University of California, Berkeley
U.C. Berkeley Division of Biostatistics Working Paper Series

Year 2006 *Paper* 203

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

Mark J. van der Laan*

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

http://biostats.bepress.com/ucbbiostat/paper203

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

Mark J. van der Laan

Abstract

An important class of models in causal inference are the so-called marginal structural models which model the comparison between counterfactual outcome distributions corresponding with a static treatment intervention, conditional on user supplied baseline covariates, based on observing a longitudinal data structure on a sample of n independent and identically distributed experimental units. Identification of a static treatment regimen specific outcome distribution based on observational data requires beyond the so-called sequential randomization assumption that each experimental unit has positive probability of following the static treatment regimen. The latter assumption is called the experimental treatment assignment assumption (ETA) (which is parameter specific). In most studies the ETA is violated for the static treatment interventions to be compared because some of the static treatment interventions cannot be followed by all experimental units due to baseline characteristics or due to the occurrence of certain events over time. For example, the development of side effects to the prescribed drug dose in a cancer patient, or the development of drug-resistance of an HIV-virus in an HIV-infected patient following the prescribed drug, describe situations in which a physician would be forced to stop the assigned treatment regimen.

In this article a generalization of marginal structural models is proposed – called intention to treat causal effect models – which does not rely on the ETA. The definition of an intention to treat causal effect requires a user-supplied definition of a time-dependent process keeping track of the possible treatment options for an experimental unit, and, if that is not available, it may be derived from a fitted treatment mechanism. The proposed intention to treat intervention enforces the static intervention until the time point at which next treatment does not belong to the set of possible treatment options, at which point the intervention is stopped. Locally

efficient estimators of the desired intention to treat causal effects are provided.

In addition causal effect models for realistic individualized treatment rules are presented which always map in the set of possible treatment options and are thereby also fully identifiable from the data; in particular it is shown that these models can be chosen to generalize marginal structural models. Analogous to Murphy et al. (2001), the corresponding locally efficient double robust inverse probability of treatment weighted estimator is presented.

1 Introduction

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

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

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 $O_1, \ldots, O_n$ of O. For simplicity, throughout this article, we will treat all random variables as discrete, but all formulas have natural continuous analogues.

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,

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

where we just 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 $j = 1, 2, \ldots$.

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 $\bar{L}(t) \equiv (L(0), \ldots, L(t))$, 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 $\bar{A}(t) \equiv (A(0), \ldots, A(t))$, $t = 0, \ldots, \tau$. 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

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

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 $L(t)$. If $\mathcal{D}(t)$ is not collected in the study, then we propose to define

$$\mathcal{D}(t) \equiv \{a(t) : g_0(a(t) \mid \bar{A}(t-1), \bar{L}(t)) > \alpha \geq 0\}$$

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.

1.1 The causal effect of a static treatment intervention.

The current literature on causal inference provides models and corresponding methods for estimation of causal effects of static treatment interventions on an outcome of interest, based on (say) sampling subjects from a particular population and recording their data over time on treatment assignment, time-dependent co-variables, and outcomes of interest.

Specifically, marginal structural models (MSM), introduced by Robins (e.g., Robins (2000a), Robins (2000b)), address the limitations of the traditional regression approach, and provide a powerful tool for causal inference in the context of longitudinal data structures. 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 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 the mean outcome under an intervention setting $\bar{A}(t) = \bar{a}(t)$ with probability 1, as a function of $\bar{a}(t)$, possibly conditional on user supplied baseline covariates. Inverse Probability of Treatment weighted (IPTW) estimators, double robust IPTW estimators, and likelihood based estimators have been proposed for the unknown causal parameters in the marginal structural model.

These methods aim to produce the results one would establish in a randomized trial randomly assigning the static treatment interventions of interest to a set of randomly sampled members of the population, *and* enforcing each subject to fully comply with the assigned static treatment intervention. As a consequence, identification of causal effects of static treatment interventions based on observational data requires, beyond the sequential randomization assumption on the treatment mechanism, also that each member of the population has a positive probability of following this static treatment intervention: this latter assumption is called the experimental treatment assignment assumption (ETA). In most studies these static treatment interventions cannot be followed by all sampled subjects due to baseline characteristics or to the occurrence of certain events over time. For example, the development

of side effects to the prescribed drug dose in a cancer patient, or the development of drug-resistance of an HIV-virus in an HIV-infected patient following the prescribed drug, describe situations in which a physician would be forced to stop the assigned treatment regimen. These situations correspond with a so called violation of the experimental treatment assignment (ETA) assumption. Theoretical or practical violation of the ETA assumption is known to result in potentially extreme bias in the Inverse-Probability of Treatment Weighted Estimators (IPTW) of the causal parameters in marginal structural models, and the full reliance of the likelihood based estimators (and DR-IPTW estimators) on model assumptions which cannot be tested from the data (Neugebauer and van der Laan (2005b)). The fact that models for static treatment interventions and their corresponding estimates aim to reproduce the results of typically unrealistic randomized trials has also been a source of philosophical criticism.

1.2 Intention to treat interventions.

In this article we propose a new class of intention to treat causal models allowing the statistical learning of intention to treat interventions which enforce the static intervention till the time point t at which the next prescribed treatment does not fall in the set of possible treatment options $\mathcal{D}(t+1)$, at which point the intervention is stopped. These causal models and their estimates aim to establish the findings of a randomized trial, applied to the same population and sample, in which each subject follows the assigned intention to treat treatment till the time point at which the prescribed treatment does not belong to the set of possible treatment options. These causal effect of individualized stopped treatment regimens are now fully identifiable from the data. As a consequence, we can develop locally efficient estimators of these causal effects without the need to assume the often unrealistic ETA assumption. In case 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 our model reduces to the marginal structural model for static treatment interventions. As a consequence, our model is a generalization of causal effect models for static treatment interventions relying on the ETA assumption to causal effect models for intention to treat interventions, which also apply when the ETA assumption is violated.

1.3 Example

Suppose that we sample subjects from an HIV-infected heavily pre-treated population which at time 0 experiences a rebound of the virus (defined by a persistent increase in viral load over the previous months) due to resistance of the virus to the prescribed drug. Suppose that $Y(12)$ is the CD4-count measured 12 months later, that the measurements $L(t), t = 0, \ldots, 12$ include viral load, CD4-count and other time-dependent characteristics of interest, and let $(A(t), t = 0, \ldots, 11)$ be the indicator process which equals zero till the time point at which the subject switches to another drug, and then jumps to 1. One might now be interested in estimation of the causal effect of time till switching on CD4 count at 12 months based on a sample of such heavily pre-treated patients who are experiencing a rebound of the virus at time 0. Specifically, we refer to Petersen et al. (2005) for a description of the SCOPE HIV-cohort, and the interest and relevance of the "when to switch question" in the HIV-AIDS research community: 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 might not necessarily imply a decrease in CD4-count. In order to be specific about the scientific question of interest it is helpful to consider randomized trials of interest. Firstly, consider the randomized trial in which one randomly assigns a switching time to each subject. In order to be able to estimate the mean outcome in the arm in which everybody is supposed to switch at time t one would need that each patient is able to fully comply with the assigned switching time t. However, suppose that some people in the population will develop an opportunistic infection or will experience side effects of the drug before time t, which make it impossible to still take the drug. Such patients cannot comply with the assigned switching time. Therefore, due to patients experiencing events which force them to switch, the causal effect of time till switch is not identifiable from the data, and, as a consequence, any of the proposed estimators suffer from potentially serious bias. However, the mean outcome of CD4 at 12 months under an intention to switch at time t is defined as the mean outcome of CD4 at 12 months if everybody who does not experience these events switches at time t, and a person who experiences an event before time t which forces a switch does switch at that time. This t-specific intention to treat mean outcome is identifiable, because every subject has a positive probability of actually following this t-specific intention to treat regimen. The difference between a t_1-specific and t_2-specific intention to treat

mean outcome measures now a causal effect of interest, which, in particular, represents a comparison between two realistic interventions.

1.4 Realistic Individualized treatment rules.

The lack of identifiability of the counterfactual distribution of the data under a static treatment intervention is due to the fact that the probability that one samples an experimental unit for which the static intervention cannot occur is larger than zero. The stopping of the static intervention at the time it can no longer be pursued resulted in the intention to treat interventions which are fully identifiable from the data. Another kind of fully identifiable intervention is an individualized treatment rule which always assigns treatments (in response to the observed history) which fall in the set of possible treatment options. The advantage of the latter kind of interventions is that its corresponding counterfactual distribution does not depend on the treatment mechanism in the study. As a consequence, causal effects comparing individualized treatment interventions are generalizable, and, in particular, a model for such causal effects also yields an optimal individualized treatment rule.

It does require some creativity with regard to proposing an interesting set of individualized treatment rules. However, it is not hard, analogue to the intention to treat interventions, to map a static treatment intervention into a corresponding individualized treatment rule which follows the static treatment intervention till the experimental unit is forced to switch at which time point one switches to a particular treatment in the set of treatment options (e.g., the one closest to the treatment assigned by the static intervention), and one sticks to this treatment till one is forced to switch again, and so on. In this manner, these individualized treatment rules are indexed by static treatment regimens, and provide approximations of the intended static treatment intervention.

Our models for causal effect of individualized treatment rules provides also an interesting double robust and locally efficient estimator of an optimal dynamic treatment regimen among a user supplied class of dynamic treatment regimens. In particular, this can be viewed as an alternative to methods for modelling and estimation of optimal dynamic treatment regimens based on a generalization of structural nested models (Robins (1989), Robins (1997), Robins (2000a) Robins (1994)), as developed in (Murphy (2003), Robins (2003)). Our proposed model for individualized treatment

rules builds on and generalizes Murphy et al. (2001), since the latter article proposes a model for a single dynamic treatment regimen conditional on baseline covariates.

1.5 Organization.

In the next section 2 we define the causal inference framework which allows us to define the causal effects of all kinds 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 effects. This framework represents a set of assumptions which do not put any restrictions on the data generating distribution, but are essential for being able to define and identify the wished causal effect from the data generating distribution. This causal inference framework states that for each experimental unit there exists intervention specific counterfactuals corresponding with setting a treatment up till time t, and that the observed data structure corresponds with observing the treatment regimen up till time t and the corresponding treatment specific process. This assumption for all t allows us now to define the causal effects of a static treatment regimen up till time t, as well, as the causal effects of dynamic treatment regimens, or, as in our case, individualized stopped treatment regimens. In addition, the causal inference framework assumes the sequential randomization assumption, a necessary, but not sufficient, assumption for identification of the causal effect of a static treatment intervention on an outcome of interest. In Section 3, given the causal inference framework, we define the intention to treat counterfactual processes, and we either assume a model for the conditional mean of the intention to treat counterfactual outcome, given some user supplied baseline co-variables, or we define our parameter of interest on the nonparametric model as the projection of the true intention to treat conditional mean outcome on the (working) model, where the L^2-projection is indexed by a weight function h. For each h, the efficient (and only) estimating function of the h-specific parameter in the nonparametric model is a particular (possibly inefficient) estimating function in the model based parameter.

For pedagogical purpose, in Section 4 we present the intention to treat causal effect model for the point treatment data structure $(W = L(0), A, Y)$, the corresponding h-specific efficient Double Robust Inverse Probability of Treatment Weighted (DR-IPTW) estimating function, and the corresponding locally efficient double robust estimator. The latter estimator is locally

efficient, in the sense that its consistency (and asymptotic linearity) relies on either correct specification of the treatment mechanism $P(A = a \mid W)$ or the regression $E(Y \mid A, W)$, and it is efficient if both are correctly specified. We will also present the likelihood based estimator and the simpler IPTW estimator, which is a special case of the DR-IPTW estimator.

In the subsequent 3 sections we will present these three types of estimators of this intention to treat causal parameter $\beta_h(P)$ or $\beta(P)$ for the general longitudinal data structure. Firstly, in Section 5 we present likelihood based estimators mapping maximum likelihood estimators of the data generating distribution in the wished parameter estimate, based on a likelihood based identifiability result for the intention to treat mean outcome. In Section 6 we derive an h-specific Inverse Probability of Treatment Weighted estimating function, and its corresponding estimator, whose consistency only relies on consistent estimation of the treatment mechanism. In Section 7 we develop the h-specific optimal estimating function, and its corresponding locally efficient estimator. The latter estimator is locally efficient, in the sense that its consistency (and asymptotic linearity) relies on correct specification of the treatment mechanism $(P(A(t) \mid \bar{A}(t-1), \bar{L}(t)) : t)$, and it is efficient if the conditional co-variable distributions $(P(L(t) \mid \bar{A}(t-1), \bar{L}(t-1)) : t)$ are correctly specified as well.

In Section 8 we present the causal model for realistic (and thereby identifiable) individualized treatment rules, and derive the corresponding locally efficient double robust inverse probability of treatment weighted estimator. Finally, Section 9 is devoted to a discussion. Some of the technical proofs are deferred to an Appendix.

2 Counterfactual Causal Inference Statistical framework:

In this section, we build on the statistical framework of counterfactuals on which marginal structural models are based. The 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 $\bar{a}(t)$ 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 assumption for all t. By applying the result in Gill and Robins (2001), Yu and van der Laan (2002)) for all t, it follows that, by construction, assuming this consistency and randomization assumptions for all t does still not put a restriction on the data generating distribution. That is, our assumptions do not put restrictions on the data generating distributions, but they 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: $A_{\bar{a}(t)}(0) = a(0), \ldots, A_{\bar{a}(t)}(t) = a(t)$. 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, $t = 0, \ldots, \tau$. Thus, $X(\tau) = (L_{\bar{a}} : \bar{a})$ denotes the collection of counterfactual processes $L_{\bar{a}}$ indexed by fully set static treatment regimens $\bar{a} = (a(0), \ldots, a(\tau))$.

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 $t = 0, \ldots, \tau$

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

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 missing-ness variable is $\bar{A}(t)$, $t = 0, \ldots, \tau$. 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 $j = 0, \ldots, t$

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

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

$$A(j) \perp X(\tau) \mid \bar{A}(j-1), \bar{L}(j). \tag{2}$$

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 $G_0(\cdot \mid X(\tau))$ of $\bar{A}$, given $X(\tau)$. We will refer to the latter as the treatment mechanism, and we denote its probability density with $g_0(\cdot \mid X(\tau))$. By the chronological ordering, and our conventions above, the τ-specific missing data structure assumption is equivalent with

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

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 (2000a), Gill and Robins (2001), Yu and van der Laan (2002)). That is, under the assumption that $g(\bar{a} \mid X(\tau)) > 0$, 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 $t = 0, \ldots, \tau$. Specifically, for each t, we have the following t-specific factorization of the likelihood of O:

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

where

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

and

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

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

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

If we assume SSRA, and the ETA assumption $g_0(\bar{a}(t) \mid X(t)) > 0$ 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 $Pr(g_0(\bar{a}(t) \mid X(t)) > 0) = 1$ does not hold for lots of static treatment regimens $\bar{a}(t)$. In this article we will define causal parameters which are identifiable without the need to assume these typically unrealistic ETA-assumptions.

2.1 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 $t \in \{0, 1 \ldots, \tau\}$, 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)).

3 Defining the parameter of interest: Intention to Treat causal effects

In practice, for many static treatment regimens $\bar{a}(t)$ we have $g(\bar{a}(t) \mid X(t)) = 0$ with positive probability: that is, there exists a subset of our population we are sampling from for which each member of this subset will not be able to complete the static treatment regimen $\bar{a}(t)$, typically due to the occurrence of events which make the treatment $a(j)$ at a certain time $j \in \{0, \ldots, t\}$ a completely wrong treatment. That is, in realistic settings we will have to acknowledge that the set of possible treatments at any point in time can be different for different subjects and that within a subject the set of possible treatments can change over time in response to time-dependent measurements. As a consequence, in realistic settings the distribution of static treatment specific counterfactuals $O_{\bar{a}(t)}$ are often not identifiable from the data. Therefore, we propose a new kind of counterfactuals indexed by static treatment regimens $\bar{a}$, which we will name intention to treat counterfactuals. Specifically, for every $\bar{a} \in \mathcal{A}$, we define the individualized stopped treatment

specific process

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

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

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

That is, $X_{d(\bar{a})}$ is the process we would have observed on the subject if the subject would follow the static treatment $\bar{a}$ till the end τ, or till time $C_{\bar{a}}$ at which next time point it corresponds with taking a treatment outside the set of options $\mathcal{D}_a(C_{\bar{a}}+1)$. 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 $\tau + 1$.

3.1 Missing Data Structure on Intention to Treat treatment specific counterfactuals:

It is of interest to understand the information the observed data provides 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 $C(\bar{a})$ (if it did not even follow $a(0)$, then it equals -1) is the amount of time the experimental unit has followed $d(\bar{a})$, where $C(\bar{a}) \in \{-1, 0, 1, \ldots, \tau\}$. Consider the indicator

$$\Delta(\bar{a}) = I(A(C(\bar{a})+1) \notin \mathcal{D}(C(\bar{a})+1) \text{ or } C(\bar{a}) = \tau). \tag{3}$$

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.

3.2 Intention to Treat Causal Effect Parameter

Let $V \subset L(0)$ be a user supplied set of baseline co-variables. Consider the model

$$E_0(Y_{d(\bar{a})} \mid V) = m(\bar{a}, V \mid \beta_0), \tag{4}$$

for some parametrization $\beta \rightarrow m(\cdot \mid \beta)$ 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 (4), 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 to not assume the model $m(\cdot \mid \beta)$, but just use it as a working model to define a smooth version of $E_0(Y_{d(\bar{a})} \mid V)$ (see Neugebauer and van der Laan (2005a)). Specifically, following Neugebauer and van der Laan (2005a), we define our parameter of interest nonparametrically as

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

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 $m(\cdot \mid \beta)$ and a weight function h. Note that, if (4) 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 $G_{\bar{A}|X}$.

Remark: In the next sections we will present an IPTW and locally efficient estimator of β_{h0} for a given h. The corresponding *class* of IPTW and locally efficient estimators of β_0 under the assumption that $m(\cdot \mid \beta)$ is a correctly specified model is obtained by letting h be arbitrary.

Before we proceed to derive the efficient influence curve of β_h at P_0 for the general longitudinal data structure, and thereby the corresponding locally efficient estimating function and estimator, we first provide a comprehensive

analysis of our intention to treat causal effect model for the much simpler point treatment data structure.

4 Intention to Treat Causal Effect Models 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: $g_0(a \mid X) \equiv Pr(A = a \mid X) = g_0(a \mid W) = Pr(A = a \mid W)$. Let $\mathcal{D} \subset W$ be a set of possible treatment options in the sense that $g_0(a \mid W) > 0$ for $a \in \mathcal{D}$.

Intention to Treat Causal Effect: Let $V \subset W$ be a user supplied set of baseline co-variables. Let $Y_{d(a)} \equiv Y I(a \notin \mathcal{D}) + Y_a I(a \in \mathcal{D})$ and $A_{d(a)} = aI(a \in \mathcal{D}) + AI(a \notin \mathcal{D})$. Let $(W, A_{d(a)}, Y_{d(a)})$ denote the data we would observe on the experimental unit if it follows the intention to treat treatment $d(a)$. The parameter of interest is $\psi_0(a, V) = \Psi(P_0)(a, V) \equiv E_{P_0}(Y_{d(a)} \mid V)$. 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 $\{m(a, V \mid \beta) : \beta\}$ for $\psi_0(a, V)$, indexed by a Euclidean parameter β. For a user supplied function h, let

$$\beta_h(P) \equiv \arg\min_{\beta} E_P \sum_a \left(\Psi(P)(a, V) - m(a, V \mid \beta)\right)^2 h(a, V) \tag{5}$$

Let $\beta_{h0} = \beta_h(P_{F_{X0}, G_0})$ be the true parameter value corresponding with the true data generating distribution $P_0 = P_{F_{X0}, G_0}$. 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 $m(\cdot \mid \beta)$ 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

$$\Delta(a) = I(A = a \text{ or } a \notin \mathcal{D})). \tag{6}$$

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 $\mathcal{D} = \{a : g_0(a \mid W) > 0\}$. 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 observing the baseline co-variables W, for each a, observing if the experimental unit followed $d(a)$, and if it did, then one observes $(A_{d(a)}, Y_{d(a)})$.

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. We present the efficient influence curve and the corresponding locally efficient estimator in the last subsection of this section. 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.

4.1 Likelihood based estimation.

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

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

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

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

$$\begin{aligned} P_{d(a)}(w, a^*, y) &= P_W(w) \left\{ I(a = a^*) P_{Y|A,W}(y \mid a, W) \right\}^{I(a \in \mathcal{D}(w))} \\ &\quad \times \left\{ g_0(a^* \mid w) P_{Y|A,W}(y \mid a^*, w) \right\}^{I(a \notin \mathcal{D}(w))}. \end{aligned}$$

One can generate the intention to treat counterfactuals $(W, A_{d(a)}, Y_{d(a)})$ in the following straightforward manner. Given the marginal distribution of W, 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 $P_{Y|W,A}(\cdot \mid W, A)$ 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 $(\hat{W}_b, \hat{A}_{d(a),b}, \hat{Y}_{d(a),b})$, $b = 1, \ldots, B$ 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 $\{m(\cdot \mid \beta) : \beta\}$ using weights $h(a, \hat{V}_b)$, $a \in \mathcal{A}$, $b = 1, \ldots, B$.

If one is only concerned with estimation of the conditional mean $E(Y_{d(a)} \mid V)$, then it suffices to directly estimate $Q_0(a, W) = E_0(Y \mid A = a, W)$ 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 $m(\cdot \mid \beta)$. That is, the likelihood based estimator of β_{h0} can be defined as

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

4.2 Inverse Probability of Treatment Weighted Estimation.

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

Result 2 *Let $\Delta(a) = I(A = a \text{ or } a \notin \mathcal{D})$. We have*

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

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

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

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

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

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

IPTW-loss based learning of intention to treat causal effect.

In this subsection we illustrate that we can estimate $\psi_0(a, V) \equiv E_0(Y_{d(a)} \mid V)$ nonparametrically by using available machine learning/data adaptive regression algorithms. The above result shows

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

Thus, for any user supplied function h, we have

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

where the loss function is defined as

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

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 $L_h(O, \psi \mid g)$ 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 $h(a, V_i)$, $a \in \mathcal{A}$, $i = 1, \ldots, n$, 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

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

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 $h(a, V_i)\Delta_i(a)/g_n(a \mid X_i)^{I(a \in \mathcal{D}_i)}$, $a \in \mathcal{A}$, $i = 1, \ldots, n$, using a machine learning algorithm.

IPTW estimation of the intention to treat causal effect

Let's now return to the estimation of the parameter β_{h0}. The above first loss function implies the following estimator of β_{h0}:

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

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}:

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

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 $h(a, V_i)\Delta_i(a)/g_n(a \mid X_i)^{I(a \in \mathcal{D}_i)}$.

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

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

By Result 2 we have that this IPTW estimating function is unbiased for β_{h0}:

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

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

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

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

4.3 The efficient influence curve: DR-IPTW estimating function

The following result provides the optimal estimating function based on the efficient infuence curve of β_h at P_0.

Result 3 *Consider the following estimating function:*

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

If $E(Y_{d(a)} \mid V) = m(a,V \mid \beta_0)$*, 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 $\beta_{h0} = \arg\min_\beta E_0 \sum_a (E_0(Y_{d(a)} \mid V) - m(a,V \mid \beta))^2 h(a,V)$*, 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 $E_0(Y_{d(a)} \mid V) = m(a,V \mid \beta_0)$, then β_0 does not depend on h so that $D_{h,DR}$ yields an estimating function for all functions h.

4.4 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

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

If $m(\cdot \mid \beta)$ 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.

5 Likelihood based estimation.

Firstly, we present the identifiability result providing the mapping from the likelihood of O to the distribution of $O_{d(\bar{a})} = (A_{d(\bar{a})}, L_{d(\bar{a})})$.

Result 4 *We have the following identifiability result:*

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

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

5.1 The likelihood based estimator.

Given a data adaptive fit Q_n of $Q_{X(\tau)0}$ and g_n of $g_{\bar{A}(\tau)|X(\tau)}$, this identifiability result, which maps (g_0, Q_0) into the distribution of $O_{d(\bar{a})}$, implies a substitution estimator $\beta_{hn}(Q_n, g_n)$ of $\beta_h(P_0)$. This substitution estimator can be evaluated/approximated with the following Monte-Carlo simulation method.

Generate intention to treat counterfactuals: The density (7) for $O_{d(\bar{a})}$ implies a simple sequential data generating experiment for generating many realizations of $O_{d(\bar{a})}$. That is, one generates sequentially from the conditional co-variable distributions of $L(j)$, given the past, setting the treatment past equal to $\bar{a}(j-1)$, till the time point $j = c$ at which

$a(c+1) \notin \mathcal{D}(c+1)$ or $c = \tau$. From then on one generates sequentially both the L as well as the future treatments, still setting the initial part of the treatment $\bar{A}(c) = \bar{a}(c)$. We denote this random variable with $\hat{O}_{d(\bar{a})}$ to indicate that we are using an estimate of the true data generating distribution.

Fitting the Intention to Treat Causal Effect Model: One generates a large collection of $\hat{O}_{d(\bar{a})}$ for a large collection (or all) $\bar{a} \in \mathcal{A}$. Now, one fits the model $m(\bar{a}, V \mid \beta)$ based on these observations $(\hat{Y}_{d(\bar{a}),b}, \bar{a}_b, \hat{V}_b)$, $b = 1, \ldots, B$, using h as weight function:

$$\beta_n \equiv \arg\min_{\beta} \sum_{b=1}^{B} (\hat{Y}_{d(\bar{a}),b} - m(\bar{a}_b, \hat{V}_b \mid \beta))^2 h(\bar{a}_b, \hat{V}_b).$$

The consistency of this estimator will rely on correct estimation of the complete data generating mechanism: i.e., both g_0 and Q_0 need to be consistently estimated. In the sequel we will present estimating function based estimators, which only rely on correct estimation of the treatment mechanism g_0.

6 Inverse Probability of Treatment Weighted Estimation.

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

Result 5 *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

$$\Delta(\bar{a}) = I(A(C(\bar{a})+1) \notin \mathcal{D}(C(\bar{a})+1) \text{ or } C(\bar{a}) = \tau). \tag{8}$$

We have

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

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

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

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

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

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

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

This proves the first statement in the result.

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

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

where $g_0(\bar{a}(c) \mid X(c))$ is defined as 1 at $c = -1$. We also used that $g_0(\cdot \mid X(c)) = g_0(\cdot \mid X(\tau))$. 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 $g_0(\bar{a}(c) \mid X(c))$ 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 the sum over $c \in \{-1, \ldots, \tau\}$ of the conditional expectations of the c-specific term, given $X(c)$, reduces to a single term corresponding with $c = C_{\bar{a}}$ 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. □

In the next subsection we use this result to define a loss function for the parameter $\psi_0(\bar{a}, V) \equiv E_0(Y_{d(\bar{a})} \mid V)$, which allows us to data adaptively estimate ψ_0 using unified loss based learning. After this subsection, we return to the estimation of the smooth parameter β_{h0}.

6.1 IPTW-Loss based learning of intention to treat causal effect.

The above result 5 shows

$$\begin{aligned} E(Y_{d(\bar{a})} \mid V) &= E\left(\frac{Y\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a})) \mid X(\tau))} \mid V\right) \\ &\equiv E\left(Y_{g_0}(\bar{a}) \mid V\right). \end{aligned}$$

Define $\psi_0(\bar{a}, V) \equiv E_0(Y_{d(\bar{a})} \mid V)$. This shows that for an arbitrary user supplied function h

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

where the loss function L_h is given by

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

As a consequence, we can estimate ψ_0 nonparametricly with the unified loss based estimation methodology of van der Laan and Dudoit (2003). For example, given an estimator g_n of g_0, one can estimate ψ_0 by data adaptively regressing $Y_{g_n,i}(\bar{a})$ on $\bar{a}, V_i$, using a machine learning algorithm, and weights $h(\bar{a}, V_i)$, $\bar{a} \in \mathcal{A}$, $i = 1, \ldots, n$. For example, we could use 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

$$L_h(O, \psi \mid g) \equiv \sum_{\bar{a} \in \mathcal{A}} \frac{\Delta(\bar{a})}{g(\bar{A}(C(\bar{a})) \mid X(\tau))} \left(Y - \psi(\bar{a}, V)\right)^2 h(\bar{a}, V),$$

which applies the IPT-weighting to the squared residuals of Y instead of applying the IPT-weighting to Y directly. In both cases h can be chosen arbitrary or be replaced by a data adaptive choice.

6.2 IPTW-estimating function for β_h.

The first IPTW-loss function implies the following least squares estimator of β_h

$$\beta_n = \arg\min_{\beta} \sum_{i=1}^{n} \sum_{\bar{a}} \left(Y_{g_n,i}(\bar{a}) - m(\bar{a}, V_i \mid \beta)\right)^2 h(\bar{a}, V_i).$$

This is a standard weighted least squares regression of $(Y_{g_n,i}(\bar{a}) : \bar{a})$ on V_i for a repeated measures type data set, where the weights are given by $(h(\bar{a}, V_i) : \bar{a})$, $i = 1, \ldots, n$. Similarly, the second IPTW-loss function implies the estimator

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

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 $h(\bar{a}, V_i)\Delta_i(\bar{a})/g_n(\bar{A}_i(C_i(\bar{a})) \mid X_i(\tau))$, $i = 1, \ldots, n$.

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

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

By Result 5 we have that this estimating function is unbiased for $\beta_0 = \beta_{h0}$:

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

6.3 Determining the followed intention to treat treatment regimens.

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

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

7 The optimal estimating function and locally efficient estimator.

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

Result 6 *Given a working model* $\{m(\bar{a}, V \mid \beta) : \beta\}$ *for* $\psi_0(\bar{a}, V) = \Psi(P_0)(\bar{a}, V) \equiv E_{P_0}(Y_{d(\bar{a})} \mid V)$ *indexed by a Euclidean parameter* β*, our parameter of interest is defined on the nonparametric model for* P_0 *as*

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

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

$$D_h(O \mid \beta_0, g_0) = \sum_{\bar{a}} \frac{\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a})) \mid X(\tau))} h(\bar{a}, V) \frac{d}{d\beta_0} m(\bar{a}, V \mid \beta_0)(Y - m(\bar{a}, V \mid \beta_0))$$

$$D^*_{h,t}(\beta_0, g_0)$$

$$= \sum_{\bar{a}} I(C(\bar{a}) \geq t) \frac{\Delta(\bar{a})}{g_0(\bar{A}(C(\bar{a})) \mid X(\tau))} h(\bar{a}, V) \frac{d}{d\beta} m(\bar{a}, V \mid \beta)(Y - m(\bar{a}, V \mid \beta_0))$$

$$t = 0, \ldots, \tau$$

$$D_h(\beta_0, Q_0, g_0) =$$

$$\sum_{t=0}^{\tau} E_{g_0,Q_0}(D^*_{h,t}(\beta_0, g_0) \mid \bar{A}(t), \bar{L}(t)) - E_{g_0,Q_0}(D^*_{h,t}(\beta_0, g_0) \mid \bar{A}(t-1), \bar{L}(t)).$$

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

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

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

If $\beta_{h0} = \arg\min_\beta E_0 \sum_a (E_0(Y_{d(\bar{a})} \mid V) - m(\bar{a}, V \mid \beta))^2 h(\bar{a}, V)$, *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 wished double robustness might only hold for point treatment. 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

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

If $m(\cdot \mid \beta)$ 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 $\beta_{h_n,DR}$ and to construct corresponding confidence intervals.

8 Causal effect models for realistic individualized treatment rules

Stopping static treatment interventions in response to individual time-dependent characteristics at the time point in which it cannot be pursued anymore is one way of defining a realistic intervention whose corresponding counterfactual distribution can be identified from the data. We named this kind of intervention an intention to treat intervention, and provided models for the effect of such interventions, and corresponding locally efficient estimators in the previous sections. An alternative realistic intervention is an individualized treatment regimen which always assigns a treatment in the set of possible treatment options. In this section we present a causal effect model for such individualized treatment interventions, and present the corresponding locally efficient double robust inverse probability of treatment weighted estimator.

An individualized treatment rule d is a vector-function $(d(0), \ldots, d(\tau))$, 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 $A(j)$, $j = 0, \ldots, \tau$.

Consistency assumption: We define the full data as the collection $X = (L_{\bar{a}} : \bar{a} \in \mathcal{A})$ of counterfactual processes $L_{\bar{a}} = (L_{\bar{a}}(0), \ldots, L_{\bar{a}}(\tau+1))$ indexed by static treatment interventions varying over the support of the marginal distribution of $A = \bar{A} = (A(0), \ldots, A(\tau))$. We also assume the temporal ordering assumption, $L_{\bar{a}}(j) = L_{\bar{a}(j-1)}(j)$, and the consistency assumption stating that $O = (A, L_A)$, which represents the longitudinal observed data structure O as a missing data structure on X with missing-ness variable A.
Dynamic treatment counterfactuals: Given this standard consistency assumption, for any rule d, L_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: $a(0) = d(0)(L(0))$, $a(1) = d(1)(a(0), \bar{L}_{a(0)}(1))$, and, in general, $a(j) = d(j)(\bar{a}(j-1), \bar{L}_{\bar{a}(j-1)}(j))$, $j = 0, \ldots, \tau$. 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. It is also of interest to note that, for each experimental unit, the rule d maps into a unique treatment regimen $\bar{a}(d, X)$, but 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. If an experimental unit follows rule d, 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 $\bar{L}_d(j) \to d(j)(\bar{L}_d(j))$.

Sequential randomization assumption: We will assume the standard sequential randomization assumption: i.e., for each $j = 0, \ldots, \tau$, $A(j)$ 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, $g_0(\cdot \mid X)$, 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

$$P(d(j)(\bar{L}_d(j)) \in \mathcal{D}_d(j), j = 0, \ldots, \tau) = 1. \tag{9}$$

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 (2000a), Gill and Robins (2001), Yu and van der Laan (2002)):

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

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

Realistic individualized treatment rules indexed by static treatment regimens: Given a static treatment regimen $\bar{a}$, one can define a dynamic treatment regimen as one which follows the static treatment regimen $\bar{a}$ till 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 $j = 0, \ldots, \tau$, 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 $\bar{a}$ till 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 $\bar{a}$), then set the next treatment equal to the element in $\mathcal{D}(t+1)$ closest to $a(t+1)$, 3) keep this treatment constant till 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 till the endpoint τ, and, 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 till 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 treatment options with $d_{\bar{a}}$, $\bar{a} \in \mathcal{A}$.

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 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 follow two types of approaches. Firstly, one could assume a model

$$E_0(Y_d \mid V) = m(d, V \mid \beta_0) \tag{10}$$

for some parametrization $(d, V) \rightarrow m(d, V \mid \beta)$ 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

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

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

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

where h is a user supplied (weight) function. If model (10) holds, then $\beta_{0h} = \beta_0$ for all h.

The article Murphy et al. (2001) models $E(Y_d \mid V)$ for a single rule d as a parametric function of V, and proposed corresponding double robust locally efficient estimators. Our model above, on the other hand, models $E(Y_d \mid V)$ as a function of d and V for d varying over a user supplied class of rules, and can therefore be viewed as a generalization.

One can map $\beta(F_{X0})$ into a corresponding optimal individualized treatment rule within each strata V:

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

Similarly, the nonparametrically defined parameter β_{h0} can be mapped into a working model based approximation of the optimal individualized treatment rule:

$$d_h(F_{X0})(V) = arg\max_{d \in \mathcal{A}^*} m(d, V \mid \beta_{h0}).$$

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 functions 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 the tangent space spanned by all scores of the treatment mechanism under the sole model assumption SRA. For details, we refer 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:

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

In the case that one defines the parameter of interest as $\beta_h(F_X) \equiv \arg\min_\beta \sum_{d,V} (E_{F_X}(Y_d \mid V) - m(d,V \mid \beta))^2 h(d,V)$, then the only full data estimating function is

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

We now need to find a 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

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

where $\bar{a} = d(\bar{L})$ is defined as $(a(0) = d_0(L(0)), a(1) = d(1)(a(0), \bar{L}(1)), \ldots, a(\tau) = (\bar{a}(\tau - 1), \bar{L}(\tau))$.

The following result establishes the wished result.

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

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

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, $a(j) = d(j)(\bar{L}_{\bar{a}(j-1)}(j))$, $j = 0, \ldots, \tau$.

We have for all h

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

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

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

and, we always have for all h

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

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

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

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

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

which completes the proof. □

Finally, we map this IPTW estimating function for β_h in 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 8 *The efficient influence curve of β_h in the (nonparametric) model for the data generating distribution P_0 at P_0 is given by $-c(\beta_{h0})^{-1}D_{h,DR}(O \mid g_0,Q_0,\beta_{h0})$, where*

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

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

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

In general, for all h,

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

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

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

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 $\beta_{h_n,DR}$.

9 Discussion

Based on our simulation studies (see e.g., Neugebauer and van der Laan (2005b)), which exposed the potentially severe bias of the IPTW estimators of causal effects of static treatment interventions due to a practical (i.e., relative to sample size) violation of the ETA assumption, we have always been very concerned with verification of the ETA assumption in our practical applications. To diagnose the presence and severity of ETA-bias of the IPTW-estimator we have developed a bootstrap simulation method estimating the bias due to the practical violation of the ETA-assumption (Wang et al. (2006)). In essence, this ETA-bias quantifies the lack of finite sample identifiability of the causal effect of interest. Unfortunately, in many data sets the bias of the IPTW estimator ETA-bias is a serious concern. Having diagnosed the impact of the ETA-bias, one is left with the question "What to do?". 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, one might decide to simply only estimate the causal effect for the data generating distribution, 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 does force one to restrict to a sub-distribution which might not be the sub-distribution of interest, and, it will require throwing away the observations not drawn from this sub-distribution. Due to the latter 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. In the treatment is time-dependent this 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 will 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 treatment intervention, then deleting the experimental units causing ETA-bias correspond with adjusting for variables on the pathway of our treatment of interest to the outcome of interest, and that is known to result in non-interpretable parameters. To summarize, static treatment interventions are typically not realistic, and, as a consequence, are typically non-identifiable, or, are 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 are not relying on the ETA assumption, but restrict attention to interventions for which the data carries information.

The proposed causal effects of intention to treat interventions aim to approximate the causal effects of static treatment interventions, generalize them, are always identifiable from the data, and are easier to learn based on finite samples, while still interpretable. By choosing the realistic individualized treatment rules appropriately the proposed causal effects of realistic individualized treatment rules also generalize causal effects of static treatment interventions, but are always fully identifiable. In addition, our models for realistic individualized treatment rules allows the user to supply its own set of realistic individualized treatment rules to be compared. In this manner, our 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 practitioner to ask the very questions which are needed to be able to obtain a collection of *identifiable and realistic* treatment regimens from data and to obtain important knowledge from subject matter experts.

For example, one might need to determine what events for a patient correspond with a reduction of treatment options. If there is not such knowledge available, then we proposed 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.

In future research we plan to implement these new methods in order to analyze various data sets of interest.

APPENDIX

Proof of Result 3.

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 $\sum_a \Delta(a)/g_0(a \mid W)^{I(a\in\mathcal{D})}S(O) = \sum_{a\in\mathcal{D}} I(A = a)/g_0 S(O) + \sum_{a\notin\mathcal{D}} S(O)$ for some S, and write the second term as a difference of two terms as well. This gives:

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

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 $Q_0(a,W) = E(Y_a \mid W)$)

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

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 $\{m(a, V \mid \beta) : \beta\}$ for $\psi_0(a, V) = \Psi(P_0) \equiv E_{P_0}(E_{P_0}(Y \mid A = a, W)I(a \in \mathcal{D}) + E_{P_0}(Y \mid A, W)I(a \notin \mathcal{D}) \mid V)$, indexed by a Euclidean parameter β. Let

$$\beta_h(P) \equiv \arg\min_\beta E_P \sum_a \left(\Psi(P)(a,V) - m(a,V\mid\beta)\right)^2 h(a,V)$$

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 assumption. Let $\beta_{h0} = \beta_h(P_0)$ denote the true parameter value.

Consider the estimator

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

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

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

where

$$D_h(O \mid \beta, g_n) = \sum_{a} \frac{\Delta(a)}{g_n(a \mid X)^{I(a \in \mathcal{D})}} h(a, V) \frac{d}{d\beta} m(a, V \mid \beta)(Y - m(a, V \mid \beta)),$$

where we use the notation $Pf \equiv \int f(o) dP(0)$. A standard M-estimator analysis shows that, in first order, we have

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

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:

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

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,

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

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

$$-P_{W0}\sum_{a\in\mathcal{D}}\frac{(g_n-g_0)(a\mid W)}{g_0(a\mid W)}h^*(a,V)(Q_0(a,W)-m(a,V\mid\beta_0)).$$

We have

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

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

$$P_{W0}\textstyle\sum_{a\in\mathcal{D}(W)}\left(\frac{I(A_i=a,W_i=W)}{p_0(W)} - \frac{g_0(a|W)}{p_0(W)}I(W_i=W)\right)h^*(a,V)\left(Q_0(a,W) - \frac{m_0(a,V)}{g_0(a|W)}\right).$$

Now, note that for a given function f $P_{W0}I(W_i=W)f(W)/p_0(W) = \sum_w I(W_i = w)f(w) = f(W_i)$. Thus,

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

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

$$-c(\beta_0)IC^*(P_0) = D_h(\beta_0, g_0, Q_0)$$

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

This completes the proof of Result 3.

Proof of result 6.

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 $D_h(\beta_0, g_0, Q) = \sum_t r_t(\bar{A}(t), \bar{L}(t)) - E_{g_0}(r_t \mid \bar{A}(t-1), \bar{L}(t))$ for $r_t = E_{Q,g_0}(D^*_{h,t}(Q, g) \mid \bar{A}(t), \bar{L}(t))$ so that each t-specific term has conditional mean zero, given $\bar{A}(t), \bar{L}(t)$ (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

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

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

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

In first order we have

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

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:

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

We note that

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

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 $\Delta(\bar{a})/g(C(\bar{a}) \mid X(\tau)) \equiv 1$ for both $g = g_n$ and $g = g_0$.

Thus,

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

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

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

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

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

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

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

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

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

where $p_n(w(l)) = \frac{1}{n} \sum_i I(W_i(l) = w(l))$, $p_n(a(l), w(l)) = \frac{1}{n} \sum_i I(A_i(l) = a(l), W_i(l) = w(l))$, $p_0(w(l)) = P(W(l) = w(l))$, and $p_0(a(l), w(l)) = P(A(l) = a(l), W(l) = w(l))$.

So we obtain

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

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

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

So,

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

where

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

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

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

This completes the proof of Result 6.

References

R. Gill and J.M. Robins. Causal inference in complex longitudinal studies: continuous case. *Ann. Stat.*, 29(6), 2001.

R.D. Gill, M.J. van der Laan, and J.M. Robins. Coarsening at random: characterizations, conjectures and counter-examples. In D.Y. Lin and T.R. Fleming, editors, *Proceedings of the First Seattle Symposium in Biostatistics*, pages 255–94, New York, 1997. Springer Verlag.

D.F. Heitjan and D.B. Rubin. Ignorability and coarse data. *Annals of statistics*, 19(4):2244–2253, December 1991.

M. Jacobsen and N. Keiding. Coarsening at random in general sample spaces and random censoring in continuous time. *Annals of Statistics*, 23:774–86, 1995.

S.A. Murphy. Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society: Series B*, 65(2):?, 2003.

S.A. Murphy, M.J. van der Laan, and J.M. Robins. Marginal mean models for dynamic treatment regimens. *Journal of the American Statistical Association*, 96:1410–1424, 2001.

R. Neugebauer and M.J. van der Laan. Locally efficient estimation of non-parametric causal effects on mean outcomes in longitudinal studies. Technical Report 134, Division of Biostatistics, University of California, Berkeley, 2005a.

R. Neugebauer and M.J. van der Laan. Why prefer double robust estimates. *Journal of Statistical Planning and Inference*, 129(1-2):405–426, 2005b.

J. Neyman. On the application of probability theory to agricultural experiments. *Statistical Science*, 5:465–480, 1990.

M.L. Petersen, S.G. Deeks, J.N. Martin, and M.J. van der Laan. History-adjusted marginal structural models: Time-varying effect modification and dynamic treatment regimens. Technical report 199, Division of Biostatistics, University of California, Berkeley, 2005.

J.M. Robins. A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512, 1986.

J.M. Robins. 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" [Math. Modelling **7** (1986), no. 9-12, 1393–1512; MR 87m:92078]. *Comput. Math. Appl.*, 14(9-12):923–945, 1987. ISSN 0097-4943.

J.M. Robins. The analysis of randomized and non-randomized aids treatment trials using a new approach in causal inference in longitudinal studies. In L. Sechrest, H. Freeman, and A. Mulley, editors, *Health Service Methodology: A Focus on AIDS*, pages 113–159. U.S. Public Health Service, National Center for Health SErvices Research, Washington D.C., 1989.

J.M. Robins. Correcting for non-compliance in randomized trials using structural nested mean models. *Communications in Statistics*, 23:2379–2412, 1994.

J.M. Robins. Structural nested failure time models. In P. Armitage, T. Colton, P.K. Andersen, and N. Keiding, editors, *The Encyclopedia of Biostatistics*. John Wiley and Sons, Chichester, UK, 1997.

J.M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology, the environment, and clinical trials (Minneapolis, MN, 1997)*, pages 95–133. Springer, New York, 2000a.

J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association*, 2000b.

J.M. Robins. Discussion of "optimal dynamic treatment regimes" by susan a. murphy. *Journal of the Royal Statistical Society: Series B*, 65(2):355–366, 2003.

D.B. Rubin. Bayesian inference for causal effects: The role of randomization. *Annals of Statistics*, 6:34–58, 1978.

S. Sinisi and M.J. van der Laan. The deletion/substitution/addition algorithm in loss function based estimation: Applications in genomics. *Journal of Statistical Methods in Molecular Biology*, 3(1), 2004.

M.J. van der Laan and S. Dudoit. Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples. Technical report, Division of Biostatistics, University of California, Berkeley, November 2003.

M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2003.

Y. Wang, M. Petersen, and M.J. van der Laan. A statistical method for diagnosing eta bias in iptw estimators. Technical report, Division of Biostatistics, University of California, Berkeley, 2006.

Z. Yu and M.J. van der Laan. Construction of counterfactuals and the G-computation formula. Technical report, Division of Biostatistics, University of California, Berkeley, 2002.