Jacobsen M, Keiding N. Coarsening at random in general sample spaces and random censoring in continuous time. *Annals of Statistics* 1995;23:774–86.
- Murphy SA. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B* 2003;65(2):331–354.
- Murphy SA, van der Laan MJ, Robins JM. Marginal mean models for dynamic treatment regimes. *Journal of the American Statistical Association* 2001;96:1410–1424.
- Neugebauer, R.; van der Laan, MJ. Technical Report 134. Division of Biostatistics, University of California; Berkeley: 2005a. Locally efficient estimation of non-parametric causal effects on mean outcomes in longitudinal studies.
- Neugebauer R, van der Laan MJ. Why prefer double robust estimates. *Journal of Statistical Planning and Inference* 2005b;129(12):405–426.
- Neyman J. On the application of probability theory to agricultural experiments. *Statistical Science* 1990;5:465–480.
- Pearl, J. Cognitive systems laboratory. Technical report. University of California, Los Angeles, Department of Computer Science; 2001.
- Petersen, ML.; Deeks, SG.; Martin, JN.; van der Laan, MJ. Technical report 199. Division of Biostatistics, University of California; Berkeley: Dec. 2005 History-adjusted marginal structural models to estimate time-varying effect modification.
- Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modelling* 1986;7:1393–1512.
- Robins JM. Addendum to: “A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect”. *Comput Math Appl* 1987;14(912):923–945. *Math. Modelling* 7 (1986), 1393–1512

- Robins, JM. The analysis of randomized and non-randomized AIDS treatment trials using a new approach in causal inference in longitudinal studies. In: Sechrest, L.; Freeman, H.; Mulley, A., editors. Health Service Methodology: A Focus on AIDS. U.S. Public Health Service, National Center for Health Services Research; Washington D.C.: 1989. p. 113-159.
- Robins, JM. Information recovery and bias adjustment in proportional hazards regression analysis of randomized trials using surrogate markers. Proceeding of the Biopharmaceutical section; American Statistical Association; 1993. p. 24-33.
- Robins JM. Correcting for non-compliance in randomized trials using structural nested mean models. Communications in Statistics 1994;23:2379-2412.
- Robins, JM. Structural nested failure time models. In: Armitage, P.; Colton, T.; Andersen, PK.; Keiding, N., editors. The Encyclopedia of Biostatistics. John Wiley and Sons; Chichester, UK: 1998. p. 4372-4389.
- Robins, JM. Marginal structural models versus structural nested models as tools for causal inference. In: Halloran, ME.; Berry, D., editors. Statistical models in epidemiology, the environment, and clinical trials. Vol. 116. Springer-Verlag; New York: 1999. p. 95-134.
- Robins, JM. Robust estimation in sequentially ignorable missing data and causal inference models. Proceedings of the American Statistical Association 1999; 2000. p. 6-10.
- Robins JM. Discussion of “optimal dynamic treatment regimes” by Susan A. Murphy. Journal of the Royal Statistical Society: Series B 2003;65(2):355-366.
- Robins, JM.; Rotnitzky, A. Recovery of information and adjustment for dependent censoring using surrogate markers. In: Jewell, N.; Dietz, K.; Farewell, V., editors. AIDS Epidemiology, Methodological issues. Birkhäuser; Boston, MA: 1992. p. 297-331.
- Robins JM, Rotnitzky A. Comment on inference for semiparametric models: some questions and an answer, by Bickel, P.J. and Kwon, J. Statistica Sinica 2001;11(4):920-935.
- Rubin DB. Bayesian inference for causal effects: the role of randomization. Annals of Statistics 1978;6:34-58.
- Sinisi S, van der Laan MJ. The deletion/substitution/addition algorithm in loss function-based estimation: Applications in genomics. Journal of Statistical Methods in Molecular Biology 2004;3(1)Article 18
- van der Laan, MJ.; Dudoit, S. Technical Report 130. Division of Biostatistics, University of California; Berkeley: Nov. 2003 Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples.
- van der Laan, MJ.; Robins, JM. Unified methods for censored longitudinal data and causality. Springer; New York: 2003.
- Wang, Y.; Petersen, M.; van der Laan, MJ. Technical Report 211. Division of Biostatistics, University of California; Berkeley: 2006. A statistical method for diagnosing ETA bias in IPTW estimators.
- Yu, Z.; van der Laan, MJ. Technical Report 122. Division of Biostatistics, University of California; Berkeley: 2002. Construction of counterfactuals and the G-computation formula.

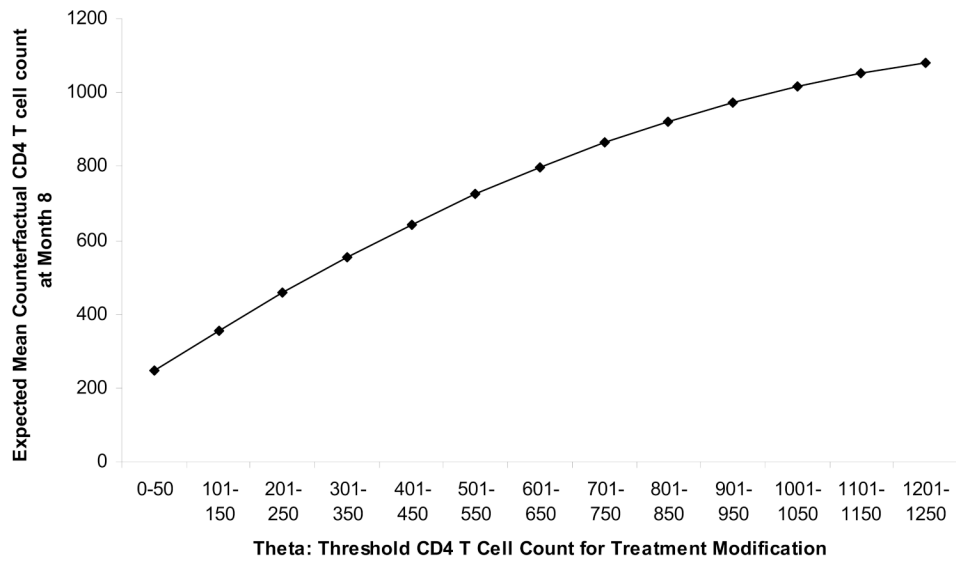


Figure 1. Mean counterfactual outcome under individualized rule where treatment modified at CD4 T cell count threshold= θ

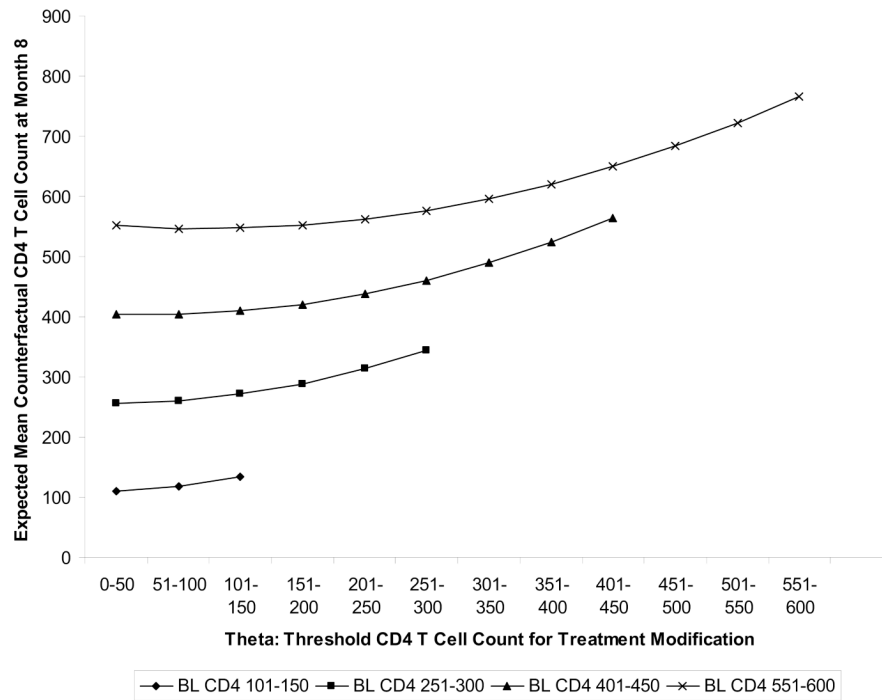


Figure 2. Mean counterfactual outcome under individualized rule where treatment modified at CD4 T cell count threshold= θ , given baseline CD4 T cell count

Table 1

Odds ratios for switching treatment based on data-adaptive fit of treatment mechanism

Covariate	Odds Ratio
Current diagnosis with an opportunistic disease	1.21
Number of protease inhibitor drugs experienced	1.11
Most recent HIV RNA level undetectable	0.44
Percent average adherence (per 10%)	0.92
Most recent CD4 T cell count (per 100 CD4 T cells)	0.92
Nadir CD4 T cell count (per 100 CD4 T cells)	1.06
Most recent HIV RNA level more than one month prior	0.90
Age (per 5 years)	0.80

Table 2

Estimated mean counterfactual CD4 T cell count 8 months after virologic failure under treatment modification at CD4 threshold θ , based on model $m(\theta|\beta) = \beta_0 + \beta_1\theta + \beta_2\theta^2$

Term	Point Estimate	95% CI
β_0	188.2	117.4, 258.9
β_1	58.9	22.7, 95.2
β_2	-0.9	-4.2, 2.3

Table 3

Estimated mean counterfactual CD4 T cell count 8 months after virologic failure under treatment modification at CD4 threshold θ , given baseline CD4 T cell count, based on model $m(\theta, V|\beta) = \beta_0 + \beta_1\theta + \beta_2\theta^2 + \beta_3\theta \times V + \beta_4\theta^2 \times V + \beta_5V$

Term	Point Estimate	95% CI
β_0	-40.8	-113.3, 31.7
β_1	-1.1	-55.5, 53.4
β_2	4.4	-1.8, 10.7
β_3	-0.9	-6.1, 4.3
β_4	-0.2	-0.7, 0.4
β_5	50.1	41.0, 59.2