The International Journal of Biostatistics

Volume 1, Issue 1

2005

Article 4

History-Adjusted Marginal Structural Models and Statically-Optimal Dynamic Treatment Regimens

Mark J. van der Laan, *Division of Biostatistics, School of
Public Health, University of California, Berkeley*

Maya L. Petersen, *Division of Biostatistics, School of Public
Health, University of California, Berkeley*

Marshall M. Joffe, *Department of Biostatistics and
Epidemiology, University of Pennsylvania School of
Medicine*

Recommended Citation:

van der Laan, Mark J.; Petersen, Maya L.; and Joffe, Marshall M. (2005) "History-Adjusted Marginal Structural Models and Statically-Optimal Dynamic Treatment Regimens," *The International Journal of Biostatistics*: Vol. 1: Iss. 1, Article 4.

DOI: 10.2202/1557-4679.1003

©2005 by the authors. All rights reserved.

History-Adjusted Marginal Structural Models and Statically-Optimal Dynamic Treatment Regimens

Mark J. van der Laan, Maya L. Petersen, and Marshall M. Joffe

Abstract

Marginal structural models (MSM) provide a powerful tool for estimating the causal effect of a treatment. These models, introduced by Robins, model the marginal distributions of treatment-specific counterfactual outcomes, possibly conditional on a subset of the baseline covariates. Marginal structural models are particularly useful in the context of longitudinal data structures, in which each subject's treatment and covariate history are measured over time, and an outcome is recorded at a final time point. However, the utility of these models for some applications has been limited by their inability to incorporate modification of the causal effect of treatment by time-varying covariates. Particularly in the context of clinical decision making, such time-varying effect modifiers are often of considerable or even primary interest, as they are used in practice to guide treatment decisions for an individual. In this article we propose a generalization of marginal structural models, which we call history-adjusted marginal structural models (HA-MSM). These models allow estimation of adjusted causal effects of treatment, given the observed past, and are therefore more suitable for making treatment decisions at the individual level and for identification of time-dependent effect modifiers. Specifically, a HA-MSM models the conditional distribution of treatment-specific counterfactual outcomes, conditional on the whole or a subset of the observed past up till a time-point, simultaneously for all time-points. Double robust inverse probability of treatment weighted estimators have been developed and studied in detail for standard MSM. We extend these results by proposing a class of double robust inverse probability of treatment weighted estimators for the unknown parameters of the HA-MSM. In addition, we show that HA-MSM provide a natural approach to identifying the dynamic treatment regimen which follows, at each time-point, the history-adjusted (up till the most recent time point) optimal static treatment regimen. We illustrate our results using an example drawn from the treatment of HIV infection.

KEYWORDS: causal inference, confounding, counterfactual, double robust estimation, dynamic treatment regimen, G-computation estimation, inverse probability of treatment weighted estimation, longitudinal data, optimal dynamic treatment regimen, HIV, antiretroviral resistance, antiretroviral therapy

Author Notes: Maya Petersen is supported by a Pre-doctoral Fellowship in the Biological Sciences from the Howard Hughes Medical Institute. Marshall Joffe is supported by NIH grants

R01-CA-095415 and 5R29HL59184. We thank the referees for their helpful comments. Maya Petersen is supported by a Pre-doctoral Fellowship in the Biological Sciences from the Howard Hughes Medical Institute.

1 Introduction.

Consider a longitudinal study in which one collects on each randomly sampled subject the chronological data structure $O = (L(0), A(0), \dots, L(K), A(K), Y = L(K+1))$, where $L(j)$ is a time-dependent covariate measured at time j , $A(j)$ is a time-dependent treatment, and Y is the final outcome of interest measured at time $K+1$, which denotes the end of the study, $j = 0, \dots, K+1$. The simplest special case (corresponding with $K = 0$) of this data structure is $O = (W, A, Y)$, where $W \equiv L(0)$ denotes baseline covariates, $A = A(0)$ denotes treatment, and Y is the final outcome of interest. Data structures such as these are particularly common in epidemiologic and clinical research, in which subjects with changing treatment regimes or exposures are followed over time and an outcome recorded. In this article we are concerned with answering questions such as “Given a subject’s past up till time t , what is the causal effect of a future treatment regimen on the outcome of interest?” and “Given a subject’s past up till time t , what is the best treatment decision at time t ?”.

Conventional approaches to estimating the causal effect of a treatment regimen will often produce biased results when applied to the longitudinal data structure described above. The most common method for handling confounders of treatment effect is to adjust for them or, in other words, to include all confounders in a regression model of the outcome on treatment. In a point-treatment study, the resulting regression coefficient for the treatment does indeed have a causal interpretation. However, in a longitudinal study where the treatment changes over time – possibly in response to observed confounders, which are also affected by past treatment – such a regression may have no causal interpretation, even if all confounders are measured and the regression model is correctly specified. In the presence of confounders that are themselves affected by treatment, a change in treatment corresponds with a change in the distribution of the confounders, and therefore results in non-comparable regressions. One way to understand this limitation of the traditional regression approach is to consider a covariate that is affected by past treatment and itself affects both outcome and future treatment (a time-dependent confounder). If one adjusts for the covariate by including it in a traditional regression model, the resulting estimate of the causal effect of treatment will not include the component of the causal effect mediated by changes in this parameter; if one fails to adjust for the covariate, estimates of causal effect of treatment will remain confounded.

Alternatives to the traditional regression approach are further motivated by a desire to define the causal parameter of interest *a priori*, based on specific hypotheses or proposed applications. Even in the absence of time-dependent confounding, the traditional regression approach yields an estimate of the causal

effect of treatment adjusted for all measured confounders, a parameter which may be of little or secondary interest. For example, even in a point-treatment study, we may be interested in the **marginal** effect of treatment on an entire population for whom a public health intervention is proposed, as opposed to the treatment effect conditional on the values of all confounders.

Marginal structural models (MSM), introduced by Robins (e.g., Robins (2000), Robins (1999)), address these limitations of the traditional regression approach, and provide a powerful tool for causal inference in the context of longitudinal data structures. Specifically, 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.

Marginal structural models allow one to condition on baseline covariates, that is, on variables which are not affected by the treatment regimen. However, MSM do not allow one to condition on time-varying covariates. For example, an MSM might be used to ask “How does the effect of treatment vary between subpopulations defined by different levels of baseline covariates?”, but not “How does the effect of future treatment differ between subpopulations defined by the different courses of the time-dependent covariates up till the current time?”. As a result of this limitation, although marginal structural models can be used to find an optimal static treatment regimen for subpopulations defined by baseline covariates, they have not been used to find optimal *dynamic* treatment regimens. We note that dynamic treatment regimens refer to treatments which change over time in response to time-varying covariates, while static treatment regimens refer to treatments that may change at some point in the future, but do not change in response to time-varying covariates. That is, given a fit of a standard MSM, one could compute the treatment regimen which maximizes (e.g.) the treatment-specific mean or median of the outcome of interest, possibly conditional on the baseline covariates used in the MSM. Such an optimal static treatment regimen may be of significant interest, particularly in the context of public health interventions that will be applied to entire populations or subpopulations. However, in clinical medicine individual treatment decisions are generally modified over time in response to changing values of multiple covariates. The effect of these time-varying covariates on choice of an optimal treatment over time is of major clinical interest.

In this article we propose a generalization of marginal structural models, which we call history-adjusted marginal structural models (HA-MSM). Our new methodology allows estimation of causal treatment effects conditional on

time-varying covariates, and is thus more suitable than standard MSM for making treatment decisions at the individual level, and, in particular, for understanding what variables in the past are strong effect modifiers. Specifically, a HA-MSM models the conditional distribution of treatment-specific counterfactual outcomes, conditional on a user-supplied subset of the observed past up till a time-point, simultaneously for all time-points. For example, the parameter of interest modelled in our HA-MSM might be the history-adjusted treatment-specific mean

$$E(Y_{\bar{a}(j-1), \underline{a}(j)} \mid \bar{A}(j-1), \bar{L}(j)),$$

where $Y_{\bar{a}}$ denotes the treatment-specific counterfactual outcome indexed by a treatment regimen $\bar{a} = (a(0), \dots, a(K))$, $\bar{L}(j) = (L(0), \dots, L(j))$ and $\bar{A}(j-1) = (A(0), \dots, A(j-1))$ denote the observed treatment and covariate history up till time j , and $\underline{a}(j) = (a(j), \dots, a(K))$ denotes a set future treatment regimen beginning at time j . HA-MSM thus allow estimation of how future treatment will affect outcome, and how this effect is modified by a subject's observed past.

We develop a class of Inverse Probability of Treatment Weighted (IPTW) and Double Robust- Inverse Probability of Treatment Weighted (DR-IPTW) estimators for the unknown parameters of the HA-MSM. Our estimating function methodology developed in this paper shows that our estimators correspond with simultaneously (over all time-points) fitting a time-point specific MSM in which treatment and covariate history up till that time-point are treated as baseline covariates and the dependence of the counterfactual outcome on future treatment is modelled.

We show that the true parameters of the HA-MSM map into a dynamic treatment regimen which follows at each time-point the conditionally optimal static treatment regimen, given the user-supplied subset of the past (we will refer to this dynamic treatment regimen as a statically-optimal dynamic treatment regimen). In other words, at each time-point, the HA-MSM allows us to identify, for subgroups defined by covariates of interest observed up till that time-point, the best static future treatment regimen. Extended over all time-points, this approach provides a dynamic treatment regimen, in which the best optimal static treatment regimen is updated at each time-point as a result of changes in time-varying covariates. Substitution of the (DR) -IPTW estimators in the HA-MSM representation for the statically-optimal dynamic treatment regimen yields (DR) -IPTW estimators of the statically-optimal dynamic treatment regimen.

This article is organized as follows. In Section 2 we introduce an example from our current research on the treatment of Human Immunodeficiency Virus-1 (HIV-1) to clarify discussion and to provide concrete motivation for the development of HA-MSM. Throughout the article, we rely on this example to

illustrate our results.

In Section 3 we review the assumptions and estimating function methodology that underlie the use of marginal structural models for causal inference. We extend this framework to define our HA-MSM for the observed data distribution implied by A1) the representation of the observed data structure as a missing data structure on a collection of treatment regimen-specific counterfactual data structures, A2) the sequential randomization assumption on the treatment mechanism, A3) the model for the history-adjusted treatment-specific mean outcome, and A4) parametric or semi-parametric working models for the treatment mechanism and partial likelihood to deal with the curse of dimensionality. We also define the statically-optimal dynamic treatment regimen implied by the history-adjusted treatment-specific mean outcome. We illustrate the HA-MSM, and the corresponding statically-optimal dynamic treatment regimen, relying on our example drawn from the treatment of HIV infection.

In Section 4 we study nonparametric (that is, within the counterfactual SRA framework defined by assumptions A1, A2, A3) identification of the history-adjusted treatment-specific mean, and present a generalization of the G-computation formula of Robins (2000) to provide the wished identifiability of history-adjusted treatment-specific distributions. Our generalized G-computation formula provides us also with an estimation procedure which we will refer to as the G-computation estimator. Our approach involves showing that our model for the observed data distribution equals the intersection of j -specific “regular” marginal structural models in which future treatment is the treatment regimen, and the past represents baseline covariates, across $j = 0, \dots, K$. This proof is given in the Appendix and provides us also with the basis for the presentation in Section 5 of the class of (Double Robust) Inverse Probability of Treatment Weighted (DR-IPTW) estimating functions, and corresponding estimators. In Section 5 we start by describing the IPTW-estimators, which are easy to implement using standard software. We then provide a more involved description of the double robust estimators.

Section 7 is devoted to a discussion in which we point out that our approach is generic, and compare it with recent advances in the literature on semi-parametric estimation of optimal dynamic treatment regimens.

2 Example: The effect of antiretroviral drugs on the CD4 T-cells of HIV-infected patients.

HIV frequently develops resistance to the antiretroviral drugs being used to treat it, resulting in a loss of viral suppression. The clinical management of patients experiencing loss of viral suppression on antiretroviral therapy is

complex, and requires deciding how long to keep a patient on his/her current non-suppressive therapy, and how to make the decision to switch to a new therapy. Loss of viral suppression does not always result in disease progression, as reflected in declining CD4 T-cells, possibly because viral mutations that confer resistance to some antiretroviral drugs also decrease the pathogenicity of the virus. However, over time viral pathogenicity may increase, as additional resistance mutations accumulate or viral fitness increases.

We use these hypotheses to illustrate the application and interpretation of HA-MSM. Specifically we address the questions: "Is the effect of non-suppressive therapy on CD4 T-cell count modified by the presence of specific viral mutations?", "Given a subject's current viral mutation profile and the length of time he/she has already been on the current non-suppressive ART regimen, how does additional time on non-suppressive therapy affect final CD4 T-cell count?", and "Given a subject's current viral mutation profile and the length of time he/she has already been on the current non-suppressive ART regimen, what is the "best" immediate treatment decision: switch therapy or delay switching?".

Consider the following data structure. Subjects are HIV-1 infected patients treated with antiretroviral therapy. On each randomly sampled subject, the observed data consist of $O = (L(0), A(0), \dots, L(K), A(K), Y = L(K+1))$. $L(j)$ is a time-dependent vector of covariates measured at time j that includes, among other variables, viral mutation profile ($S(j)$). Y is the outcome of interest, final CD4 T-cell count, measured at time $K+1$ which denotes the end of the study, $j = 0, \dots, K+1$. (We note that, in practice, a more interesting and relevant outcome is CD4 T-cell count at a fixed interval in the future as the patient moves forward in time. However, for pedagogical purposes we focus here on CD4 T-cell count measured at a fixed time-point, simplifying the necessary notation and allowing us to defer discussion of censored outcomes until Section (7).) Finally, $A(j)$ is a time-dependent treatment reflecting whether or not an individual has switched to a new therapy or not. Specifically, $A(j) = 1$ for each time-point j for which the subject remains on his/her original treatment. When a subject switches to a new treatment, $A(j)$ jumps to 0 and remains there; we assume that a subject can only switch to a new therapy once. Thus, $A^*(j) \equiv \sum_{l=0}^j A(l)$ is the time the subject has spent on his/her original non-suppressive therapy by time j and $A^*(K)$ represents the total time a subject spends on the original non-suppressive therapy during the course of the study until outcome is evaluated at time $K+1$, or equivalently, the time until switching. We refer to this delay time between loss of virologic treatment suppression and initiation of a new treatment regimen as 'exposure time'.

For simplicity we assume that all random variables are discrete valued. Thus, CD4 T-cell count is treated as an ordinal variable, corresponding to stages of immunologic deterioration, and mutation profile is treated as a binary

variable, corresponding to the presence or absence of mutations known to alter viral fitness.

In this example the treatment $A(j)$ does not contain the actual drug combination the patient is using. Thus this type of causal analysis aims to estimate the causal effect of setting the switching time, but it allows the drug combination assignment to be randomly distributed.

As always, the interpretation of causal effects needs to be done in a careful manner. For example, switching late might be correlated with an assignment of a more aggressive drug combination, and thereby result in an improved final CD4-count. In this hypothetical situation a positive causal effect of switching late is partly due to the indirect effect of switching late on the assignment of drugs.

3 History-adjusted marginal structural model for the observed data.

In this section, we begin by outlining 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). We briefly review 'standard' marginal structural models, including the assumptions on which they rely. We then extend this framework to define the causal parameter of interest estimated by the HA-MSM methodology, the history-adjusted treatment-specific mean outcome, and discuss modelling this parameter.

Review of marginal structural models and the counterfactual framework for causal inference. Marginal structural models define causal effects using the counterfactual framework. Specifically, the causal effect of a treatment on an outcome for an individual is defined as the difference in outcome that would be observed if the same individual received and did not receive the treatment. Since only one of these outcomes is observed, they are termed counterfactual. Marginal structural models are models of these treatment-specific counterfactual outcomes. For example, a marginal structural model in the HIV example might regress the counterfactual CD4 T-cell count for the population on exposure time. Note that such a model differs from a simple regression of observed final CD4 T-cell count on observed exposure time.

To define marginal structural models, one needs to assume that counterfactuals exist, or in other words, that it makes sense to talk about what an

individual's outcome would have been under a treatment different from that which he/she actually received. More formally, for each possible treatment regimen one needs to assume the existence of the treatment regimen-specific counterfactual longitudinal random variable, whose distribution represents the population distribution of the complete (i.e., covariates and outcome) longitudinal data structure in the counterfactual world in which each subject follows that particular treatment regimen. In addition, one assumes that the observed longitudinal data structure on a subject equals the counterfactual longitudinal data structure corresponding with the treatment regimen the subject actually followed. For example, we assume that the final CD4 T-cell count observed for a subject who switched therapies after six months is the same as the counterfactual final CD4 T-cell count for that subject under a six month exposure time. This assumption is called the consistency assumption.

These assumptions allow us to view a subject's observed longitudinal data as a missing data structure, where the full data is a collection of treatment-specific counterfactual data structures corresponding to a set of values of the treatment regimen, and the missingness variable is the actual treatment regimen followed by the subject. In terms of our HIV example, the full data structure consists of the CD4 T-cell count and viral mutation trajectories over time that would have been observed on a randomly sampled subject of the population under each possible exposure time. In reality, we only observe the CD4 T-cell count and mutation trajectories under one exposure time per subject; the values that these covariates would have had if the subject had followed a different exposure time are missing. A marginal structural model is thus a model on the full data distribution that represents assumptions about how the treatment-specific counterfactual distributions change with treatment.

In censored (e.g., missing) data structures, some kind of conditional independence assumption on the conditional distribution of the censoring variable, given the full data, needs to be assumed in order to identify parameters of the full data distribution. 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). The analogue of the coarsening at random assumption for the counterfactual missing data structure is the so called sequential randomization assumption (SRA) on the treatment mechanism (that is, the conditional distribution of treatment regimen, given the full data), which states that, at each time-point, conditional on the observed past, the treatment at this time-point is conditionally independent of the full data (i.e., the subject's unknown characteristics). Applied to our HIV example, this assumption states that, given the observed past (such as past CD4 T-cell counts, viral mutations, and treatment history), at each time point whether an individual switches to a new antiretroviral therapy does not depend on unobserved characteristics of the individual.

Under SRA (or CAR, in general), the likelihood of the observed data factorizes in two terms: the conditional density of the treatment regimen given the full data (the treatment mechanism), and an observed data partial likelihood defined as the product over time of the conditional distributions of covariates and outcomes, given the observed past (Gill et al. (1997)). From a likelihood perspective, asymptotically optimal inference regarding the MSM-parameters can be based solely on this partial likelihood. Under the SRA and the so-called experimental treatment assignment assumption (ETA), which states that a treatment decision is not a deterministic function of the past, it is possible to identify the treatment-specific counterfactual distributions from the observed data partial likelihood, through the G -computation formula (Robins (2000), Gill and Robins (2001), Yu and van der Laan (2002)). Since the time-dependent covariates can be high-dimensional, implementation of this maximum likelihood estimator can be quite involved. In addition, because the curse of dimensionality requires strong modelling assumptions, the maximum likelihood estimator is generally quite sensitive to model misspecification.

Therefore, estimating function methodology specifically designed to estimate the parameter of interest (instead of estimating the whole observed data density), provides a powerful practical (finite sample) alternative to maximum likelihood estimation (MLE). Double robust inverse probability of treatment weighted (DR-IPTW) estimators based on a general estimating function methodology outlined in van der Laan and Robins (2002) have been developed and studied in detail (Robins (1999), Neugebauer and van der Laan (2005)). For a comprehensive description of the double robust estimation method corresponding with Neugebauer and van der Laan (2005) we refer to van der Laan and Robins (2002) (e.g., Chapter 1, section 1.6). The double robust estimators proposed in Neugebauer and van der Laan (2005) rely on either correct specification of the partial likelihood or on correct specification of the treatment mechanism, and are therefore always more nonparametric than MLE and Inverse Probability of Treatment Weighted (IPTW)-estimation. For finite sample comparisons of the three types of methods (i.e., MLE, IPTW, DR-IPTW) for fitting MSM under a variety of settings, we refer to Yu and van der Laan (2003), Neugebauer and van der Laan (2005), and Bryan et al. (2003).

Data structure and assumptions for the history-adjusted marginal structural model. We now introduce the data structure and assumptions we will use to define the history-adjusted marginal structural model. In order to avoid the need for the introduction of technical (and possibly quite involved) measurability conditions, we will assume that all random variables are discrete valued. We note, however, that all our formulas have valid representations for continuous-valued random variables.

Given a time-dependent process $X(t)$, we will adopt the notation $\bar{X}(t) = (X(s) : s \leq t)$, and $X(t_1, t_2) = (X(s) : t_1 \leq s \leq t_2)$. Let $\bar{A} = \bar{A}(K) =$

$(A(0), \dots, A(K))$ denote the multivariate treatment regimen, or the treatment history through time K , and let $\mathcal{A}$ denote the set of possible treatment regimens over the course of the study. We will also use the notation $\mathcal{A}(m)$ for all possible outcomes for $\bar{A}(m)$, $m = 0, \dots, K$: thus $\mathcal{A} = \mathcal{A}(K)$.

For each $\bar{a} \in \mathcal{A}$, we assume the existence of a counterfactual process $(L_{\bar{a}}(j) : j = 0, \dots, K + 1)$, which denotes the treatment-specific process one would have observed if the subject had followed treatment regimen $\bar{A} = \bar{a}$. It is assumed that $L_{\bar{a}}(j) = L_{\bar{a}(j-1)}(j)$, the temporal ordering assumption. In other words, the values of the covariates are not affected by treatment that occurs after they are measured (note that we assume that $A(j)$ is measured after $L(j)$). This defines the treatment-specific outcomes $Y_{\bar{a}} = L_{\bar{a}}(K + 1)$. It is further assumed that the observed $L(j)$ equals the treatment-specific $L_{\bar{a}}(j)$ corresponding with the treatment the subject actually took: that is, $L(j) = L_{\bar{A}(j-1)}(j)$, $j = 0, \dots, K + 1$. Under this consistency assumption (CA), we have the following relation between the collection of counterfactual data structures and the observed data:

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

In the special single-time-point case $O = (W, A, Y)$ we have $O = (W, A, Y_A)$, where Y_a represents the treatment-specific counterfactual outcome of Y .

Applied to our HIV example, $\bar{A} = \bar{A}(K) = (A(0), \dots, A(K))$ denotes a subject's antiretroviral history at each time point (as above, $A(j) = 1$ for each time-point that the subject remains on his/her original therapy, and jumps to 0 when the subject switches to a new regimen). We remind the reader that $A^*(j) \equiv \sum_{l=0}^j A(l)$ is the time the subject has spent on his/her original non-suppressive therapy by time j and thus represents a summary of $\bar{A}(j)$, and $A^*(K)$ represents the total time a subject spends on the original failing therapy during the course of the study, or equivalently, the exposure time. Note that in our example, the set of possible treatment histories will not include any treatment histories in which subjects switch treatment more than once. For each possible exposure time, $(L_{\bar{a}}(j) : j = 0, \dots, K + 1)$ denotes the counterfactual CD4 T-cell count and viral mutation trajectory that would have been observed if the subject had followed that exposure time $\bar{A} = \bar{a}$. $Y_{\bar{a}} = L_{\bar{a}}(K + 1)$ denotes the counterfactual CD4 T-cell count observed at the end of the study under the exposure time indicated by $\bar{a}$. Under the temporal ordering and consistency assumptions, the observed data for a subject can be represented as his/her actual exposure time and the counterfactual CD4 T-cell count and viral mutation process corresponding to that exposure time.

This representation (1) of the observed data as a known function of the full data structure $X = (L_{\bar{a}} : \bar{a} \in \mathcal{A})$, and the missingness variable $\bar{A}$, shows that the distribution of O is indexed by the distribution X and the conditional distribution of $\bar{A}$, given X , which we will denote with $F_{X,0}$ and $G_0(\cdot | X)$,

respectively. The actual probability distribution (i.e., density) of $\bar{A}$, given X , will be denoted with $g_0(\cdot | X)$.

We assume the sequential randomization assumption (SRA) which states that $A(m)$ is conditionally independent of X , given $\bar{A}(m-1), \bar{L}(m)$, $m = 0, \dots, K$:

$$g_0(A(m) | \bar{A}(m-1), X) = g_0(A(m) | \bar{A}(m-1), \bar{L}(m)), \quad m = 0, \dots, K. \quad (2)$$

We also assume the experimental treatment assignment assumption (ETA), which states that:

$$g_0(a(m) | \bar{A}(m-1), \bar{L}(m)) > 0$$

for all possible $a(m) \in \{a_*(m) : (\bar{A}(m-1), a_*(m)) \in \mathcal{A}(m)\}$ (3)

This assumption states that, for each time-point in the study, conditional on any possible observed past, each possible treatment at time m which is compatible with a past $\bar{A}(m-1)$ (i.e. is part of a possible longitudinal treatment regimen) needs to have positive probability.

Density of observed data. It follows immediately that, under the SRA, the density of O factorizes as

$$p_{F_{X_0}, g_0}(O) = g_0(\bar{A} | X) \prod_{j=0}^{K+1} p_0(L(j) | \bar{L}(j-1), \bar{A}(j-1)), \quad (4)$$

where we remind the reader that $Y = L(K+1)$, and the conditional distribution of $\bar{A}$, given X , is given by (by SRA)

$$g_0(\bar{A} | X) = \prod_{j=1}^K g(A(j) | \bar{A}(j-1), \bar{L}(j)).$$

We refer to g_0 as the treatment mechanism. Under the SRA and ETA, we have

$$\begin{aligned} p_0(L(j) = l(j) | \bar{L}(j-1) = \bar{l}(j-1), \bar{A}(j-1) = \bar{a}(j-1)) \\ = P_{F_{X_0}}(L_{\bar{a}}(j) = l(j) | \bar{L}_{\bar{a}}(j-1) = \bar{l}(j-1)) \\ \equiv Q_{F_{X_0}, j}(l(j) | \bar{l}(j-1), \bar{a}(j-1)). \end{aligned} \quad (5)$$

That is, the second product $Q_{F_{X_0}} = \prod_{j=0}^{K+1} Q_{F_{X_0}, j}$ in the representation (4) of the observed data density, evaluated at $\bar{A} = \bar{a}$, yields the density of the counterfactual process $L_{\bar{a}}$. This latter result is typically referred to as the G -computation formula for the distribution of $L_{\bar{a}}$ (Robins (2000)). Generalizations of the G -computation formula to continuous valued processes are provided in (Gill and Robins (2001), Yu and van der Laan (2002)).

Definition of history-adjusted treatment-specific mean outcome.

We now proceed with defining the HA-MSM causal parameter of interest as a parameter of the distribution of the underlying $(\bar{A}, X)$.

Let $\underline{a}(j) \equiv (a(j), \dots, a(K))$ denote a future treatment regimen, starting at time j , $j = 0, \dots, K$. Let $S(j) \subset L(j)$ be a sub-vector of $L(j)$, $j = 0, \dots, K$, and let $\bar{S}_{\bar{a}} \subset \bar{L}_{\bar{a}}$ denote the corresponding treatment-specific counterfactual process. Let $V(j) \subset (\bar{A}(j-1), \bar{L}(j))$ denote a subset of the observed treatment and covariate history up till time j . For example, $V(j)$ might include viral mutation profile at time j , and whether or not a patient has already switched to a new treatment.

For each time-point $j \in \{0, \dots, K\}$, conditional on an observed subset $(V(j))$ of the (at time j) available past $(\bar{L}(j), \bar{A}(j-1))$, we consider the counterfactual random variables $Y_{\bar{A}(j-1), \underline{a}(j)}$, indexed by a possible future treatment regimen $\underline{a}(j) \in \underline{a}_*(j) : (\bar{A}(j-1), \underline{a}_*(j)) \in \mathcal{A}$. Note that if the set of possible treatments at time j is a function of past treatment up till time j , past treatment must be included in $V(j)$. For example, if only a single change in treatment is allowed over the course of the study, possible future treatment regimens at time j will depend on whether a subject has already switched treatment before time j , and at least $A(j-1)$ must be included in $V(j)$. In general, $V(j)$ is a user-supplied subset of the observed past which plays the role of the adjustment variable in our HA-MSM and should be implied by the scientific question of interest one wishes to answer.

In this article we are concerned with inference for the following parameter of the distribution of $(\bar{A}, X)$:

$$\theta_0(j, \underline{a}(j) \mid V(j)) \equiv E_0(Y_{\bar{A}(j-1), \underline{a}(j)} \mid V(j)), \quad j = 0, \dots, K. \quad (6)$$

In words, we are interested in the expected counterfactual outcome if the study population were to follow their observed treatment history up till time j , followed by future treatment history $\underline{a}$, conditional on a subset of possibly time-varying covariates, which may include treatment history, measured up till time j .

We make the convention that $\bar{A}(-1)$ is empty ($j = 0$ refers to the time when our treatment of interest starts), so that $\theta_0(0, \underline{a}(0) \mid \bar{S}(0)) = E(Y_{\underline{a}(0)} \mid \bar{S}(0))$ is the conditional mean of the treatment-specific mean outcome $Y_{\bar{a}}$ conditional on a subset $S(0)$ of the baseline covariates, as modelled in a regular MSM. Note that, in the case that all random variables are discrete valued, given an infinite sample of $(\bar{A}, X)$, for each $j = 0, \dots, K$, one can evaluate $\theta_0(j, \underline{a}(j) \mid v(j))$ as a simple average of outcomes $Y_{\bar{A}(j-1), \underline{a}(j)}$ among all observations satisfying $V(j) = v(j)$. In the continuous case, $\theta_0(j, \underline{a}(j) \mid V(j))$ is defined as a conditional expectation of the random variable $Y_{\bar{A}(j-1), \underline{a}(j)}$, conditional on the random variable $V(j)$.

To clarify both the notation and the interpretation of the parameter of

interest, we rely again on our HIV example. Future antiretroviral treatment after time j , denoted $\underline{a}(j)$, consists of a vector of treatment decisions $(a(j), \dots, a(K))$, where $a(l) = 0$ if subject has switched treatment at or before time l , otherwise $a(l) = 1$. This vector of future treatment decisions exists for each subject beginning at each time-point in the study, $j = 0, \dots, K$. We define $a^*(j)$ similar to $A^*(j)$, where $a^*(j) \equiv \sum_{l=j}^K a(l)$ is the future time (after time-point j) that the subject will spend on his/her original failing therapy before the outcome is measured at time $(K + 1)$, and thus represents a summary of $\underline{a}(j)$. We note that $a^*(j) = 0$ implies that the subject has switched therapies at or before time j . The viral mutation profile at time j (presence or absence of mutations conferring altered viral fitness) is denoted $S(j)$, a subset of the full covariate history measured over time, $L(j)$, and the counterfactual mutation profile that would be observed under treatment history $\bar{a}$ is denoted $S_{\bar{a}}$. For each time-point j , we are interested in the mean final CD4 T-cell count among the population who have not yet switched therapy ($A(j - 1) = 1$) if this population had followed their observed treatment history up till time j (not switching), and then a specified treatment history after time j , conditional on treatment history and viral mutation profile at time j (in our example, $V(j) = \bar{A}(j - 1), S(j)$). Note that we condition on not having already switched therapy as this is the only population for which future treatment changes can occur, and thus the only population for which counterfactuals indexed by future treatment exist.

In Section 4 we show that, under the consistency assumption (1), sequential randomization assumption (2), and experimental treatment assignment assumption (3), θ_0 is indeed identifiable from the observed data distribution. However, multiple time-points and possible treatment regimens ensure that θ_0 is very high-dimensional in most cases. Even if we assume that the treatment mechanism $g_0(\bar{A} \mid X)$ is known, as is the case in a sequentially randomized trial, by the curse of dimensionality, only extremely large sample sizes will allow us to estimate θ_0 with reasonable precision. To make estimation possible using finite samples, further modelling assumptions on $\theta_0(j, \underline{a}(j) \mid V(j))$ are typically necessary to reduce the dimensions of the parameter being estimated.

Modelling the history-adjusted treatment specific mean outcome. Specifically, we will assume a parametrization

$$\theta_0(j, \underline{a}(j) \mid V(j)) = m_{\beta_0}(j, \underline{a}(j) \mid V(j)), \quad (7)$$

$j = 0, \dots, K$, where m_{β} is a known function parameterized by a parameter β , and β_0 denotes the true parameter value.

In other words, to reduce the dimensions of the problem we assume some model for the dependence of the mean counterfactual outcome on time, future treatment regimens, and the observed history of a subset of (possibly time-varying) covariates. This can be contrasted with the standard MSM approach,

in which we assume some model for the dependence of the mean counterfactual outcome on treatment history and (not time-varying) covariates at baseline. We note that, if one models $m_\beta(0, \underline{a}(0) \mid V(0))$, and leaves the model for the other time-points unspecified, then our HA-MSM reduces to a regular MSM. This shows that our class of HA-MSM includes the regular MSM as a special case.

We note that, contrary to a standard MSM, which models a parameter of the full data (i.e., counterfactual) distribution, a HA-MSM models a parameter of the joint distribution of the full data *and* the treatment regimen. This is due to the fact that the way treatment $\bar{A}(j-1)$ has been assigned in response to the time-dependent confounders $\bar{L}(j-1)$ does typically affect the resulting subpopulation defined by strata of $V(j) = v(j)$, and thereby its corresponding mean of counterfactual outcomes $Y_{\bar{A}(j-1), \underline{a}(j)}$. To illustrate, we return to our example. While the effect of switching therapy at time t is defined using the full (counterfactual) data, we evaluate this effect within a specific subpopulation. Membership in this subgroup, and thus expected outcome in this subgroup, can be affected by time-dependent confounders. For example, since changes in CD4 T-cell count over time may affect when a patient is switched to a new therapy, the subpopulation that has not switched therapy by time t may include a disproportionately high number of individuals who have maintained high CD4 T-cell counts up till that time-point. Thus, the expected final CD4 T-cell count in this subgroup may be higher than the expected final CD4 T-cell count among a random sample of the cohort if they had switched therapy at time t . We note, however, that if $V(j) = (\bar{A}(j-1), \bar{L}(j))$ equals the complete observed past, then, by SRA, $\theta_0(j, \underline{a}(j) \mid V(j) = (\bar{a}^*(j-1), \bar{l}(j))) = E(Y_{\bar{a}^*(j-1), \underline{a}(j)} \mid \bar{L}_{\bar{a}^*}(j))$ reduces to a parameter of the full data distribution only.

A possible example of a HA-MSM model with $V(j) = (\bar{A}(j-1), S(j))$ and $S(j)$ being a univariate time-dependent covariate is given by

$$m_\beta(j, \underline{a}(j) \mid \bar{A}(j-1), \bar{S}(j)) = \beta_0 + \beta_1 S(j) + \beta_2 A^*(j-1) + \beta_3 a^*(j), \quad (8)$$

where we remind the reader that $A^*(j-1)$ is a summary measure of $\bar{A}(j-1)$ and $a^*(j)$ is a summary of $\underline{a}(j)$. Note that time elapsed since the beginning of the study (j) can also be incorporated explicitly as an additional covariate. Applied to the HIV example, the above model assumes that the conditional expectation of the counterfactual CD4 T-cell count given observed viral mutation history and treatment history is a linear function of history on non-suppressive therapy, the viral mutation profile at time j , and the future time (after time j) on non-suppressive therapy. We note that, in the HIV example, the sample model presented (8) would apply only to individuals who have not already switched therapy by time j (in other words, it will incorporate an indicator that $A(j-1) = 1$), since it is only among this subpopulation that counterfac-

tual outcomes indexed by future treatment decisions exist. We note further that β_2 cannot be interpreted as the causal effect of exposure through $j - 1$ on the outcome; rather, it is merely an associational parameter representing the association of previous exposure with the counterfactual CD4 T-cell count $Y_{\bar{A}(j-1), \underline{a}(j)}$.

Further modelling assumptions in observational studies - double robust model. In a sequentially randomized trial our proposed IPTW and Double Robust-IPTW (DR-IPTW) estimators (using the, by design, known treatment mechanism g_0) are consistent and asymptotically linear estimators of the causal parameter β_0 under weak regularity conditions. However, in an observational study our IPTW-estimators require a consistent estimate of the unknown treatment mechanism $g_0(\bar{A} | X)$, and our DR-IPTW estimators require either a consistent estimate of the unknown g_0 or a consistent estimate of the (always) unknown partial likelihood $Q_{F_{X_0}} = (Q_{F_{X_0,j}} : j = 0, \dots, K + 1)$.

Let $\mathcal{G}$ denote the working model for $g_0(\bar{A} | X)$, and let $\mathcal{Q}$ denote the working model for the partial likelihood $Q_{F_{X_0}}$. Note that these working models also imply models for each j -specific conditional distribution $g_0(A(j) | \bar{A}(j-1), X)$, $j = 0, \dots, K$, and $Q_{F_{X_0,j}}$, $j = 0, \dots, K + 1$. Our IPTW estimator only relies on correct specification of the model for $g_0(\bar{A} | X)$, while our double robust IPTW estimator relies on correct specification of either the model $\mathcal{G}$ or the model $\mathcal{Q}$. We assume that we estimate the unknown parameters in these working models with a sensible maximum likelihood procedure. For example, one could estimate g_0 and $Q_{F_{X_0}}$ with the maximum likelihood estimators

$$g_n = \arg \max_{g \in \mathcal{G}} \prod_{i=1}^n g(\bar{A}_i | X_i)$$

$$Q_n = \arg \max_{Q \in \mathcal{Q}} \prod_{i=1}^n \prod_{j=0}^{K+1} Q_j(L_i(j) | \bar{L}_i(j-1), \bar{A}_i(j-1)),$$

but, if necessary, one would use variations involving regularizing the likelihood, and fine tuning methods based on (e.g.) cross-validation.

We conclude this section with some remarks regarding the definition and modelling of the history-adjusted treatment specific mean.

The history-adjusted treatment specific mean under a user supplied treatment mechanism.

The group of subjects in a strata $V(j) = v(j)$ identifies a (self-selected) subgroup of the population depending on the treatment assignment mechanism. Thus the conditional treatment-specific mean adjusted for $V(j)$, as modelled by the HA-MSM, is modelling the effect of a future treatment regimen on mean outcome for sub-populations affected by the actual treatment mechanism, so

that different treatment mechanisms naturally result in different HA-MSM parameters.

As pointed out to us by a referee, one might also define the parameter of interest as the conditional mean of $Y_{\bar{A}(j-1), \underline{a}(j)}$, given $V(j) = v(j)$, for a population which has been subject to another user-supplied treatment assignment mechanism g^* instead of the actual treatment mechanism g in the population we are sampling from. Since our estimating functions are unbiased under $P_{F_X, g}$, our estimating functions multiplied with $g^*(\bar{A} | X)/g(\bar{A} | X)$ are unbiased under P_{F_X, g^*} ; this operation can be viewed simply as additional IPTW weighting. Use of IPTW-weights g^*/g allows us to immediately map our class of IPTW or DR-IPTW estimating functions into corresponding estimating functions for the g^* -specific history-adjusted treatment-specific mean. We refer to Section 2.3.4 of Chapter 2 in van der Laan and Robins (2002) for the corresponding DR-IPTW g^* -specific class of estimating functions.

Modelling strategies for the history-adjusted treatment specific mean.

The model for θ_0 adopted can be arbitrarily flexible. One can augment the main term model above with interactions between $S(j)$ and $\bar{A}(j-1)$, $S(j)$ and $\underline{a}(j)$ and so on, and introduce polynomial interactions between time j and the other variables. Alternatively, one can assume for each time-point j a separate model with separate coefficients β^j . In this case $\beta = (\beta^0, \dots, \beta^K)$ represents a concatenation of all these j -specific coefficient vectors, and our estimating equations for β correspond with separate estimating equations for β^j , $j = 0, \dots, K$, which can be solved independently from each other.

An important motivation for causal modelling is to provide a summary measure for a true nonparametric causal parameter. The true nonparametric causal parameter may not be path-wise differentiable, and thereby not estimable at the regular rate $1/\sqrt{n}$. A causal model can be viewed as a projection of the true nonparametric causal parameter on a model for the parameter that can be estimated with reasonable precision. Therefore, a parametric HA-MSM can be viewed as a way of defining a smooth parameter of the nonparametric parameter $E(Y_{\bar{A}(j-1), \underline{a}(j)} | V(j))$ in the bigger model that makes no assumptions on this parameter. This is comparable to using least squares linear regression to obtain an estimate of a projection of the true regression on the linear working model. In this case, choice of the parametric form of the HA-MSM and choice among the class of estimating functions defines a particular projection, where the choice of estimating function corresponds with selecting the weights in a Hilbert space norm (Neugebauer and van der Laan (2003)).

However, if the goal is to estimate the true history-adjusted treatment-

specific mean, rather than its projection on a specified model, model selection should be employed to select among various candidate models based on the appropriate trade-off between bias and variance. A general cross-validation methodology has been developed in (van der Laan and Dudoit (2003)), and can be used for this purpose. Specifically, in Wang and van der Laan (2004) we apply this unified cross-validation methodology to carry out model selection in point-treatment marginal structural models. In the future, we plan to apply such data-adaptive learning methods to HA-MSM's as well.

3.1 A statically-optimal dynamic treatment regimen.

We remind the reader that a dynamic treatment regimen is a collection of functions $(V(j)) \rightarrow d_j(V(j))$ of a subset $(V(j))$ of the available history $(\bar{A}(j-1), \bar{L}(j))$, indexed by the time-point j . The output of this j -specific function d_j represents a particular treatment decision at time j , $j = 0, \dots, K$. In other words, a dynamic treatment regimen gives a recommended treatment decision at each time-point, based on treatment history and a subset of covariates measured up to the time-point. We will represent such a dynamic treatment regimen with a vector $d = (d_0, \dots, d_K)$, where the j -th component equals the function d_j , $j = 0, \dots, K$. Suppose that one aims to find dynamic treatment regimens which maximize the mean outcome: that is, one wishes to maximize $E[Y_d]$, where Y_d is the outcome we would have observed on the whole population if we would have assigned dynamic treatment regimen d to each subject.

An example of a dynamic treatment regimen is a decision rule about when to switch to a new antiretroviral treatment, given time spent on current non-suppressive treatment, current viral mutation profile, and possibly additional covariates (such as the drug class in the patient's current regimen, the patient's immunologic status, etc.). We are interested in how different decision rules, based on these covariates, would affect the final mean CD4 T-cell count.

We note that static treatment regimens $(a(0), \dots, a(K))$ represent a special case of a dynamic treatment regimen, in which the decision rule d_j which yields a treatment decision at each time-point j is simply a constant for each time-point, and does not depend on $V(j)$. (Of course, $a(j)$ can still vary over time).

We will show that the history-adjusted treatment-specific mean outcome yields an interesting optimal dynamic treatment regimen. The type of optimality will be defined below. First, we define a history-adjusted optimal static treatment regimen for the future as

$$\underline{a}_0(j \mid V(j)) \equiv \arg \max_{\underline{a}(j) \in \mathcal{A}(\bar{A}(j-1))} \theta_0(j, \underline{a}(j) \mid V(j)),$$

where $\mathcal{A}(\bar{A}(j-1)) \equiv \{\underline{a}_* : (\bar{A}(j-1), \underline{a}_*(j)) \in \mathcal{A}\}$ denotes all future treatment regimens compatible with past treatment $\bar{A}(j-1)$, $j = 0, \dots, K$. In words,

the history-adjusted optimal static treatment regimen at time j is the future treatment regimen (after time j), which maximizes the mean counterfactual outcome of the population within subgroups defined by a subset of observed covariate history (where covariate history can include treatment history).

We now define the following dynamic treatment regimen:

$$d_{0j}(V(j)) \equiv a_0(j \mid V(j)), j = 0, \dots, K.$$

That is, $d_{0j}(V(j))$ equals the first treatment action of the history- (at time j) adjusted optimal static treatment regimen $\underline{a}_0(j \mid V(j))$. We note that the dynamic treatment regimen d_0 is a parameter, which happens to be identified by the parameter θ_0 . Applying this dynamic treatment regimen to a subject over time involves calculating the subject's optimal static future treatment regimen $\underline{a}(j \mid V(j))$ given a subset of his/her history ($V(j)$) at each time-point j , and carrying out the initial treatment action of this latter regimen, $j = 0, \dots, K$.

We refer to this dynamic treatment regimen d_0 as a *statically-optimal dynamic treatment regimen*, to distinguish it from an optimal dynamic treatment regimen as modelled in (Murphy (2003), Robins (2003a)). The following properties of d_0 follow directly from its definition.

- For the single time-point case $K = 0$ (so that the data is $O = (L(0), A(0), Y)$), we have that the dynamic treatment regimen $d_0(V(0))$ is optimal among all dynamic treatments $V(0) \rightarrow d(V(0))$ which are only functions of $V(0) \subset L(0)$. Thus, in the single time-point setting, the statically optimal dynamic treatment regimen is also the best dynamic treatment regimen based on the same covariates.
- In general, at any time-point j , given any observed past ($V(j)$), the future dynamic treatment regimen $\underline{d}_{0j} \equiv (d_{0j}, \dots, d_{0K})$ is better than the $V(j)$ -adjusted optimal future static treatment regimen $\underline{a}_0(j \mid V(j))$. Formally, for $j = 0, \dots, K$, we have

$$E_0(Y_{\bar{A}(j-1), \underline{d}_{0j}} \mid V(j)) \geq \max_{\underline{a}(j)} E_0(Y_{\bar{A}(j-1), \underline{a}(j)} \mid V(j)).$$

In other words, updating the optimal history-adjusted static treatment regimen at each time-point after j based on new covariate values, and modifying treatment decisions accordingly, results in a better final outcome than simply following the optimal history-adjusted static treatment regimen at time j .

Finally, we note that when HA-MSM are used to model an outcome at a fixed interval m in the future rather than an outcome at a fixed time point such as the end of the study (e.g. $E[Y(j+m)]$ rather than $E[Y(K+1)]$), the

statically optimal dynamic treatment regimen identified by HA-MSM aims to optimize a different quantity than the optimal dynamic treatment regimens of Murphy and Robins.

A concrete example to show the non-optimality of the history adjusted statically-optimal dynamic treatment regimen.

Consider a simple longitudinal data structure with four binary variables $A(0), L(1), A(1), Y$. $A(0)$ is an indicator of being treated with a drug at time 0, $L(1)$ is an indicator of developing a comorbid condition at time 1 that will lead to fatal adverse effects if the drug is administered in the future, $A(1)$ is an indicator of being treated with the drug at time 1, and Y is the outcome, an indicator of survival. We propose a setting in which a good treatment decision at time 1 is heavily based on $L(1)$; among patients with $L(1)=1$, the best treatment decision $A(1)$ is to stop treatment (due to the adverse effects), while among patients with $L(1) = 0$ the best decision is to continue treatment. In this scenario, the initial component of the optimal static treatment regimen at time 0 can equal “do not treat” ($A(0) = 0$), because the static regimen does not allow the option of being able to respond to $L(1)$ in the future. As a consequence, following the first component “do not treat” of the statically-optimal treatment strategy at time 0 would deviate from the truly optimal dynamic treatment regimen.

Suppose that $P(L(1) = 1 \mid A(0)) = A(0) * \alpha_1 + (1 - A(0)) * \alpha_2$, That is, α_1 is the percentage of subjects who develop the comorbid condition among the subjects who were treated (i.e., $A(0) = 1$), while α_2 is the percentage of subjects who develop the comorbid condition among the untreated. Let $P(Y = 1 \mid L(1) = 1, A(0), A(1)) = (1 - A(1)) * A(0) * \beta_1$, and $P(Y = 1 \mid L(1) = 0, A(0), A(1)) = A(0) * \beta_2 + A(1) * \beta_3$. That is, β_1 is the percentage of subjects who survive among the subjects who were treated at time 0, developed comorbidity at time 1, and were taken off of treatment. A subject who receives treatment after developing comorbidity always dies, as a result of adverse effects, as does a subject who is never treated. Finally, β_2 and β_3 represent contributions to the survival percentage due to treatments $A(0)$ and $A(1)$, respectively, among subjects who do not develop comorbidity. Given values for these parameters, for each of the four static treatment regimens $\bar{a} \in \{0, 1\}^2$, we can calculate $P(Y_{\bar{a}} = 1)$ from the G -computation formula $P(Y_{\bar{a}} = 1, L_{\bar{a}}(1) = l_1) = P(Y = 1 \mid L(1) = l_1, A(1) = a(1), A(0) = a(0))P(L(1) = l_1 \mid A(0) = a(0))$.

It can be verified from the G -computation formula for $P(Y_d = 1)$ (i.e., set $A(1) = d(L(1))$ in the above formula) for a dynamic treatment regimen, that the optimal dynamic treatment regimen is given by $A(0) = 1$ and $A(1) =$

$I(L(1) = 0)$: that is, treat at time 0, and take the subject off treatment if he/she develops adverse effects at time 1. To provide a concrete example in which the statically-optimal dynamic treatment regimen (which consists of taking the first action of the optimal static treatment regimen identified at each time-point) results in an outcome inferior to the optimal dynamic treatment regimen, it remains to find values of the parameters so that the first component of the optimal static treatment regimen at time 0, given by $\arg \max_{\bar{a}} P(Y_{\bar{a}} = 1)$, equals zero.

Consider the values: $\alpha_1 = 0.7$, $\alpha_2 = 0.1$, $\beta_1 = 0.2$, $\beta_2 = 0.6$, $\beta_3 = 0.4$. These parameter values result in the following survival probabilities:

$$\begin{aligned} P(Y_{01} = 1) &= 0.36 \\ P(Y_{dopt} = 1) &= 0.44 \\ P(Y_{d*} = 1) &= 0.36, \end{aligned}$$

where d^* denotes the statically-optimal dynamic treatment regimen, as identified by the HA-MSM. Given these parameter values, then, the statically-optimal dynamic treatment regimen results in a lower probability of survival than the optimal dynamic treatment regimen. For convenience, we provide the reader with the R-code for the calculation of these counterfactual probabilities in an Appendix.

4 Identifiability of the history-adjusted treatment-specific mean outcome.

Consider our observed data model defined by the CA (1), the SRA (2), and the ETA (3). In this section we will establish an identifiability result of general parameters of the conditional distribution of $Y_{\bar{A}(j-1), \underline{a}(j)}$, given $(V(j))$, $j = 0, \dots, K$. This generalizes the G -computation formula of Robins (Robins (2000)) concerning the identifiability of the distribution of $Y_{\bar{a}}$, possibly conditional on baseline covariates, which corresponds with our result at $j = 0$.

The fundamental idea is to apply the G -computation formula of Robins (2000) at a fixed time point j , treating $W(j) \equiv (\bar{A}(j-1), \bar{L}(j))$ as baseline covariates, and the future treatment and future covariate process as the time-dependent data structure with $K-j$ time-points at which treatment changes, and a final outcome at time $K-j+1$, where we will condition on $V(j) = v(j)$. The validity of the j -specific G -computation formula relies on the j -specific CA in terms of the full data being the collection of j -specific counterfactuals $L_{\bar{A}(j-1), \underline{a}(j)}$ indexed by future treatment regimens $\underline{a}(j) = (a(j), \dots, a(K))$ (whose index set may depend on $V(j)$ through $\bar{A}(j-1)$), the j -specific SRA, and the j -specific ETA, where these j -specific assumptions are defined in the

Appendix: see (13), (14), and (15). The latter three assumptions (for all $V(j)$) and the j -specific model assumption $\theta_0(j, \underline{a}(j) \mid V(j)) = m_\beta(j, \underline{a}(j) \mid V(j))$ (for all $\underline{a}(j)$ and $V(j)$) define a j -specific causal inference model $\mathcal{M}_j$ for the distribution of the observed data. We will refer to the latter model as the j -specific MSM for the observed data distribution.

If the set of possible future treatments beginning at time j depends on treatment history up till time j , then $V(j)$ must include $\bar{A}(j-1)$ in order that conditioning on $V(j)$ fixes $\bar{A}(j-1)$ so that the set $\mathcal{A}(\bar{A}(j-1))$ of future treatment regimens is a fixed set; obviously, this is a necessity for the G -computation formula to represent a sensible statement.

In the Appendix we prove that our HA-MSM model $\mathcal{M}$ for the observed data structure implied by the consistency assumption (1), the SRA (2), the ETA (3), and the model (7) on the HA-MSM parameter equals the intersection of the j -specific MSM $\mathcal{M}_j$, $j = 0, \dots, K$.

4.1 The generalized G -computation formula.

This allows us now to apply the G -computation formula of Robins (2000) to the j -specific causal inference model for the observed data structure O implied by the j -specific CA (13), j -specific SRA (14), and the j -specific ETA (15), conditional on baseline covariates $V(j) = v(j)$, for each $j = 0, \dots, K$.

This proves the following result.

Theorem 1 (Identifiability) *Let $j \in \{0, \dots, K\}$ be given. Define $V(j) \subset (\bar{A}(j-1), \bar{L}(j))$. Let $W(j) \equiv (\bar{A}(j-1), \bar{L}(j))$. We adopt the notation $\bar{L}(j, m) = (L(j), \dots, L(m))$ for $m \geq j$. Consider our causal inference model for the observed data structure O defined by the sole assumptions (1), (2) and (3).*

We have that

$$\begin{aligned} P(W(j) = w(j), \bar{L}_{\bar{A}(j-1), \underline{a}(j)}(j+1, K+1) = \bar{l}(j+1, K+1) \mid V(j) = v(j)) \\ = P(w(j) \mid v(j)) \prod_{m=j+1}^{K+1} P(l(m) \mid \bar{a}(j, m-1), \bar{l}(j+1, m-1), w(j)), \end{aligned} \quad (9)$$

where we used short-hand notations $P(w(j) \mid v(j))$ for the conditional probability $P(W(j) = w(j) \mid V(j) = v(j))$, and $P(l(m) \mid \bar{a}(j, m-1), \bar{l}(j+1, m-1), w(j))$ for the conditional probability $P(L(m) = l(m) \mid \bar{A}(j, m-1) = \bar{a}(j, m-1), \bar{L}(j+1, m-1) = \bar{l}(j+1, m-1), W(j) = w(j))$, $m = 1, \dots, K+1$.

The generalized G -computation estimator. This identifiability result (1) provides us with a likelihood-based approach for estimation of (and inference for) θ_0 , analogue to the G -computation estimation procedures proposed to estimate the parameters of a MSM (see e.g., van der Laan and Robins (2002),

Neugebauer and van der Laan (2005), Yu and van der Laan (2003)). To be specific, let $j = 0, \dots, K$ be given. Firstly, we note that a fit of the conditional distribution of $W(j)$, given $V(j)$, and the MLE Q_n based on the model $\mathcal{Q}$ provides us with a fit of the right-hand side in (9). Of course, if $V(j) = (\bar{L}(j), \bar{A}(j-1))$ then the conditional distribution $p(w(j) | v(j))$ cancels out in (9). Secondly, at a fixed choice $\underline{a}(j)$, the right-hand side in (9) suggests immediately a sequential data generating mechanism for drawing $W(j)$, $L(j+1)$, $\dots$, $L(K+1)$. Such a Monte-Carlo simulation provides us with arbitrarily many (approximate) realizations of $(W(j), \bar{L}_{\bar{A}(j-1), \underline{a}(j)}(j+1, K+1))$, given $V(j) = v(j)$, which, in particular, results in many realizations of $W(j)$, $Y_{\bar{A}(j-1), \underline{a}(j)}$, given $V(j) = v(j)$. We can carry out this Monte-Carlo simulation procedure for a set of future treatment regimens $\underline{a}(j)$, for each time-point j , and for a set of $v(j)$ -values. Finally, one fits the HA-MSM based on this pooled (across time-points, across future treatment regimens, and across $v(j)$ -values) data set.

5 IPTW-estimators.

We will first derive the IPTW-estimating functions for the parameters in the j -specific marginal structural model $\mathcal{M}_j$, which assumes $E(Y_{\bar{A}(j-1), \underline{a}(j)} | V(j)) = m_{\beta}(j, \underline{a}(j) | V(j))$, at a fixed time point j , treating the past $W(j) = (\bar{A}(j-1), \bar{L}(j))$ as baseline covariates, the future treatment and future covariate process as the time-dependent data structure with time-points $t = j, \dots, K$ at which treatment changes, and a final outcome at time $t = K+1$, conditional on $V(j)$. These estimating functions are presented in van der Laan and Robins (2002) (page 319-322). Subsequently, we take the linear combination of these j -specific classes of estimating functions as our class of estimating functions for β_0 .

5.1 The j -specific IPTW-estimating functions.

The class of j -specific IPTW-estimating functions for the j -specific causal inference model $\mathcal{M}_j$ for the observed data density has been derived in the literature (van der Laan and Robins (2002), at formula (6.9), page 320)

$$\begin{aligned} IC_0^j(O | g_0, h^j, \beta_0) &\equiv \frac{h^j(\underline{A}(j), V(j), \bar{A}(j-1))}{g_0(\underline{A}(j) | X^j, \bar{A}(j-1))} (Y_{\underline{A}(j)}^j - m_{\beta_0}(j, \underline{a}(j) | V(j))) \\ &= \frac{h^j(\underline{A}(j), V(j), \bar{A}(j-1))}{g_0(\underline{A}(j) | X, \bar{A}(j-1))} (Y - m_{\beta_0}(j, \underline{A}(j) | V(j))), \end{aligned}$$

where h^j can be any function of $\underline{A}(j), V(j), \bar{A}(j-1)$. Here

$$g_0(\underline{A}(j) | X, \bar{A}(j-1)) = \prod_{l=j}^K g_0(A(l) | \bar{A}(l-1), L(l)).$$

Unbiased IPTW-estimating functions. It is straightforward to show that, under the j -specific ETA, these IPTW-estimating functions are indeed unbiased at correctly specified treatment mechanism:

$$E_{P_{g_0, Q_0}} IC_0^j(O \mid g_0, h^j, \beta_0) = 0 \text{ for all } h^j.$$

A typical effective and convenient recommendation for the choice h^j is given by (see formula (6.13), van der Laan and Robins (2002))

$$h^{j*}(\underline{A}(j), V(j), \bar{A}(j-1)) \equiv g^*(\underline{A}(j) \mid V(j), \bar{A}(j-1)) \frac{d}{d\beta} m_\beta(j, \underline{A}(j) \mid V(j)) \Big|_{\beta=\beta_0}.$$

Here

$$g^*(\underline{A}(j) \mid V(j), \bar{A}(j-1)) \equiv \prod_{l=j}^K g^*(A(l) \mid \bar{A}(l-1), V(j)),$$

where $g^*(A(l) \mid \bar{A}(l-1), V(j))$ can be any conditional density of $A(l)$, given $\bar{A}(l-1), V(j)$ (all that matters for the unbiasedness of the IPTW-estimating functions is that it is only a function of $\bar{A}(l), V(j)$).

5.2 The IPTW-estimating functions.

Our history-adjusted marginal structural model for the observed data equals intersection of all the j -specific marginal structural models $\mathcal{M}_j$ for the observed data. As a consequence, the class of all estimating functions of β_0 equals the sum-space of the j -specific classes of all estimating functions: this sentence can be formalized by replacing “the class of all estimating functions” by “the orthogonal complement of the nuisance tangent space” (see van der Laan and Robins (2002) for a complete description of this general estimating function methodology). In particular, the sum-space of the j -specific IPTW-estimating functions generates a rich class of IPTW estimating functions. This provides us now with our proposed class of IPTW estimating functions for β_0 indexed by a multivariate choice $\vec{h} = (h^1, \dots, h^K)$:

$$IC_0(O \mid g_0, \vec{h}, \beta) \equiv \sum_{j=0}^K IC_0^j(O \mid g_0, h^j, \beta).$$

Note that these estimating functions depend on the unknown treatment mechanism g_0 .

5.3 IPTW estimators.

Given the maximum likelihood estimator g_n of the treatment mechanism, and a possibly data dependent $\vec{h}_n$, we would estimate β_0 with the solution β_n of

the corresponding IPTW-estimating equation in β given by:

$$\begin{aligned} 0 &= \sum_{i=1}^n IC_0(O_i \mid g_n, \bar{h}_n, \beta) \\ &= \sum_{i=1}^n \sum_{j=0}^K \frac{h_n^j(\underline{A}_i(j), V_i(j), \bar{A}_i(j-1))}{g_n(\underline{A}_i(j) \mid X_i, \bar{A}_i(j-1))} (Y_i - m_\beta(j, \underline{A}_i(j) \mid V_i(j))). \end{aligned}$$

If we use the choice h_n^{j*} , then this reduces to the following IPTW-estimating equation in β :

$$0 = \sum_{i=1}^n \sum_{j=0}^K \frac{g_n^*(\underline{A}_i(j) \mid V_i(j), \bar{A}_i(j-1))}{g_n(\underline{A}_i(j) \mid X_i, \bar{A}_i(j-1))} \frac{d}{d\beta} m_\beta(j, \underline{A}_i(j) \mid V_i(j)) (Y_i - m_\beta(j, \underline{A}_i(j) \mid V_i(j))).$$

Fortunately, this estimator can be implemented with standard software. Specifically, it corresponds with the following weighted least squares estimator

$$\beta_n = \arg \min_{\beta} \sum_{i=1}^n \sum_{j=0}^K \{Y_i - m_\beta(j, \underline{A}_i(j) \mid V_i(j))\}^2 w_n(i, j),$$

where the $n \times (K+1)$ -weight-matrix w_n is given by

$$w_n(i, j) = \frac{g_n^*(\underline{A}_i(j) \mid V_i(j), \bar{A}_i(j-1))}{g_n(\underline{A}_i(j) \mid X_i, \bar{A}_i(j-1))}, \quad i = 1, \dots, n, \quad j = 0, \dots, K.$$

For example, if the model $m_\beta(j, \underline{A}_i(j) \mid V_i(j))$ is a generalized linear model, then we can compute β_n by applying the R -function $glm()$ to the pooled sample in which each subject i contributes $K+1$ lines of data with weights $w_n(i, j)$, $j = 0, \dots, K$.

5.4 The j -specific DR-IPTW-estimating functions.

The j -specific DR-IPTW-estimating functions for the j -specific marginal structural model for the observed data are given by (van der Laan and Robins (2002), page 322)

$$\begin{aligned} IC^j(O \mid g, Q_X, h^j, \beta_0) &\equiv IC_0^j(O \mid g, h^j, \beta_0) \\ &- \sum_{l=j}^K E_{Q_X, g}(IC_0^j(O \mid g, h^j, \beta_0) \mid \bar{A}(l), \bar{L}(l)) - E_{Q_X, g}(IC_0^j(O \mid g, h^j, \beta_0) \mid \bar{A}(l-1), \bar{L}(l)), \end{aligned} \quad (10)$$

where h^j can vary arbitrarily.

Double robustness of estimating functions. These estimating functions are double robust in the following sense: if (the possibly misspecified) g_1 satisfies the ETA assumption (3), then

$$E_{P_{Q_0, g_0}} IC^j(O \mid g_1, Q_1, h^j, \beta_0) = 0 \text{ if } g_1 = g_0 \text{ or } Q_1 = Q_0. \quad (11)$$

This result does not exactly correspond with the double robustness stated in van der Laan and Robins (2002) for the j -specific double robust estimating functions. This is due to the fact that in the j -specific marginal structural model for the observed data, the marginal distribution of $W(j)$ is considered part of the full-data distribution, while in our original model, this marginal distribution depends on g_0 and Q_{X0} . However, due to the fact that the conditional expectation $E_{Q_{X,g}}(IC_0^j(O \mid g, h^j, \beta_0) \mid \bar{A}(l), \bar{L}(l))$ for $l \geq j$ never requires knowing the marginal distribution of $W(j)$ (one always conditions on it), it happens to be the case that the j -specific double robustness still yields the wished double robustness (11).

5.5 The class of DR-IPTW-estimating functions.

For the j -specific marginal structural model for the observed data the results in (van der Laan and Robins (2002)) show that the orthogonal complement of the nuisance tangent space at $P_{Q_{X,g}}$ is the class of all DR-IPTW estimating functions (10) with all parameters evaluated at $P_{Q_{X,g}}$. Our HA-MSM model for the observed data equals the intersection of all the j -specific MSM models for the observed data. As a consequence, the sum-space of all the j -specific orthogonal complements of the nuisance tangent space at $P_{Q_{X,g}}$ is contained in the orthogonal complement of the nuisance tangent space in our observed data model. (It requires a technical proof to actually establish equality of these two spaces.) Therefore a rich subspace (or the whole space) of the orthogonal complement of the nuisance tangent space at $P_{Q_{X,g}}$ in our model can be represented as:

$$T_{\text{nuis}}^\perp(P_{Q_{X,g}}) \equiv \left\{ \sum_{j=0}^K IC^j(O \mid g, Q_X, h^j, \beta) : (h^0, \dots, h^K) \right\},$$

where β represents the parameter of the data generating distribution $P_{Q_{X,g}}$. This provides us now with our proposed class of DR-IPTW estimating functions for β_0 indexed by a multivariate choice $\vec{h} = (h^1, \dots, h^K)$:

$$IC(O \mid g_0, Q_0, \vec{h}, \beta) \equiv \sum_{j=0}^K IC^j(O \mid g_0, Q_0, h^j, \beta).$$

Note that these estimating functions for β_0 depend on nuisance parameters (g_0, Q_{X0}) .

The previously stated j -specific double robustness (11) immediately carries over to the wished double robustness of these estimating functions: if (the possibly misspecified) g_1 satisfies the ETA assumption (3), then

$$E_{P_{Q_0, g_0}} IC(O \mid g_1, Q_1, \vec{h}, \beta_0) = 0 \text{ if } g_1 = g_0 \text{ or } Q_1 = Q_0.$$

5.6 Double robust IPTW estimators.

Given estimators g_n, Q_n of the treatment mechanism g_0 and partial likelihood $Q_{F_{X0}}$, and possibly data dependent $\vec{h}_n$, a DR-IPTW estimator β_n of β_0 is defined as a solution of the estimating equation in β given by:

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

For the actual implementation of this double robust estimator we refer to van der Laan and Robins (2002); Neugebauer and van der Laan (2005); Yu and van der Laan (2003). Alternative implementations of so called generalized double robust IPTW estimators are presented in Robins (1999) and van der Laan and Robins (2002).

Remark: Different types of double robustness. The double robust estimator requires specification of models for Q_X and g , which raises the question if these models can be or should be chosen compatible with the HA-MSM (7) on θ_0 . Since the Q_X and g play the role of nuisance parameters in estimating functions, and one typically assumes HA-MSM to obtain a smooth summary measure of the true history-adjusted treatment-specific mean, we believe it is appropriate to use models for these nuisance parameters which are non-compatible with the HA-MSM (7) for θ_0 . The double robust estimators in (van der Laan and Robins (2002); Neugebauer and van der Laan (2005); Yu and van der Laan (2003)) are based on separate models for the F_X -part Q_X and g -part of the density of the observed data structure, and use Monte-Carlo simulation to evaluate the conditional expectations defining the projection part of the double robust estimating functions. On the other hand the generalized double robust estimators (van der Laan and Robins (2002); Robins (1999)) use direct estimates of the F_X components of the conditional expectations not compatible with a full data density. This is an important distinction between these two types of double robust estimation procedures. For example, if we view the HA-MSM as a smooth parameter of the true history-adjusted treatment-specific mean defined in the model, without making any assumptions on θ_0 (so that the non-compatibility is a non-issue), then the double robust estimators are truly double robust, while the generalized double robust estimators are not truly double robust (but are double robust from a practical point of view, as argued in Robins and Rotnitzky (2001)).

6 Asymptotic linearity of β_n .

We first discuss the asymptotic linearity properties of the double robust IPTW estimator β_n without presentation of the formal conditions, and discuss infer-

ence for β_0 based on these results. We remind the reader that an estimator β_n is asymptotically linear at data generating distribution P with influence curve $IC(O | P)$ if $\beta_n - \beta = 1/n \sum_{i=1}^n IC(O_i | P) + o_P(1/\sqrt{n})$, which implies, in particular, that $\sqrt{n}(\beta_n - \beta)$ is asymptotically normally distributed with mean zero and covariance matrix $\Sigma = \text{VAR}(IC(O | P))$. As a consequence, asymptotic linearity allows statistical inference based on $\beta_n \approx N(\beta, \Sigma_n/n)$, where $\Sigma_n = 1/n \sum_{i=1}^n \hat{IC}(O_i) \hat{IC}(O_i)^\top$.

Let $P_0 = P_{F_{X_0}, G_0}$ denote the true data generating distribution. Recall the following Hilbert space terminology: $L_0^2(P_0)$ is the Hilbert space of functions of O with finite variance and mean zero endowed with the covariance inner product $\langle v_1, v_2 \rangle_{P_0} \equiv \int v_1 v_2 dP_0$. Let $T_{G,j} = T_{G,j}(P_0)$ denote the tangent space for the parameter $g(\underline{A}(j) | X, \bar{A}(j-1))$ implied by $\mathcal{G}$ at the true data generating distribution P_0 : that is, it is the closure of the linear extension in $L_0^2(P_0)$ of the scores at P_0 from all correctly-specified one-dimensional parametric submodels for the distribution $g(\underline{A}(j) | X, \bar{A}(j-1))$ implied by our model $\mathcal{G}$ for $g(\underline{A} | X, \bar{A}(j-1))$, $j = 0, \dots, K$.

In our causal model for the observed data distribution implied by Assumptions CA (1), SRA (2), ETA (3), and HA-MSM (7), it is possible to derive the following double robust asymptotic linearity result for β_n , where its influence curve depends on the correctness of the working models $\mathcal{G}$ for g_0 and $\mathcal{Q}$ for the partial likelihood $Q_0 = Q_{F_{X_0}}$. The derivation of these asymptotic linearity results requires a few formal empirical process conditions, and rate conditions on g_n, Q_n as estimators of possibly misspecified g^1, Q^1 (analogue to Theorem 2.4 in van der Laan and Robins (2002), and our Theorem 2 in the Appendix below).

Both Correct: If both working models for g_0, Q_0 are correctly specified, then β_n is asymptotically linear with influence curve

$$-d/d\beta E_0 IC(O | g_0, Q_0, \vec{h}, \beta) \big|_{\beta=\beta_0}^{-1} IC(O | g_0, Q_0, \vec{h}, \beta_0).$$

Censoring mechanism Correct: If the working model for g_0 is correct, and Q_n converges to a possibly misspecified Q^1 , then β_n is asymptotically linear with influence curve

$$-c_0^{-1} \sum_{j=0}^K IC^j(O | g_0, Q^1, h^j, \beta_0) - \Pi(IC^j(\cdot | g_0, Q^1, h^j, \beta_0) | T_{G,j}), \quad (12)$$

where $c_0 \equiv \frac{d}{d\beta} \sum_{j=0}^K IC^j(O | g_0, Q^1, h^j, \beta) \big|_{\beta=\beta_0}$, $\Pi(\cdot | T_{G,j})$ is the projection operator onto the sub-Hilbert space $T_{G,j} \subset L_0^2(P_0)$, $T_{G,j}$ is the tangent space spanned by the scores of the treatment mechanism working model for $g_0(\underline{A}(j) | X)$ implied by $\mathcal{G}$.

Double Robust: If either the working model for g_0 is correct, or the working model for Q_0 is correct, then β_n is asymptotically linear.

The influence curve in the latter double robust (i.e., union) model is of the form as presented under the second item, but with the projection operator replaced by a more involved expression (also depending on the asymptotic behavior of Q_n) as defined in Theorem 2.5 in van der Laan and Robins (2002).

Since each j -specific influence curve in the expression (12) of the true influence curve involves a subtraction of a j -specific projection, we suggest that, in practice, it is sensible to use $-c_0^{-1}IC(O \mid g_0, Q^1, \vec{h}, \beta_0)$ as a conservative influence curve.

Thus, if the model $\mathcal{G}$ is correctly specified, then one can (typically) conservatively estimate the covariance matrix of the multivariate normal limiting distribution of $\sqrt{n}(\beta_n - \beta_0)$ as follows:

$$\Sigma_n \equiv \frac{1}{n} \sum_{i=1}^n IC(O_i \mid g_n, Q_n, \vec{h}_n, \beta_n) IC(O_i \mid g_n, Q_n, \vec{h}_n, \beta_n)^\top.$$

Asymptotic inference for β_0 can now be based on $\sqrt{n}(\beta_n - \beta_0) \approx N(0, \Sigma_n)$.

In general, the bootstrap provides us with an asymptotically valid tool for statistical inference in the model where one assumes a correctly specified model $\mathcal{G}$ for g_0 , and in the union model in which one only assumes that either $\mathcal{G}$ or $\mathcal{Q}$ is correctly specified.

7 Discussion.

The HA-MSM presented in this paper represent an important generalization of marginal structural model methodology. Marginal structural models are well-established as powerful tools for causal inference, particularly in the setting of longitudinal data. Unlike traditional methodologies, MSM allow for consistent estimation of causal effects in the presence of time-dependent confounding, and facilitate the estimation of causal effects specified *a priori* by the investigator, rather than driven by the confounding structure in the data. As a generalization of 'standard' MSM, HA-MSM retain these attractive properties.

At the same time, HA-MSM overcome a significant limitation of standard MSM. Unlike standard MSM, HA-MSM allow adjustment by time-dependent covariates, making possible estimation of time-dependent effect modification. In addition, while MSM only provide marginally optimal static treatment regimens, HA-MSM provide statically-optimal dynamic treatment regimens. Thus HA-MSM identify treatment decision rules that can be based on time-dependent covariates, a crucial application in the biomedical sciences.

Though our focus in this article is on HA-MSM modelling the mean, our methods immediately generalize further to HA-MSM modelling other parameters of the conditional distribution of the treatment-specific outcome, given a subset of the observed past. Our results can be generalized to longitudinal data structures allowing for general time-dependent outcome processes (e.g., survival) and informative censoring, and to HA-MSM of general parameters of the conditional distribution of the treatment-specific outcome, given a subset of the observed past.

In particular, the statistical framework and results in this article immediately generalize to the following general longitudinal data structure (see Chapter 6 in van der Laan and Robins (2002)), which allows one to define the outcome Y of interest as a survival time. Suppose that the observed data structure on a randomly sampled subject is defined as $O = (T, \bar{A}(T), \bar{L}(T))$, where T is a possibly random endpoint such as a survival time, $A(t)$ denotes treatment at time t , and $L(t)$ denotes time-dependent measurements of covariates. Let $R(t) = I(T \leq t)$ so that T is identified by this time-dependent process $R(t)$. Suppose that the outcome Y of interest is a function of this data structure: $Y = f(\bar{R}(T), \bar{L}(T))$ for some function f . For example, $Y = T$. For all treatment regimens $\bar{a}$, we define the treatment-specific end-point $T_{\bar{a}}$ and the truncated treatment-specific process $L_{\bar{a}}(t) = L_{\bar{a}}(\min(t, T_{\bar{a}}))$. Let $Y_{\bar{a}} = f(T_{\bar{a}}, \bar{L}_{\bar{a}}(T_{\bar{a}}))$ be the corresponding counterfactual outcome. Let $\bar{R}_{\bar{a}}$ and $\bar{L}_{\bar{a}}$ denote the full sample paths of these counterfactual processes. The full data structure is now defined as the vector

$$X = (\bar{L}_{\bar{a}}, \bar{R}_{\bar{a}} : \bar{a}),$$

which thus includes $Y_{\bar{a}}$. Under the consistency assumption, the observed data structure can now be viewed as a missing data structure on X

$$O = (\bar{A}(T), \bar{L}_{\bar{A}}, \bar{R}_{\bar{A}}).$$

A history-adjusted MSM can be defined similarly with respect to a subset of the past; for example, one could model $E(Y_{\bar{A}(t-1), \underline{a}(t)} | \bar{V}(t))$, where $\bar{V}(t) \subset (\bar{A}(t-1), \bar{L}(t))$. This HA-MSM possibly needs to take into account that an observed past could imply the outcome $Y_{\bar{A}(t-1), \underline{a}(t)}$ (e.g. the subject has already died), which can be handled by using indicator functions of the past in the HA-MSM. The derivation of the t -specific CA, SRA, and ETA follows in an identical manner from the CA, SRA and ETA for our observed data structure. Thus, the observed data model implies the t -specific MSM model, conditional on $\bar{V}(t)$, for the observed data defined by the t -specific CA, SRA, ETA, and MSM, for all t . It remains to map this t -specific MSM to its corresponding t -specific IPTW and DR-IPTW estimating functions, (based on the general IPTW and DR-IPTW mappings as provided in van der Laan and

Robins (2002) and used in this article). Finally, we define the class of IPTW and DR-IPTW estimating functions as the sum-space of the t -specific IPTW and DR-IPTW estimating functions. We also remark that, by allowing $\bar{A}$ to play the role of a general action process possibly including censoring actions, this methodology generalizes to any censored longitudinal data structure (see Chapter 6, van der Laan and Robins (2002)).

To summarize, the estimating function methodology presented in this article involves combining known classes of estimating functions for j -specific marginal structural models for the observed data distribution, and is therefore generic, and avoids reinventing the wheel for each variation which can occur in the choice of HA-MSM and data structure.

In conclusion, we note that important and innovative progress has been made in the literature on semiparametric estimation of optimal dynamic treatment regimens: that is, on methods which aim to model the optimal dynamic treatment regimen of interest, while leaving the nuisance parameters as unspecified as possible. Murphy (2003) models the mean of the difference between the counterfactual outcomes for treatment regimens $\bar{a}(m)$ and $\bar{a}(m-1)$, conditional on the observed past up till time m , and setting future treatment equal to the optimal treatment regimen. She refers to these causal effects as "regret-functions", and provides semi-parametric estimators of the unknown parameters identifying the regret functions. She shows that these regret functions, through backwards solving, yield the optimal dynamic treatment regimen and its corresponding mean. Robins (2003b) proposes a class of regimen-specific structural nested mean models, generalizing the model of Murphy (2003). The regimen-specific structural nested mean models represent a generalization of the original structural nested mean models, inspired by Murphy (2003), and are well suited to the estimation of and inference for an optimal dynamic treatment regimen. Standard structural nested mean models were originally introduced and analyzed by Robins (Robins (1989), Robins (1997), Robins (2000) Robins (1994)). The latter models model the causal effect of treatment at time j (also referred to as a "blip" in treatment), conditional on the observed past, setting future treatment to 0, and naturally lend themselves for *robust* testing of the null hypothesis of no-treatment effect. Murphy (2003)'s class of models, and the optimal regimen-specific structural nested mean models in Robins (2003b), on the other hand, model the causal effect of treatment at time j , conditional on the observed past, *setting future treatment optimal*. Robins (2003a) develops double robust estimators of the unknown parameters in the optimal-regimen specific structural nested mean models. The estimators in the above articles require solving estimating equations which are non-differentiable in the parameters of interest, which makes the practical implementation challenging. Future research still needs to extend the approaches of Robins (Robins (1989), Robins (1997), Robins (2000),

Robins (1994)) and Murphy (2003) to allow adjustment by a *subset* of the observed past, as is accomplished with HA-MSM.

In contrast to the above approaches for modelling and estimating optimal dynamic treatment regimens, HA-MSM provide a methodology for estimating a history-adjusted statically-optimal dynamic treatment regimen that can be implemented using standard software. It will be of interest for future research to compare practically and theoretically statically-optimal dynamic treatment regimens, as immediately obtained from our fitted HA-MSM, and the fully optimal dynamic treatment regimen as obtained with the methods mentioned above.

8 Appendix 1: HA-MSM model equals intersection of time-specific MSM.

The j -specific causal inference framework/model.

In this model we represent the observed data structure O as

$$O = (W(j) \equiv (\bar{A}(j-1), \bar{L}(j)), A(j), L(j+1), \dots, A(K), Y = L(K+1)),$$

and we condition on $V(j)$.

j -specific counterfactuals. Given the previously defined counterfactual processes $X = (L_{\bar{a}} : \bar{a})$, we define $L_{\underline{a}(j)}^j \equiv L_{\bar{A}(j-1), \underline{a}(j)}$ for all possible future treatment regimens $\underline{a}(j) \in \mathcal{A}(\bar{A}(j-1))$. In other words, for each possible future treatment history given an individual's subset of baseline covariates at time j ($V(j)$), $L_{\underline{a}(j)}^j$ is the covariate process that would be observed if an individual followed his/her observed treatment history up till time j , and then followed treatment regimen $\underline{a}$ from time j onwards. Note that for $l > j$, $L_{\underline{a}(j)}^j(l) = L_{\bar{a}(j,l-1)}^j$, where we use the notation $\bar{a}(s, t) = (a(s), \dots, a(t))$ for $t \geq s$. In other words, when $L_{\underline{a}(j)}^j$ is evaluated at a time-point l after j , it corresponds to the counterfactual covariate process under a treatment history consisting of a subject's observed treatment up till time j and specified treatment history $\bar{a}$ from time-point j to point l .

We now define the collection of treatment-specific counterfactual processes as $X^j \equiv (L_{\underline{a}(j)}^j : \underline{a}(j) \in \mathcal{A}(\bar{A}(j-1)))$. Thus, for each time-point j , the full data for an individual consist of the counterfactual covariate processes corresponding to the individual's observed treatment history up till time j and all of the individual's possible future treatment histories after time j (given observed history up to j). Given $V(j)$, $\mathcal{A}(\bar{A}(j-1))$ is a fixed set of future treatment regimens, so that the full data X^j is well defined.

Firstly, we will show that our original CA, SRA, and ETA imply, conditional on $V(j)$, the j -specific CA, j -specific SRA, and the j -specific ETA, for

all $V(j)$, and thereby a j -specific causal inference model for the observed data, for all $j = 0, \dots, K$.

Verification of j -specific Consistency Assumption. We have that $L(l) = L_{\bar{A}(j-1), A(j), \dots, A(l-1)}(l) = L_{\bar{A}(j)}^j(l)$, $l = j+1, \dots, K+1$, which proves the j -specific consistency assumption:

$$O = (W(j), A(j), L_{\bar{A}(j)}^j(j+1), \dots, A(K), Y = L_{\bar{A}(j)}^j(K+1)). \quad (13)$$

In words this assumption implies that the observed data are simply the observed baseline covariates (including treatment history) up till time j , the observed treatment at each time-point from j to K , and the counterfactual covariate process at each time point $j+1$ to $K+1$ corresponding to the observed treatment history up till that time-point.

Verification of j -specific Sequential Randomization Assumption. The j -specific SRA states that $P(A(l) \mid \bar{A}(j, l-1), X^j) = P(A(l) \mid \bar{A}(l-1), \bar{L}(l))$, $l = j, \dots, K$. In words, the assumption states that treatment at each time-point after time-point j is independent of the full data (indexed by time-point j), given the observed treatment and covariate history up till that time-point. We have

$$\begin{aligned} P(A(l) \mid \bar{A}(j, l-1), X^j) &= P(A(l) \mid \bar{A}(l-1), W(j), X^j) \\ &= P(A(l) \mid \bar{A}(l-1), \bar{L}(l)), \end{aligned} \quad (14)$$

where the latter equality (14) is a consequence of our SRA assumption (2), as we will show now.

Recall that $X = (X_{\bar{a}} : \bar{a})$ is the collection of counterfactual time-dependent processes corresponding with setting complete treatment regimens $\bar{a} = (a(0), \dots, a(K))$, and $X^j = (X_{\bar{A}(j-1), \underline{a}(j)} : \underline{a}(j))$ is the corresponding collection of time-dependent counterfactuals leaving the first j -treatment values random at $\bar{A}(j-1)$. Suppose we condition on $\bar{A}(l-1), \bar{L}(l)$ for a $l \geq j$. By our SRA assumption w.r.t. X , we know that conditional on $\bar{A}(l-1), \bar{L}(l)$, $A(l)$ is independent of X . However, conditional $\bar{A}(l-1), \bar{L}(l)$, the random components $\bar{A}(j-1)$ in the counterfactuals of X^j are fixed, so that X^j is a subvector of the random vector X . By the fact that, given an arbitrary random vector (Z_1, Z_2, Z_3) , $Z_1 \perp (Z_2, Z_3)$ ($\perp$ stands for independence) implies $Z_1 \perp Z_2$ and $Z_1 \perp Z_3$, we have, conditional on $\bar{A}(l-1), \bar{L}(l)$ for a $l \geq j$, that $A(l) \perp X$ implies $A(l) \perp X^j$. This proves the wished j -specific SRA (14).

Verification of j -specific Experimental Treatment Assignment Assumption. The j -specific experimental treatment assignment assumption states that

$$g_0(a(l) \mid \bar{A}(j, l-1), W(j), \bar{L}(j+1, l)) > 0 \quad (15)$$

for all possible $a(l) \in \{a_* : (\bar{A}(l-1), a_*(l)) \in \mathcal{A}(l)\}$. However, $g_0(a(l) \mid \bar{A}(j, l-1), W(j), \bar{L}(j+1, l)) = g_0(a(l) \mid \bar{A}(l-1), \bar{L}(l))$, which is positive by the original experimental treatment assignment assumption (3).

The j -specific MSM.

Let $V(j) \subset W(j)$ be given. We now note that

$$E(Y_{\bar{A}(j-1), \underline{a}(j)} | V(j)) = E(Y_{\underline{a}(j)}^j | V(j)). \quad (16)$$

Thus, our history-adjusted marginal structural model implies at time j the following j -specific marginal structural model:

$$E(Y_{\underline{a}(j)}^j | V(j)) = m_{\beta_0}(j, \underline{a}(j) | V(j)). \quad (17)$$

The j -specific CA (13), j -specific SRA (14), j -specific ETA (15), and MSM (17) define now a j -specific causal inference model $\mathcal{M}_j$ for the observed data density.

Showing equality.

Above, we verified that our HA-MSM model $\mathcal{M}$ for the observed data defined by CA, SRA, ETA and the HA-MSM model for θ_0 is contained in the j -specific MSM model $\mathcal{M}_j$ for the observed data, defined by the j -specific CA (13), j -specific SRA (14), j -specific ETA (15), and MSM (17): in notation, we showed $\mathcal{M} \subset \cap_j \mathcal{M}_j$. In order to prove that $\mathcal{M} = \cap_j \mathcal{M}_j$, we need to show that if $P \in \mathcal{M}_j$ for all j , then $P \in \mathcal{M}$. This happens to be trivial, since $CA(0) = CA$, $SRA(0) = SRA$, and $ETA(0) = ETA$, while it is obvious that the intersection of all j -specific restrictions on θ_0 implies the overall HA-MSM model on θ_0 . This proves that $\mathcal{M} = \cap_j \mathcal{M}_j$.

9 Appendix 2: Asymptotics assuming consistent estimation of the treatment mechanism.

We will use the notation $Pf = \int f(o)dP(o)$ for a given function f and probability distribution P for O . Let μ_n^0 be an initial estimator of β_0 , and let $\mu_n^1 = \mu_n^0 - P_n c_n^{-1} IC(\cdot | Q_n, G_n, h_n, \mu_n^0)$ be the corresponding one-step estimator, where $c_n = \frac{d}{d\mu} E_0 IC(O | Q_n, G_n, h_n, \mu) \big|_{\mu=\mu_n^0}$: that is, μ_n^1 is the estimator one obtains at the first step in the Newton-Raphson algorithm for solving the estimating equation $P_n IC(\cdot | Q_n, G_n, h_n, \mu) = 0$, using μ_n^0 as initial starting point. The next theorem also applies to the IPTW-estimators (just let $IC = IC_0$).

Theorem 2 Consider the observed data model $\mathcal{M}(\mathcal{G})$ for the observed data structure O defined by Assumptions CA (1), SRA (2), ETA (3), HA-MSM (7), and model $\mathcal{G}$ for the conditional distribution G of $\bar{A}$, given X . Let $O_1, \dots, O_n$

be n i.i.d. observations of $O \sim P_0 = P_{F_{X_0}, G_0} \in \mathcal{M}(\mathcal{G})$. Consider a one-step estimator of the parameter $\mu_0 = \beta_0 \in \mathbb{R}^k$ of the form $\mu_n^1 = \mu_n^0 - P_n IC(\cdot \mid Q_n, G_n, c_n, h_n, \mu_n^0)$.

Let $f_n(\mu) \equiv P_n IC(\cdot \mid Q_n, G_n, h_n, \mu)$. Assume (we write $f \approx g$ for $f = g + o_P(1/\sqrt{n})$)

$$c_n^{-1}\{f_n(\mu_n^0) - f_n(\mu_0)\} \approx \mu_n^0 - \mu_0. \quad (18)$$

In addition, assume

(i) $c_n^{-1} IC(\cdot \mid Q_n, G_n, h_n, \mu_n^0)$ falls in a P_0 -Donsker class with probability tending to 1.

(ii) Let $IC_n(\cdot) = c_n^{-1} IC(\cdot \mid Q_n, G_n, h_n, \mu_n^0)$. For some (h, Q^1) , we have

$$\| IC_n(\cdot) - c^{-1} IC(\cdot \mid Q^1, G_0, h, \mu_0) \|_{P_0} \rightarrow 0,$$

where the convergence is in probability, and c denotes the limit of the matrix c_n .

(iii) Define for a G_1 and each $j = 0, \dots, K$

$$\Phi^j(G_1) = P_0 c^{-1} IC^j(\cdot \mid Q^1, G_1, h, \mu_0).$$

For notational convenience, let

$$\begin{aligned} IC_n^j(G) &\equiv c_n^{-1} IC^j(\cdot \mid Q_n, G, h_n, \mu_0), \\ IC^j(G) &\equiv c^{-1} IC^j(\cdot \mid Q^1, G, h, \mu_0). \end{aligned}$$

Assume

$$P_0\{IC_n^j(G_n) - IC_n^j(G_0)\} \approx \Phi^j(G_n) - \Phi^j(G_0).$$

(iv) $\Phi^j(G_n)$ is an asymptotically efficient estimator of $\Phi^j(G_0)$ for the SRA model $\mathcal{G}^j$ (implied by $\mathcal{G}$) containing the true $g_0(\underline{A}_j \mid X, \bar{A}(j-1))$ with tangent space $T_{G,j}(P_0)$

Then μ_n^1 is asymptotically linear with influence curve $-c^{-1} IC$, where

$$IC \equiv IC(\cdot \mid Q^1, G_0, h, \mu_0) - \sum_{j=0}^K \Pi(IC^j(\cdot \mid Q^1, G_0, h_j, \mu_0) \mid T_{G,j}(P_0)).$$

If $Q^1 = Q_{X_0}$, then this influence curve equals $-c^{-1} IC(\cdot \mid Q_0, G_0, h, \mu_0)$. In particular, if $h = h_{opt}$ so that $-c^{-1} IC(O \mid Q_{X_0}, G_0, h_{opt}, \mu_0)$ equals the efficient influence curve $S_{eff}^*(O \mid P_0)$ of μ at P_0 , then μ_n^1 is asymptotically efficient.

The proof is analogue to Theorem 2.4 in van der Laan and Robins (2002), and will therefore be omitted. In particular, the form of the influence curve follows from the analogue representation of the influence curve of the one-step estimator in the j -specific causal model.

Discussion of asymptotic linearity Theorem 2. We will discuss the assumptions of Theorem 2 and illustrate that the assumptions are natural.

Condition (18) is a natural condition, which is illustrated as follows. By definition, we have

$$c_n = f'_n(\mu_n^0) \equiv d/d\mu f_n(\mu)|_{\mu=\mu_n^0}.$$

In this notational setting, condition (18) translates to

$$\{f'_n(\mu_n^0)\}^{-1}\{f_n(\mu_n^0) - f_n(\mu_0)\} = \mu_0 - \mu_n^0 + o_P(1/\sqrt{n}). \quad (19)$$

Under regularity conditions (e.g., a Taylor expansion of $f_n(\mu)$ at μ_n^0), one expects to have a first-order expansion

$$f_n(\mu_0) - f_n(\mu_n^0) = f'_n(\mu_n^0)(\mu_n^0 - \mu_0) + o(|\mu_n^0 - \mu_0|). \quad (20)$$

If the second-order term is $o_P(n^{-1/2})$ and the determinant $f'_n(\mu_n^0)$ is bounded away from zero uniformly in n , then this expansion proves (19) and thus (18). If $f_n(\mu)$ is continuously differentiable in μ at μ_0 with bounded derivative, then the second order term is $o_P(|\mu_n^0 - \mu_0|)$, while if it is twice continuously differentiable, then the second-order term is $O(|\mu_n^0 - \mu_0|^2)$. We can conclude that condition (18) typically requires a convergence rate of the initial estimator μ_n^0 .

Condition (i) can often be arranged by choosing truly lower-dimensional working models $\mathcal{G}$ and $\mathcal{Q}$ when estimating G_0 and Q_{X0} . This condition formally represents the “asymptotic curse of dimensionality” since if one uses as working models $\mathcal{G}(SRA)$ (that is, one only assumes SRA) or an unspecified working model for Q_X , then the class of functions of $IC(\cdot | Q_n, G_n, h_n, \mu_n^0)$ of O generated by varying Q_n, G_n over all possible parameter values will typically be very large (meaning that for finite samples the first-order asymptotics is irrelevant) or not even be a Donsker class. Condition (ii) is a weak consistency condition requiring that G_n be consistent and Q_n converge to something. This condition will hold if our model $\mathcal{M}(\mathcal{G})$ contains the true P_0 . Regarding condition (iii), we have

$$\begin{aligned} P_{F_{X0}, G_0} IC_n(G_n) - IC_n(G_0) &= P_{F_{X0}, G_0} IC_n(G_n) \\ &= P_{F_{X0}, G_0 - G_n} IC_n(G_n), \end{aligned}$$

where the latter equality holds because of the protection of unbiasedness of the estimating function against misspecification of Q_{X0} at the data-generating distribution being P_{F_{X0}, G_n} . Similarly, $P_{F_{X0}, G_0} IC(G_n) - IC(G_0) = P_{F_{X0}, G_0 - G_n} IC(G_n)$. Thus, condition (iii) requires that second-order terms involving integrals of differences $(G_n - G_0)(Q_n - Q^1)$ are $o_P(1/\sqrt{n})$.

Thus, if G_n converges to G_0 at a rate $n^{-1/2}$, then this condition typically only requires consistency of the other nuisance parameter estimates Q_n , but if

G_n converges at a very low rate to G_0 , then the rate of Q_n needs to compensate for this by converging at an appropriate rate.

Condition (iv) just requires that one estimates G_0 with an efficient procedure such as maximum likelihood estimation. Condition (iv) is not needed to establish that μ_n^1 is asymptotically linear, but it is needed to obtain the elegant formula of its influence curve in terms of projection operators.

10 Appendix 3: R code for calculation of counterfactual probabilities.

Example showing non-optimality of the statically-optimal dynamic treatment regimen.

```
proby<-function(a0,l1,a1,beta1,beta2,beta3)
{ if (l1==1) out<-(a1==0)*(beta1*a0)
  if (l1==0) out<-(a0==1)*beta2+(a1==1)*beta3
  out
}

probl1<-function(a0,alpha1,alpha2) {
out<-(a0==1)*alpha1+(a0==0)*alpha2 out }

probycf<-function(a0,a1,alpha1,alpha2,beta1,beta2,beta3) {
out<-proby(a0,1,a1,beta1,beta2,beta3)*probl1(a0,alpha1,alpha2)+
proby(a0,0,a1,beta1,beta2,beta3)*(1-probl1(a0,alpha1,alpha2))
out }

rulea1<-function(l1) { out<-(l1==0) out }

probycfrule<-function(a0,alpha1,alpha2,beta1,beta2,beta3)
{ a11<-rulea1(1) a10<-rulea1(0)
out<-proby(a0,1,a11,beta1,beta2,beta3)*probl1(a0,alpha1,alpha2)
+proby(a0,0,a10,beta1,beta2,beta3)*(1-probl1(a0,alpha1,alpha2))
out
}

alpha1<-0.7
alpha2<-0.2
beta1<-0.2
beta2<-0.7
beta3<-0.5

probycf(1,1,alpha1,alpha2,beta1,beta2,beta3)
probycf(1,0,alpha1,alpha2,beta1,beta2,beta3)
probycf(0,1,alpha1,alpha2,beta1,beta2,beta3)
probycf(0,0,alpha1,alpha2,beta1,beta2,beta3)
probycfrule(1,alpha1,alpha2,beta1,beta2,beta3)
```

`probycfrule(0,alpha1,alpha2,beta1,beta2,beta3)`

References

- J. Bryan, Z. Yu, and M.J. van der Laan. Analysis of longitudinal marginal structural models. *Biostatistics*, 5(3):361–380, 2003.
- R. Gill and J.M. Robins. Causal inference for complex longitudinal data: continuous variables. *Ann. Stat.*, 29(6):1785–1811, 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):331–584, 2003.
- 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, 2003. URL www.bepress.com/ucbbiostat/paper134/.
- R. Neugebauer and M.J. van der Laan. Why prefer double robust estimates. *Journal of Statistical Planning and Inference*, 129(1-2):405–426, 2005.
- J. Neyman. On the application of probability theory to agricultural experiments. *Statistical Science*, 5:465–480, 1990.
- J. M. Robins and A. Rotnitzky. Comment on the Bickel and Kwon article, "Inference for semiparametric models: Some questions and an answer". *Statistica Sinica*, 11(4):920–936, 2001.
- 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. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association: Section on Bayesian Statistical Science*, pages 6–10, 1999.
- 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, 2000.
- 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, 2003a.
- J.M. Robins. Semantics of causal dag models and the identification of direct and indirect effects. In N. Hjort P. Green and S. Richardson, editors, *Highly Structured Stochastic Systems*, pages 70–81. Oxford University Press, Oxford, 2003b.
- D.B. Rubin. Bayesian inference for causal effects: the role of randomization. *Annals of Statistics*, 6:34–58, 1978.
- 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 130, Division of Biostatistics, University of California, Berkeley, 2003. URL www.bepress.com/ucbbiostat/paper130/.

- M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2002.
- Y. Wang and M.J. van der Laan. Data adaptive estimation of the treatment specific mean. Technical Report 159, Division of Biostatistics, University of California, Berkeley, 2004. URL [paper159/](#).
- Z. Yu and M.J. van der Laan. Construction of counterfactuals and the g-computation formula. Technical Report 122, Division of Biostatistics, University of California, Berkeley, 2002. URL [paper122/](#).
- Z. Yu and M.J. van der Laan. Double robust estimation in longitudinal marginal structural models. Technical Report 132, Division of Biostatistics, University of California, Berkeley, 2003. URL [paper132/](#).