The International Journal of Biostatistics

Volume 6, Issue 2

2010

Article 8

CAUSAL INFERENCE

Dynamic Regime Marginal Structural Mean Models for Estimation of Optimal Dynamic Treatment Regimes, Part I: Main Content

Liliana Orellana, *Instituto de Cálculo, Universidad de
Buenos Aires*

Andrea Rotnitzky, *Universidad Torcuato Di Tella and
Harvard School of Public Health*

James M. Robins, *Harvard School of Public Health*

Recommended Citation:

Orellana, Liliana; Rotnitzky, Andrea; and Robins, James M. (2010) "Dynamic Regime Marginal Structural Mean Models for Estimation of Optimal Dynamic Treatment Regimes, Part I: Main Content," *The International Journal of Biostatistics*: Vol. 6: Iss. 2, Article 8.

DOI: 10.2202/1557-4679.1200

Dynamic Regime Marginal Structural Mean Models for Estimation of Optimal Dynamic Treatment Regimes, Part I: Main Content

Liliana Orellana, Andrea Rotnitzky, and James M. Robins

Abstract

Dynamic treatment regimes are set rules for sequential decision making based on patient covariate history. Observational studies are well suited for the investigation of the effects of dynamic treatment regimes because of the variability in treatment decisions found in them. This variability exists because different physicians make different decisions in the face of similar patient histories. In this article we describe an approach to estimate the optimal dynamic treatment regime among a set of enforceable regimes. This set is comprised by regimes defined by simple rules based on a subset of past information. The regimes in the set are indexed by a Euclidean vector. The optimal regime is the one that maximizes the expected counterfactual utility over all regimes in the set. We discuss assumptions under which it is possible to identify the optimal regime from observational longitudinal data. Murphy et al. (2001) developed efficient augmented inverse probability weighted estimators of the expected utility of one fixed regime. Our methods are based on an extension of the marginal structural mean model of Robins (1998, 1999) which incorporate the estimation ideas of Murphy et al. (2001). Our models, which we call dynamic regime marginal structural mean models, are specially suitable for estimating the optimal treatment regime in a moderately small class of enforceable regimes of interest. We consider both parametric and semiparametric dynamic regime marginal structural models. We discuss locally efficient, double-robust estimation of the model parameters and of the index of the optimal treatment regime in the set. In a companion paper in this issue of the journal we provide proofs of the main results.

KEYWORDS: dynamic treatment regime, double-robust, inverse probability weighted, marginal structural model, optimal treatment regime, causality

Author Notes: This work was supported by NIH grant R01 GM48704.

1 Introduction

Dynamic treatment regimes are set rules for sequential decision making based on patient covariate history. In the last decade, major developments in methods for estimating the effects of dynamic treatment regimes from observational data emerged, specifically with a series of papers by Robins (1986, 1989, 1993, 1997). In these papers Robins considers the use of structural nested models to estimate a variety of causal contrasts. In his 1993 paper, Robins raised the possibility of estimating dynamic treatment regime effects by censoring subjects the first time they fail to adhere to the dynamic regime. Murphy, van der Laan and Robins (2001) implemented this idea and applied it to the estimation of the mean of a potential outcome under one dynamic treatment regime, possibly conditional on baseline covariates.

An important public health question related to the management of chronic diseases is to determine the optimal dynamic treatment regime among those in a set of simple regimes that can be enforced in practice. Recently, the task of estimating the optimal dynamic treatment regime from observational data has seen major methodological developments, specifically with the seminal paper of Murphy (2003) and the subsequent work on doubly-robust g-estimation of optimal regime structural nested mean models of Robins (2004). Until these papers, the Biostatistical community had regarded the problem of finding optimal treatment regimes from longitudinal observational databases as a very hard, nearly intractable, problem for the following reasons: *i*) methods for estimation of treatment effects have to appropriately control for high dimensional time dependent confounders (i.e. time varying risk factors that affect future treatments) that are themselves predicted by past treatments; standard longitudinal regression methods which adjust for time dependent risk factors generally yield biased estimators of the time dependent treatment effects (Robins, 1997) and, *ii*) the determination of an optimal treatment strategy is a high dimensional sequential decision problem. Specifically, because the treatment to be prescribed at each time k is decided based on updated information, the set of all potential dynamic regimes, from which the optimal regime needs to be identified, is very large.

Murphy and Robins' methods have one important limitation; they optimize over a class of regimes in which the decision maker has access to very rich information on the patient's covariate history, this being comprised by an increasing sigma field of either a subset or all of the time dependent confounders for treatment. With such rich information, the optimal rule may well be a very complicated function of past covariate history, one that may be hardly enforceable in most health care settings. The methods we will de-

velop in this paper are suitable for finding the optimal regime out of a smaller class of regimes in which treatment decisions are based on limited information about the patient's covariate history. Our methods are based on an extension of the marginal structural mean (MSM) model of Robins (1998, 1999) which incorporate estimation ideas of Murphy et al. (2001). Our model, which we call dynamic regime marginal structural mean model, is specially suitable for estimating the optimal treatment regime in a moderately small class of enforceable regimes of interest.

The idea of using dynamic regime MSM models as a device to estimate optimal treatment regimes was first introduced in Andrea Rotnitzky's NIH grant proposal submitted in November 2005, where it was also indicated how double-robust locally efficient inverse probability weighted estimators of the model parameters could be constructed. This proposal had been developed in collaboration with Liliana Orellana and James Robins. The methods were further investigated and extended in Liliana Orellana's Ph.D. thesis at the Department of Biostatistics of Harvard School of Public Health, during the period 2005-2007. During the same period similar ideas were independently developed by Mark van der Laan and were reported in van der Laan (2006), van der Laan and Petersen (2007) and later illustrated in Bembom and van der Laan (2008) for estimation of the optimal treatment regimes from sequentially randomized trial data. In this paper we report in detail the results developed in Orellana's Ph.D. thesis (2007). Some of these results without proofs were included in the overview paper Robins, Orellana and Rotnitzky (2008). These results differ from and extend the results in van der Laan (2006) and van der Laan and Petersen (2007) and Bembom and van der Laan (2008) in a number of ways. First, these papers considered only parametric models for the dependence of the mean of the counterfactual outcome on the dynamic treatment regime. Here we consider also the more flexible semiparametric MSM models. Second, van der Laan and Petersen derived, as we do here, a class of double-robust estimators of the dynamic regime MSM parameters but did not discuss the efficient choice in the class. Here we derive this efficient choice and propose an estimator that has smallest asymptotic variance among all double-robust (DR) estimators regardless of which of the two working models postulated to construct the DR estimator is correct. Third, of the three aforementioned articles, only Bembom and van der Laan (2008) discuss the construction of confidence regions about the index of the optimal regime in the class, but they do it assuming the class is uncountable and indexed in a continuum. Here we consider the important special case in which the class is finite, a case which raises non-trivial technical challenges.

The methods for estimation and inference of the parameters of dynamic regime MSM models can be derived from the general theory for inference in models for coarsened at random data developed by Robins and Rotnitzky (1992). We are aware that the general theory in that paper (and later exposed and specialized to many examples in a number of papers (e.g. Robins, Rotnitzky and Zhao, 1994, Robins, 1998 and 1999) and books (van der Laan and Robins, 2003, Tsiatis, 2006)) is difficult to follow for the average reader of methodological statistical journals. With this in mind, in this paper we present the proposed methods in an expository manner, providing discussion and examples of the key assumptions they rely on, and of the modelling steps and algorithms required to compute the proposed estimators. Furthermore, we provide a step-by-step derivation of certain optimality and double-robustness properties that are a consequence of the Robins-Rotnitzky general theory for inference in coarsened at random models. In addition, with the intention of making the presentation self-contained in a companion paper in this issue of the journal (Orellana, Rotnitzky and Robins, 2010) we provide proofs of key results that appeared elsewhere in earlier papers of Robins and colleagues.

This paper is organized as follows. In Section 2 we introduce notation, the data structure, the concept of dynamic treatment regimes and the definition of the potential variables. In Section 3 we discuss a set of assumptions that allows identification of the expected utility in the hypothetical world in which everybody follows a given dynamic regime and we derive the expression of this expected utility as a functional of the observed data law. In Section 4 we introduce dynamic regime parametric and semiparametric marginal structural mean models (MSM). In Section 5 we propose augmented (AIPTW) and non-augmented inverse probability of treatment weighted (IPTW) estimating equations for the parameters of the dynamic regime MSM models. We derive their asymptotic distribution and we use this asymptotic distribution to guide the derivation of a class of locally efficient AIPTW estimators indexed by a function b . In Section 6 we motivate and discuss the double-robustness property of each locally efficient AIPTW estimator in the class. We discuss double-robust inference (variance estimation and confidence regions) for the parameters of the dynamic MSM models and for the optimal treatment regime. In Section 7 we derive the optimal index b and propose an estimator that is double-robust and locally efficient in the class of all AIPTW estimators of the parameters of the dynamic regime MSM models. In Section 8 we provide a discussion of some limitations of our proposal.

2 General Formulation

2.1 The Setup

Suppose that a registry contains observational data about a group of patients, all of whom were followed from an agreed upon baseline event until at least the end of K time intervals or until their death time, whichever came first. Patients came to the clinic once during each of the K equidistant intervals to have various clinical and laboratory measurements made. Assume that clinic visits happened at the end of the interval. Furthermore, assume that no patient missed a clinic visit. Treatment decisions, i.e. whether to start, switch, discontinue or alter the dose of a treatment, were made by the physicians after examining the patient's laboratory and clinical data results. Thus, treatment decisions were made soon after each clinic visit and at no other moment. Assume that the subjects in the registry are a random sample from a large population of interest.

Each record of the database contains the patient's information recorded over the entire follow-up period. This is comprised by the variables

$$L_0, A_0, R_1, T_1, L_1, A_1, \dots, R_k, T_k, L_k, A_k, R_{K+1}, T_{K+1}, L_{K+1}$$

where R_k is a binary indicator for being at Risk, that takes the value 0 if the patient has experienced an event of interest by time k and 1 otherwise (this event often being death, but possibly also being the onset of a disease, such as the onset of a symptom defining AIDS), the variable T_k denotes the minimum between the time to the event of interest and k (here and throughout the unit measure for time is 1 unit = length of one inter-visit interval, and time to event is measured since time of start of follow-up), L_k are the clinical and laboratory variables measured during the k^{th} clinic visit if the patient was alive, A_k is the subsequent treatment prescription which we assume takes values in a finite set $\mathcal{A}_k$. In the database, the entries for L_k and A_k at a time k after the patient experienced the event of interest, e.g. death, are set to any agreed upon value such as NA , $*$ or the last measured values of these variables. The chosen convention is inconsequential for any analysis method that, like the one presented in this article, disregards the entries of these variables after the occurrence of event. To avoid notational burden, we will assume that the entries for the post-event variables L_k and A_k are set equal to the last available values for these variables. One final outcome measurement L_{K+1} is available for all subjects in the registry that have not yet experienced the event at time $K + 1$.

Throughout we use the following conventions:

$$A_{-1} \equiv 0, O_k \equiv (R_k, T_k, L_k),$$

overbars with a subscript, say k , denote the present variable, i.e. at time k , and all its past values, e.g.

$$\overline{O}_k \equiv (O_0, \dots, O_k)$$

and variables without a time index denote the entire variable history, i.e.

$$O \equiv \overline{O}_{K+1} \text{ and } A = \overline{A}_K.$$

Furthermore, we use capital letters, such as $O_k = (R_k, T_k, L_k)$, to refer to random variables or vectors, i.e., variables which can take on different values for different subjects. We use small letters, such as $o_k = (r_k, t_k, l_k)$, to refer to the possible values of the corresponding capital letter random variable.

2.2 The Treatment Regime

A dynamic treatment regime is a sequential rule for determining, at each time k , the next treatment prescription A_k . The rule may depend on part or all of the recorded health information about the patient's health up to and including time k . Formally, a dynamic treatment regime is defined by a collection of maps

$$\overline{o}_k \mapsto g_k(\overline{o}_k) \in \mathcal{A}_k, k = 0, \dots, K.$$

Because we are interested in evaluating the effect of treatment regimes on the health experience of the patient up to the development of the event of interest, in principle, we can define the treatment regime arbitrarily for times k after the occurrence of the event. For example, a treatment rule is meaningless if the patient is dead and in principle, we could define $g_k(\overline{o}_k)$ arbitrarily at any month k after death. However, because our methods require that we use the same convention for coding data after the event time in the hypothetical world in which everyone followed regime g as in the actual world, and in view of the aforementioned convention for coding variables after the event time, we set

$$g_k(\overline{o}_k) = g_{k-1}(\overline{o}_{k-1}) \text{ if } r_k = 0. \quad (1)$$

Remark 1: Some authors define dynamic treatment regimes as a collection of functions mapping each $(\overline{o}_k, \overline{a}_{k-1})$ to an element of $\mathcal{A}_k$. However, for the purposes of evaluating the effect of a dynamic regime that it is enforced

since $k = 0$, this definition is redundant because when a given dynamic regime is enforced, the treatments received up to time $k - 1$ are functions of the outcomes obtained up to time $k - 1$, so the domain of the decision rule at time k depends only on past outcomes.

Example 1: Consider a registry that contains information about a cohort of HIV-infected patients followed since their first doctor's visit after infection. Suppose that patients in the registry return monthly to the clinic for clinical and laboratory tests. Suppose all patients have been followed until month $K + 1 = 37$ from baseline or until their death time, whichever came first and that death occurrence was recorded in continuous time. Patient information includes CD4 cell count and viral load recorded at the clinic visit as well as the treatment prescription. Assume that treatment decisions, i.e. whether or not to give highly active antiretroviral-therapy (HAART), are made at and only at, a clinic visit. Formally, the variables entered into the registry corresponding to month k , $k = 0, \dots, K + 1 = 37$, are $O_k = (R_k, T_k, L_k)$ where T_k is equal to the minimum between i) death time, ii) the time to the occurrence of an AIDS defining event and iii) k , $R_k = 1$ if the subject is alive and AIDS free at the beginning of month k and $R_k = 0$ otherwise. The variables A_k and $L_k = (CD4_k, V_k)$ are defined as follows. At a month k at which the subject is at risk $CD4_k$ is the subject's CD4 cell count, V_k is the logarithm of his/her viral load and A_k denotes the treatment prescription right after measuring $CD4_k$ and V_k (0 if no HAART is given, 1 otherwise). At any month k at which $R_k = 0$, L_k and A_k are set to the values of the health outcomes and treatment indicator recorded at the last visit $k' < k$ at which $R_{k'}$ was 1. A treatment regime that specifies that the patient must start HAART as soon as it is detected that his/her CD4 count is or was at or below 350 cell count/ μL , and must continue on HAART afterwards is a dynamic treatment regime defined by the set of functions $\{g_k : k = 0, \dots, K\}$ where

$$g_k(\bar{O}_k) = \begin{cases} 0 & \text{if } R_k = 1 \text{ and } \min(\overline{CD4_k}) > 350 \text{ cell count}/\mu L \\ 1 & \text{if } R_k = 1 \text{ and } \min(\overline{CD4_k}) \leq 350 \text{ cell count}/\mu L \\ g_{k-1}(\bar{O}_{k-1}) & \text{if } R_k = 0 \end{cases} \quad (2)$$

2.3 The Potential Variables

To define the causal effects of distinct dynamic treatment regimes we use the notion of potential, also called counterfactual, variables introduced by Neyman (1923) and Rubin (1978) for time independent treatments and later extended by Robins (1986, 1987) to the context of time dependent treatments.

Associated with each vector $a \equiv (a_0, \dots, a_K)$, with $a_k \in \mathcal{A}_k$, we conceptualize a vector of the potential outcomes and event indicators

$$O_{(a)} = (L_0, R_{(a_0)1}, T_{(a_0)1}, L_{(a_0)1}, \dots, R_{(\bar{a}_K)K+1}, T_{(\bar{a}_K)K+1}, L_{(\bar{a}_K)K+1})$$

of a subject if he/she had followed, possibly contrary to fact, the treatment pattern $A = a$. The set

$$\mathcal{O} = \{O_{(a)} : a_k \in \mathcal{A}_k, k = 0, \dots, K\}$$

denotes all possible vectors of potential outcomes and event indicators for a random patient from the population. This set includes vectors $O_{(a)}$ that encode potential outcomes under the scenario in which treatment alterations occur after the event of interest, obviously a meaningless situation if the event is death. These meaningless variables will never be used in the methods that we will describe here and we include them in the set $\mathcal{O}$ only to avoid the extra notational burden needed to exclude them from it.

The notation defining $\mathcal{O}$ makes the implicit Stable Unit Treatment Value Assumption (SUTVA) (Cox, 1958, Rubin, 1978) that one subject's potential variables do not depend on the treatment patterns followed by other subjects. SUTVA implies the following assumption connecting the vector O with one specific member of $\mathcal{O}$.

Consistency Assumption (C). For each $k = 1, \dots, K + 1$, the vector $(R_{(\bar{A}_{k-1})k}, T_{(\bar{A}_{k-1})k}, L_{(\bar{A}_{k-1})k})$ is equal to (R_k, T_k, L_k) .

Assumption C stipulates that the outcomes and event indicators recorded in the registry are the potential outcomes and event indicators corresponding to the treatment pattern actually followed by the patient. Assumption C implies that $O = t(\mathcal{O}, A)$ for the map $t : (\mathcal{O}, A) \mapsto O_{(A)}$ that assigns to each $(\mathcal{O}, A)$ the counterfactual string $O_{(A)}$ of $\mathcal{O}$.

Define $A^g = (A_0^g, \dots, A_K^g)$ to be the treatment sequence if, possibly contrary to fact, the subject had obeyed the dynamic regime g . Then $O^g \equiv O_{(A^g)}$ is the vector of outcomes if the subject had followed regime g .

The vectors $(\mathcal{O}, O, A)$ and $(\mathcal{O}, O^g, A^g)$ are random vectors on an underlying probability space $(\Omega, \mathcal{F}, \mathbb{P})$ whose probability laws we denote with P and P_g . The marginal distribution of $\mathcal{O}$, throughout denoted as $P_{\mathcal{O}}$, is the same under either P or P_g because the potential variables in the set $\mathcal{O}$ are, like age and gender, fixed patient characteristics and hence unaffected by the treatments actually followed by the subject. In contrast, the distribution of (O, A) and (O^g, A^g) are different. Specifically, if $I_{\mathcal{U}}(\cdot)$ denotes the indicator function that takes the value 1 if $\cdot$ is in $\mathcal{U}$ and 0 otherwise then

$$\mathbb{P}(A_k^g = a_k | \bar{O}_k^g = \bar{o}_k, \bar{A}_{k-1}^g = \bar{a}_{k-1}) = I_{\{g_k(\bar{o}_k)\}}(a_k)$$

because in the world in which everybody follows regime g , the treatment assignment at time k is equal to $g_k(\bar{O}_k)$ with probability 1. In contrast,

$$\mathbb{P}(A_k = a_k | \bar{O}_k = \bar{o}_k, \bar{A}_{k-1} = \bar{a}_{k-1}) \equiv \lambda_k(a_k | \bar{o}_k, \bar{a}_{k-1}) \quad (3)$$

is the probability that in the actual world treatment a_k is prescribed at time k for a patient with observed past $\bar{O}_k = \bar{o}_k$ and $\bar{A}_{k-1} = \bar{a}_{k-1}$. In observational studies, $\lambda_k(a_k | \bar{o}_k, \bar{a}_{k-1})$ is an unknown function of $a_k, \bar{o}_k$ and $\bar{a}_{k-1}$. Of course, the vectors O and O^g , being deterministic functions of $(\mathcal{O}, A)$ and $(\mathcal{O}, A^g)$ respectively, have different distributions. Throughout, $P^{\text{marg}}, P_g^{\text{marg}}$ denote the marginal distributions of (O, A) and (O^g, A^g) .

3 Identification of the Expected Utility if Everybody Follows Regime g

Suppose that $u(o, a)$ is a user specified utility function of some or all the components of (o, a) that quantifies the health benefits at time $K+1$ of a person who had outcome history $O = o$ and treatment history $A = a$. For instance, in Example 1 of Section 2.2, $u(O, A)$ may be survival time T_{K+1} , or quality-of-life adjusted survival time $\sum_{k=0}^{K+1} u_k(\bar{CD4}_k, \bar{V}_k) R_k$ where $u_k(\bar{CD4}_k, \bar{V}_k)$ is a user-specified utility function that quantifies the health status at time k of a patient with recorded CD4 history $\bar{CD4}_k$ and recorded log-viral load history $\bar{V}_k$.

If Z denotes a subset or possibly all of the baseline vector L_0 defining subpopulations of interest, then the conditional mean

$$m(z, g) \equiv E\{u(O^g, A^g) | Z = z\}$$

is the expected utility for the subpopulation of subjects who have baseline covariates Z equal to z , in the hypothetical world in which they all follow the treatment regime g . Comparison of the expected utility $m(z, g)$ for different regimes g of interest quantifies the causal effects for the subpopulation with $Z = z$. Thus, $m(z, g)$ is a target parameter for inference when evaluation of the causal effect of regime g is of substantive interest.

The consistency assumption C implies that in the database we will have the records of (O^g, A^g) for all the study participants that actually obeyed regime g (since for them, their recorded data (O, A) agrees with their potential data (O^g, A^g)). However, data on (O^g, A^g) will be unavailable for those that did not obeyed regime g . This fundamental missing data problem implies that $m(z, g)$ will generally not be identified from P^{marg} unless certain assumptions are made.

3.1 Identifying Assumptions

In this subsection we state and discuss two assumptions that are standard in the causal inference literature which, when made in conjunction with assumption C, imply the identifiability of $m(z, g)$ from the law P^{marg} .

Sequential Randomization (SR): A_k is conditionally independent of $\mathcal{O}$ given $\bar{O}_k$ and $\bar{A}_{k-1}$, for each $k = 0, \dots, K$.

Positivity Assumption (PO): $\mathbb{P}[\lambda_k(A_k^g | \bar{O}_k^g, \bar{A}_{k-1}^g) > 0] = 1$ for all $k = 0, \dots, K$.

Assumption SR would be true in a sequentially randomized trial in which treatment at each time k was randomly assigned with probabilities depending on the subject's recorded history. In an observational study, the assumption cannot be empirically tested (Robins, 1997). Thus, investigators will need to use their subject matter knowledge to collect data on many relevant variables so that the assumption is at least approximately correct. Murphy et al. (2001) studied conditions for identification of $m(z, g)$ when L_k can be decomposed into (S_k, V_k) , $u(O, A) = S_{K+1}$ and the regime g is such that g_k at each time k depends on $\bar{O}_k$ only through the subvector $(S_0, \dots, S_k)$. They showed that in this special case, a slightly weaker version of the SR assumption which stipulates independence of A_k solely with the $S(\cdot)$ parts of $\mathcal{O}$, together with the PO assumption, suffices for identification of $m(z, g)$.

Assumption PO is tantamount to assuming that if in the scenario in which regime g were to be enforced in the entire population, subjects with outcome history $\bar{o}_k$ and treatment history $\bar{a}_{k-1}$ were to exist and were to be assigned to a_k at time k then, in the observational world (i.e. the setting under which the study data are collected) there must also exist subjects with outcome history o_k , treatment history $\bar{a}_{k-1}$ who take treatment a_k . The intuition behind assumption PO is simple. If we want to learn about the distribution of the outcome in the hypothetical world in which regime g was implemented from data in the actual observational study, then every subject in our study population must have a positive chance of following regime g . To see this, consider the subpopulation, say $\mathcal{P}$, comprised of subjects that if regime g were enforced up to time $k - 1$, they would follow it and would have $\bar{O}_k^g = \bar{o}_k$ and $\bar{A}_{k-1}^g = \bar{a}_{k-1}$. Suppose that a subset of subjects in $\mathcal{P}$ actually followed regime g up to time $k - 1$ in the observational study. The sequential randomization assumption implies that those from this subset that would continue to follow regime g at time k , (i.e. those with $A_k = g_k(\bar{o}_k)$) are representative of the others in this subset that depart from regime g at this cycle. However, suppose that none in this subset has $A_k = g_k(\bar{o}_k)$. Then, in our dataset all subjects from the sub-

population $\mathcal{P}$ have departed from regime g at time k or earlier and as such we have no representative from this subpopulation that follows regime g for the entire follow-up period. Then, unless we make further untestable assumptions relating the distribution of $\bar{O}^g$ in subpopulation $\mathcal{P}$ with the distribution of $\bar{O}^g$ in other subpopulations, we have no way to learn from the observational data the outcome distribution that subjects from $\mathcal{P}$ would have if they followed regime g for the entire follow-up period. In Section 4 of the companion paper (Orellana, Rotnizky and Robins, 2010) we provide an example illustrating this point.

Another point of note about assumption PO is that it is an assumption only about the recorded treatments that represent true treatment assignments. That is, if $R_k = 0$, i.e. if the absorbing event has occurred by time k , then the assumption does not actually place restrictions on the probabilities of the artificially defined treatment variables A_k . To see this note that if subjects with past $\bar{o}_k$ exist in both the hypothetical and the observational world, then the next recorded treatment will automatically take the same value in both worlds when the event indicator r_k is 0 because, by convention, this next treatment will agree with the previous recorded treatment value a_{k-1} .

3.2 Identification Lemma

Under the convention $0/0 = 0$, assumption PO implies that the functions

$$\omega_k(\bar{o}_k, \bar{a}_k) \equiv \prod_{j=0}^k \frac{I_{\{g_j(\bar{o}_j)\}}(a_j)}{\lambda_j(a_j|\bar{o}_j, \bar{a}_{j-1})} \quad (4)$$

and

$$\underline{\omega}_{k,k+l}(\bar{o}_{k+l}, \bar{a}_{k+l}) \equiv \prod_{j=k+1}^{k+l} \frac{I_{\{g_j(\bar{o}_j)\}}(a_j)}{\lambda_j(a_j|\bar{o}_j, \bar{a}_{j-1})}$$

are well defined for all $k = 0, \dots, K$ and all $l > 0$ such that $k + l \leq K$. In the product (4), k can be replaced by $k^* \equiv k^*(\bar{o}_k, \bar{a}_k)$, the minimum between k and the last visit time prior to the occurrence of the event of interest when $\bar{O}_k = \bar{o}_k$ and $\bar{A}_k = \bar{a}_k$. This is so because when $r_k = 0$, *i*) $I_{\{g_k(\bar{o}_k)\}}(a_k) = I_{\{g_k(\bar{o}_{k-1})\}}(a_{k-1})$ (by (1)) and, *ii*) $\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1}) = I_{\{a_{k-1}\}}(a_k)$ (because, by convention, A_k at a time k after the absorbing event has occurred is set to the last previously prescribed treatment).

The following key Lemma due to Robins (1986, 1987) (see also Robins, 1997, and Murphy, van der Laan and Robins, 2001), implies that under assumptions C, SR and PO, the law P_g^{marg} of the outcomes and treatments that

would be recorded in the hypothetical world in which everyone followed regime g is identified by P^{marg} , the law of the data recorded in the actual study. The proof for Lemma 1 is given in Section 2.1 of the companion paper (Orellana, Rotnitzky and Robins, 2010).

Lemma 1. Under assumptions C, SR and PO,

1. If (O, A) is a vector in $\mathbb{R}^v$, then for any Borel set B of $\mathbb{R}^v$ it holds that for all $k = 0, \dots, K$

$$\begin{aligned} & E \left[I_B(O, A) \omega_{k-1, K}(\bar{O}_K, \bar{A}_K) \mid \bar{O}_k, \bar{A}_{k-1} = \bar{g}_{k-1}(\bar{O}_{k-1}) \right] \\ &= E \left[I_B(O^g, A^g) \mid \bar{O}_k, \bar{A}_{k-1} = \bar{g}_{k-1}(\bar{O}_{k-1}) \right] \text{ w.p.1.} \end{aligned}$$

2. P_g^{marg} is absolutely continuous with respect to P^{marg} and $\omega_K(\bar{O}_K, \bar{a}_K)$ is a version of the Radon-Nikodym derivative $(dP_g^{\text{marg}}/dP^{\text{marg}})(\bar{O}_{K+1}, \bar{a}_K)$.

In part 1 of the Lemma, we used the notation

$$\bar{g}_k(\bar{o}_k) \equiv (g_0(o_0), g_1(\bar{o}_1), \dots, g_k(\bar{o}_k)).$$

Part 2 of the Lemma implies that under assumptions C, SR and PO, if p_g^{marg} and p^{marg} denote (versions of) the densities of P_g^{marg} and P^{marg} (with respect to some dominating measure), then

$$p_g^{\text{marg}}(o, a) = \omega_K(\bar{o}_K, \bar{a}_K) p^{\text{marg}}(o, a)$$

or equivalently

$$p_g^{\text{marg}}(o, a) = \prod_{j=0}^K I_{\{g_j(\bar{o}_j)\}}(a_j) \prod_{j=1}^{K+1} p^{\text{marg}}(o_j \mid \bar{o}_{j-1}, \bar{a}_{j-1}) p^{\text{marg}}(o_0)$$

where throughout we remove the subscript indicating the variables intervening in any given density; the lower case letters where they are evaluated indicates the random variable that the density correspond to, e.g. we write $p^{\text{marg}}(o_j \mid \bar{o}_{j-1}, \bar{a}_{j-1})$ instead of $p_{O_j \mid \bar{O}_{j-1}, \bar{A}_{j-1}}^{\text{marg}}(o_j \mid \bar{o}_{j-1}, \bar{a}_{j-1})$.

In contrast, part 1 of Lemma 1 informally states that among subjects that in the actual world followed regime g until time $k-1$ and had observed outcome history $\bar{o}_k$ until time k , the joint density (with respect to some dominating measure) of the outcome and treatments (O^g, A^g) that they would have if they continued to follow regime g until time K is

$$\prod_{j=k}^K I_{\{g_j(\bar{o}_j)\}}(a_j) \prod_{j=k+1}^{K+1} p^{\text{marg}}(o_j \mid \bar{o}_{j-1}, \bar{a}_{j-1} = \bar{g}_{j-1}(\bar{o}_{j-1})).$$

Note that the counterfactual law P_g^{marg} has the same conditional law of O_j given $\bar{O}_{j-1}, \bar{A}_{j-1}$ as the observational law P^{marg} for all $j = 0, \dots, K+1$, but it differs from P^{marg} in that the conditional probability of A_j given $\bar{O}_j, \bar{A}_{j-1}$ is the mass point probability $I_{\{g_j(\bar{O}_j)\}}(A_j)$ rather than the conditional probability $\lambda_j(A_j|\bar{O}_j, \bar{A}_{j-1})$. This replacement is not surprising since P_g^{marg} is the law of the treatments and outcomes in the hypothetical world in which everybody follows treatment regime g , and in such a world, A_j is a deterministic, i.e. non-random, function of $\bar{O}_j$. On the other hand, sequential randomization and positivity ensures that the conditional law of O_j given $\bar{O}_{j-1}, \bar{A}_{j-1}$ is the same as the conditional law of O_j^g given $\bar{O}_{j-1}^g, \bar{A}_{j-1}^g$.

Because Z is a subset of the baseline covariates O_0 and because $O_0 = O_0^g$, Part 2 of Lemma 1 also implies that under the assumptions C, SR and PO,

$$\begin{aligned} m(z, g) &\equiv E\{u(O^g, A^g) | Z = z\} \\ &= E\{\omega_K(\bar{O}_K, \bar{A}_K) u(O, A) | Z = z\} \end{aligned} \quad (5)$$

$$= \sum_{\substack{a_k \in \mathcal{A}_k \\ k=0, \dots, K}} \int u(o, a) \prod_{j=0}^K I_{\{g_j(\bar{O}_j)\}}(a_j) \prod_{j=1}^{K+1} dP^{\text{marg}}(o_j | \bar{o}_{j-1}, \bar{a}_{j-1}) dP^{\text{marg}}(o_0 | z) \quad (6)$$

In equation (5) the numerator of $\omega_K(\bar{O}_K, \bar{A}_K)$ is different from 0 for subjects that actually followed treatment regime g and the denominator of $\omega_K(\bar{O}_K, \bar{A}_K)$ is the probability that a subject with counterfactual $\mathcal{O}$ follows the treatment regime g . Thus, $\omega_K(\bar{O}_K, \bar{A}_K)$ censors subjects that did not follow regime g . The weights are just the right ones to represent the censored subjects with the appropriate uncensored subjects. In fact they effectively produces a stratified redistribution to the right operation in which non-compliers to regime g are censored the first time they do not comply and their contribution is redistributed among those that have the same covariate and treatment history and that remain compliers. This redistribution produces the right estimand because, by the sequential randomization assumption SR, among those with the same past, compliance status at a given time is the result of a random mechanism that is independent of the future health outcomes that the subjects would experience if they were to comply with regime g . The redistribution process is tantamount to creating clones of each all-time complier. The compliers and their clones form a pseudo-population with the same distribution of O^g as in the study population, in which everybody followed regime g .

The equality (6) proves that $m(z, g)$ is identified by P^{marg} because the expectation on the right hand side of the equality is an integral with respect to the measure P^{marg} . The rightmost member of the equalities is the so-called G-computation-algorithm (Robins, 1986). After recursive applications of Fubini's theorem (see proof of Lemma 1, Section 2.1 companion paper), this formula has the following alternative expression

$$E[u(O^g, A^g) | Z = z] = E[\phi_1(O_0) | Z = z] \quad (7)$$

where $\phi_1(o_0)$ is defined recursively from,

$$\phi_{K+1}(\bar{o}_K) \equiv E[u(O, A) | \bar{O}_K = \bar{o}_K, \bar{A}_K = \bar{g}_K(\bar{o}_K)]$$

and if $K > 0$,

$$\phi_{k+1}(\bar{o}_k) \equiv E[\phi_{k+2}(\bar{o}_{k+1}) | \bar{O}_k = \bar{o}_k, \bar{A}_k = \bar{g}_k(\bar{o}_k)], \quad k = K-1, \dots, 0. \quad (8)$$

Lemma 1 entails the following further elaborations, proved in Section 2.2 of the companion paper (Orellana, Rotnitzky and Robins, 2010), on the meaning of the function $\phi_{k+1}(\bar{o}_k)$ under assumptions C, SR and PO.

- a) In the hypothetical world in which treatment regime g is implemented in the entire population, $\phi_{k+1}(\bar{o}_k)$ is the utility mean among subjects who have history $\bar{O}_k^g = \bar{o}_k$ up to time k , i.e.

$$\phi_{k+1}(\bar{o}_k) = E[u(O^g, A^g) | \bar{O}_k^g = \bar{o}_k].$$

- b) $\phi_{k+1}(\bar{o}_k)$ is the utility mean among subjects that in the observational world had covariate history $\bar{O}_k = \bar{o}_k$ at time k and that actually followed regime g up to time $k-1$ if, possibly contrary to fact, these subjects were to continue to follow regime g from time k until time K , i.e.

$$\phi_{k+1}(\bar{o}_k) = E[u(O^g, A^g) | \bar{O}_k = \bar{o}_k, \bar{A}_{k-1} = \bar{g}_{k-1}(\bar{o}_{k-1})].$$

Factoring the law of (O, A) as $dP^{\text{marg}} = \mathcal{L}_O \mathcal{L}_A$ where

$$\mathcal{L}_O = \left[\prod_{j=1}^{K+1} dP_{O_j | \bar{O}_{j-1}, \bar{A}_{j-1}}^{\text{marg}} \right] dP_{O_0}^{\text{marg}} \text{ and } \mathcal{L}_A = \prod_{j=0}^K \lambda_j(A_j | \bar{O}_j, \bar{A}_{j-1}) \quad (9)$$

formula (6) implies that, under C, SR and PO, the function $m(\cdot; g) \equiv E[u(O^g, A^g) | Z = \cdot]$ depends on P^{marg} only through the laws entering the $\mathcal{L}_O$ -part of its factorization. We will return to this point in Section 6.

4 Dynamic Regime Marginal Structural Mean Models

In order to choose the best (enforceable) course of action for subjects with baseline covariates $Z = z$, one will want to find the regime $g_{x_{opt}(z)}$ maximizing $m(z; g_x) \equiv E[u(O^{g_x}, A^{g_x}) | Z = z]$ among all g_x in a given class of candidate regimes $\mathcal{R} = \{g_x : x \in \mathcal{X}\}$. That is, one will want to find

$$x_{opt}(z) \equiv \arg \max_{x \in \mathcal{X}} m(z; g_x)$$

where throughout, without loss of generality, we assume that higher values of the utility function are preferable.

The treatment regime $g_{x_{opt}(z)}$ is the optimal treatment regime among regimes in the class $\mathcal{R}$ for subjects with baseline values $Z = z$. Murphy (2003) and Robins (2004) considered the problem of estimating optimal dynamic treatment regimes from observational longitudinal data. In these authors' work, the class $\mathcal{R}$ is very large; it is comprised of g 's such that the g_k 's are arbitrary functions of (the increasing sequence of) sigma fields $\mathcal{F}_k = \sigma(O_0^*, O_1^*, \dots, O_k^*)$ where the O_j^* 's are (possibly different) subsets of the time dependent covariates O_j . However, as indicated in the introduction, often the class $\mathcal{R}$ of candidate (enforceable) regimes is significantly smaller, usually comprised of functions g_k that can depend only on limited covariate history information. For example, consider the pressing question in AIDS research of which is the optimal threshold CD4 count value at which to start prescribing HAART to HIV positive subjects. In this setting, interest lies in finding the optimal regime in the set $\mathcal{R}$ in which g_x is defined as in (2) but with 350 replaced by x and $\mathcal{X}$ is the interval $[200, 600]$. This interval includes the set of scientifically relevant CD4 count/ μL threshold values. The class $\mathcal{R}$ of candidate regimes is much smaller than the class $\mathcal{R}^*$ of all dynamic regimes which includes rules that can possibly depend on past CD4 count in any complicated way. Murphy (2003) and Robins' (2004) methods can estimate the optimal treatment regime in the class $\mathcal{R}^*$ but not in the class $\mathcal{R}$.

In what follows we will consider estimation of $x_{opt}(z)$ when the set $\mathcal{R}$ of regimes of interest is possibly smaller than that considered in Murphy's and Robins' work. We will assume that $\mathcal{R}$ can be indexed by the elements of a set $\mathcal{X}$ that is either discrete, i.e. finite or countable, or a subset of $\mathbb{R}^s$ with no isolated points.

Although under assumptions C, SR and PO for all g_x , $x_{opt}(z)$ is identified by P^{marg} for each z , in practice we cannot hope to estimate it well with the argmax of the individually estimated values of $m(z, g_x)$. This is so because:

i) when z is a vector with two or more continuous components estimation of $m(\cdot, g_x)$ using smoothing techniques will be practically unfeasible due to the curse of dimensionality, and *ii*) even if z took values in a small finite set, in most applications there will be few subjects in the registry that follow any given regime g_x and hence each individual $m(z, g_x)$ will be estimated with large variability. To ameliorate this difficulty we propose conducting inference under models for $m(z, g_x)$ that allow one to combine information from many regimes g_x and across subjects with different values of the baseline covariates Z . One possibility is to consider parametric models of the form

$$m(z, g_x) = h_{\text{par}}(z, x; \beta^*) \quad (10)$$

where $h_{\text{par}}(z, x; \cdot)$ is a known smooth function of a $p \times 1$ parameter β and β^* is unknown. For instance, in example 1 of Section 2.2 we might choose

$$\begin{aligned} h_{\text{par}}(z, x; \beta) = & \beta_1(x - x^*) + \beta_2(x - x^*)^2 + \beta_3(x - x^*)z + \quad (11) \\ & \beta_4(x - x^*)^2 z + \beta_5 + \beta_6 z \end{aligned}$$

with x^* the index of an arbitrary regime, e.g. $x^* = 350$. This model postulates that the function $m(z, g_{x^*})$ is equal to $\beta_5 + \beta_6 z$ and that the difference $m(z, g_x) - m(z, g_{x^*})$ is a quadratic function of $(x - x^*)$ with coefficients which depend linearly on z .

For the purposes of estimating $x_{\text{opt}}(z)$, model (10) is unnecessarily stringent. Specifically, for any fixed index x^* the optimal index $x_{\text{opt}}(z)$ satisfies

$$x_{\text{opt}}(z) = \arg \max_{x \in \mathcal{X}} \{m(z, g_x) - m(z, g_{x^*})\}.$$

Thus, for inference about $x_{\text{opt}}(z)$ it suffices to model the difference $m(z, g_x) - m(z, g_{x^*})$. However, model (10) places parametric restrictions in the form of the dependence on z of not only the differences $m(z, g_x) - m(z, g_{x^*})$ but also of the function $m(z, g_{x^*})$. This latter unnecessary parametric restriction may lead to invalid inference about $x_{\text{opt}}(z)$. For instance, model (11) makes the unnecessary assumption that $m(z, g_{x^*})$ is a linear function of z . If in fact, $m(z, g_{x^*})$ is not linear in z , then estimators of $x_{\text{opt}}(z)$ computed under this assumption will generally be inconsistent. This remark suggests that in order to reduce the chance of invalid inference about $x_{\text{opt}}(z)$ due to model misspecification and to retain the possibility of borrowing information across treatments and baseline covariates, we consider (following Robins, 1999) flexible semi-parametric models of the form

$$m(z, g_x) = h_{\text{sem}}(z, x; \beta^*) + q(z) \quad (12)$$

where $q(\cdot)$ is an unknown function of z only, β^* is an unknown $p \times 1$ parameter and for each z and x , $h_{\text{sem}}(z, x; \beta)$ is a known smooth function of β satisfying $h_{\text{sem}}(\cdot, x^*; \cdot) = 0$ for the index x^* of an arbitrary, user-specified, regime. For example,

$$h_{\text{sem}}(z, x; \beta) = \beta_1(x - x^*) + \beta_2(x - x^*)^2 + \beta_3(x - x^*)z + \beta_4(x - x^*)^2 z \quad (13)$$

The condition $h_{\text{sem}}(\cdot, x^*; \cdot) = 0$ implies that $q(z)$ is equal to $m(z, g_{x^*})$, the expected utility function for the regime x^* , and hence that $h_{\text{sem}}(z, x; \beta^*) = m(z, g_x) - m(z, g_{x^*})$. Thus, to estimate $x_{\text{opt}}(z)$ it suffices to estimate the value of x maximizing $h_{\text{sem}}(z, x; \beta^*)$.

So long as $h_{\text{sem}}(z, x; \beta) = 0$ for some value of the vector β , model (12) is guaranteed to be correctly specified under the null hypothesis that all treatments regimes g_x are equally effective. Thus, inference under model (12) are guaranteed to result in valid α level tests of indifference between treatment regimes. In contrast, inference under model (10) do not generally share this robustness property. Which model (10) or (12) should one adopt for estimation of β , and hence of $x_{\text{opt}}(z)$, raises the usual bias/efficiency trade-off considerations: if model (10) is indeed correctly specified then, in general, estimators of $x_{\text{opt}}(z)$ that are nearly efficient under model (10) will be more efficient than estimators of $x_{\text{opt}}(z)$ that are nearly efficient under model (12); however the former will be inconsistent if model (10) is incorrect and model (12) is correct while the latter will be consistent under this circumstance. In our opinion, unless firm scientific background warrants model (10), model (12) should be preferable for inference about $x_{\text{opt}}(z)$ as it provides the data analyst a better chance to carry out valid inference about $x_{\text{opt}}(z)$.

Relaxing the positivity assumption: Under (10) or (12), $m(z, g_x)$, and consequently $x_{\text{opt}}(z)$, may be identified by P^{marg} for all $x \in \mathcal{X}$ even when the positivity assumption PO fails for some g_x . For instance, under model (10) identification of β^* suffices for identification of $m(z, g_x)$ for all $x \in \mathcal{X}$. Yet, for example, when $h_{\text{par}}(z, x; \beta)$ is defined as in (11), β^* is identified provided $m(\cdot; g_x)$ is identified for three different values of x in $\mathcal{X}$. Thus, to identify $m(z, g_x)$ for all $x \in \mathcal{X}$ it suffices that the PO assumption is satisfied for all g_x with $x \in \mathcal{X}_{\text{pos}}$ where $\mathcal{X}_{\text{pos}}$ is any subset of $\mathcal{X}$ comprised of at least three different x values. Throughout, for any subset $\mathcal{C}$ of $\mathcal{X}$, $\text{PO}(\mathcal{C})$ will stand for the assumption that stipulates that PO holds when g is replaced by g_x for all $x \in \mathcal{C}$. Furthermore, we will let $\mathcal{X}_{\text{pos}}$ denote any subset of $\mathcal{X}$ such that when $\text{PO}(\mathcal{X}_{\text{pos}})$ holds, then β^* is identified. We will assume that $\mathcal{X}_{\text{pos}}$ is discrete (finite or countable) or otherwise that it is a subset of $\mathbb{R}^s$ without isolated points.

We refer to models defined by the C, SR and $\text{PO}(\mathcal{X}_{pos})$ assumptions and the restrictions (10) / (12) as dynamic regime parametric/semiparametric marginal structural mean models (DYR-Par-MSM and DYR-Sem-MSM).

5 Inference

By definition, under model DYR-Par-MSM, assumptions C, SR and $\text{PO}(\mathcal{X}_{pos})$ hold. Thus, in view of (5), the model is equivalently defined by these assumptions and the restriction

$$E \left[\omega_K^x (\overline{O}_K, \overline{A}_K) \{u(O, A) - h_{\text{par}}(x, Z; \beta^*)\} | Z \right] = 0 \text{ for all } x \in \mathcal{X}_{pos} \quad (14)$$

where $\omega_K^x (\overline{O}_K, \overline{A}_K)$ is defined like $\omega_K (\overline{O}_K, \overline{A}_K)$ in (4) but with g_x instead of g .

Likewise, model DYR-Sem-MSM is defined by assumptions C, SR and $\text{PO}(\mathcal{X}_{pos})$ and

$$E \left[\omega_K^x (\overline{O}_K, \overline{A}_K) \{u(O, A) - h_{\text{sem}}(x, Z; \beta^*)\} | Z \right] = q(Z) \quad (15)$$

for all $x \in \mathcal{X}_{pos}$, $q(\cdot)$ unknown.

Assumptions C and SR do not impose restrictions on the law P^{marg} of (O, A) (Gill, van der Laan and Robins, 1997). Thus, as models for P^{marg} , models DYR-Par-MSM and DYR-Sem-MSM are defined as the set of laws P^{marg} satisfying the $\text{PO}(\mathcal{X}_{pos})$ assumption and the restrictions (14) and (15) respectively. Note that, although not immediately apparent, (15) does indeed impose restrictions on P^{marg} because the random variables $\omega_K^x (\overline{O}_K, \overline{A}_K) \times \{u(O, A) - h_{\text{sem}}(x, Z; \beta^*)\}$ depend on x yet the condition (15) imposes that they have the same conditional expectation given Z regardless of the value of x .

Because the data available to us are a random sample of (O, A) from the law P^{marg} , then to conduct inference about β^* we can ignore the fact that β^* has a causal interpretation, and regard it simply as a parameter in the model for P^{marg} determined by the restriction (14) or (15) under consideration and the $\text{PO}(\mathcal{X}_{pos})$ assumption. We take this point of view in our construction of estimators of β^* and denote the models for the law P^{marg} determined by the $\text{PO}(\mathcal{X}_{pos})$ assumption and restriction (14) / (15) as DYR-Par-MSM-obs and DYR-Sem-MSM-obs models.

Restrictions (14) and (15) are not directly applicable for construction of estimators of β^* because the weights $\omega_K^x (\overline{O}_K, \overline{A}_K)$ depend on the unknown

conditional treatment probabilities $\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1})$, $k = 0, \dots, K$, defined in (3) which cannot be non-parametrically estimated well using smoothing techniques due to the curse of dimensionality. One strategy for dimension reduction, that we adopt in this paper, is to consider estimation of β^* under models DYR-Par-MSM-obs* and DYR-Sem-MSM-obs* defined like DYR-Par-MSM-obs and DYR-Sem-MSM-obs but with the additional assumption

$$\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1}) = \lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1}; \gamma^*), k = 0, \dots, K \quad (16)$$

where for each $a_k, \bar{o}_k$ and $\bar{a}_{k-1}$, $\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1}; \cdot)$ is a known smooth function and γ^* is an unknown $q \times 1$ parameter vector. For instance if $\mathcal{A}_k = \{0, 1\}$, we may consider a logistic regression model

$$\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1}; \gamma) = \begin{cases} \frac{\exp\{\gamma'_k e_k(\bar{o}_k, \bar{a}_{k-1}) a_k\}}{1 + \exp\{\gamma'_k e_k(\bar{o}_k, \bar{a}_{k-1})\}} & \text{if } r_k = 1 \\ I_{\{a_{k-1}\}}(a_k) & \text{if } r_k = 0 \end{cases}$$

for some $q \times 1$ -vector-valued, user-specified, function $e_k(\bar{o}_k, \bar{a}_{k-1})$ and $\gamma' = (\gamma_0, \dots, \gamma_{q-1})$. The choice of $I_{\{a_{k-1}\}}(a_k)$ for λ_k when $r_k = 0$ is made in accordance with our convention for assigning values to A_k corresponding to times k after the event has occurred, which stipulates precisely that A_k is equal to A_{k-1} if the event has already occurred by time k . We define λ_k in this way so that when we compute the denominator of $\omega_K^x(\bar{O}_K, \bar{A}_K)$ from the data recorded in the study database, the factors $\lambda_k(A_k|\bar{O}_k, \bar{A}_{k-1}; \gamma)$ for the times k after the event has occurred will be equal to 1, and hence will not contribute to the calculation of $\omega_K^x(\bar{O}_K, \bar{A}_K)$.

5.1 Estimators of β^* and of $x_{opt}(z)$

To define our estimators of β^* we will use the following conventions: the expression $\mu - a.e. (\mathcal{X}_{pos})$ next to a statement stands for the assertion that the statement holds for all x in $\mathcal{X}_{pos}$ except on a subset of $\mathcal{X}_{pos}$ of μ -measure 0, where μ is the counting measure if $\mathcal{X}_{pos}$ is discrete (finite or countable) and μ is the Lebesgue measure on $\mathbb{R}^s$ otherwise; $\mathcal{G}$ stands for the sigma field of all subsets of $\mathcal{X}_{pos}$ if $\mathcal{X}_{pos}$ is discrete, and for the Borel sigma-field of $\mathcal{X}_{pos}$ otherwise; P_X denotes a measure on $(\mathcal{X}_{pos}, \mathcal{G})$ which is mutually absolutely continuous with the restriction $\mu_{\mathcal{X}_{pos}}$ of μ to $\mathcal{X}_{pos}$. Even though our derivations will not require that P_X be a specific one, when $\mathcal{X}_{pos}$ is bounded, we will take P_X to be the uniform measure. The formalism of defining a generic P_X is needed to be able to accommodate settings with unbounded sets $\mathcal{X}_{pos}$. Nevertheless we will argue in Remark 2 below that the choice of P_X is essentially inconsequential.

For any function $b(x, z)$, we define

$$b_{par}(x, z) \equiv b(x, z) \text{ and } b_{sem}(x, z) \equiv b(x, z) - \int_{\mathcal{X}_{pos}} b(x, z) dP_X(x).$$

Because our estimators and their properties can be derived arguing essentially identically in both models DYR-Par-MSM-obs* and DYR-Sem-MSM-obs*, then to avoid redundancy, throughout we use the subscript $\cdot$ in every instance in which the definitions and results being derived hold indistinctly if $\cdot$ is replaced by *par* or by *sem*.

For $\cdot = par$ and $\cdot = sem$ define

$$\begin{aligned} U(x; \beta, b) &\equiv b(x, Z) \{u(O, A) - h(x, Z; \beta)\} \\ S(\beta, \gamma, b) &\equiv \int_{\mathcal{X}_{pos}} \omega_K^x(\gamma) U(x; \beta, b) dP_X(x) \end{aligned}$$

and

$$\begin{aligned} S_{aug}(\gamma, d) &\equiv \sum_{k=0}^K \sum_{a_k \in \mathcal{A}_k} \{I_{\{a_k\}}(A_k) - \lambda_k(a_k | \bar{O}_k, \bar{A}_{k-1}, \gamma)\} d_k(\bar{O}_k, a_k, \bar{A}_{k-1}) \\ &= \sum_{k=0}^K \{d_k(\bar{O}_k, \bar{A}_k) - E_\gamma[d_k(\bar{O}_k, \bar{A}_k) | \bar{O}_k, \bar{A}_{k-1}]\} \end{aligned}$$

where $b(\cdot, \cdot)$ is any, possibly vector valued, function for which the integrals where it intervenes are well defined, $d_k(\cdot, \cdot)$ is an arbitrary, possibly vector-valued, function, $E_\gamma[\cdot]$ denotes, for the adequate k , conditional expectation under the law $\lambda_k(a_k | \bar{o}_k, \bar{a}_{k-1}; \gamma)$,

$$\omega_k^x(\bar{o}_k, \bar{a}_k; \gamma) \equiv \frac{\prod_{j=0}^k I_{\{g_{x,j}(\bar{o}_j)\}}(a_j)}{\prod_{j=0}^k \lambda_j(a_j | \bar{o}_j, \bar{a}_{j-1}; \gamma)}, \quad 0 \leq k \leq K$$

and

$$\omega_k^x(\gamma) \equiv \omega_k^x(\bar{O}_k, \bar{A}_k; \gamma).$$

The following Lemma, whose proof is given in Section 2.3 of the companion paper (Orellana, Rotnitzky and Robins, 2010), provides the key result for the derivation of our estimators.

Lemma 2. Under (16)

1. $E[S_{aug}(\gamma^*, d)] = 0$ for all functions d ;
2. restriction (14) holds $\mu - a.e. (\mathcal{X}_{pos})$ if and only if $E[S_{par}(\beta^*, \gamma^*, b)] = 0$ for all b and,
3. restriction (15) holds $\mu - a.e. (\mathcal{X}_{pos})$ if and only if $E[S_{sem}(\beta^*, \gamma^*, b)] = 0$ for all b .

We are now ready to describe our estimators of β^* and of $x_{opt}(z)$. Our estimators require two inputs from the data analyst, the first input is a (column) vector-valued function $b(x, z)$ of the same dimension as β^* , the second input is a collection of (column) vector-valued functions $\{d_k(\bar{o}_k, \bar{a}_k)\}_{k=0, \dots, K}$, each function also being of the same dimension as β^* . Each choice of b and $\{d_k\}_{k=0, \dots, K}$ will result in a different estimator which we denote with $\hat{\beta}_{par}(b, d)$ or $\hat{\beta}_{sem}(b, d)$ depending on the model under consideration. Different choices of b and $\{d_k\}_{k=0, \dots, K}$ result in estimators with limiting mean zero normal distributions but with different asymptotic variances. The following algorithm describes the steps required to compute $\hat{\beta}_{par}(b, d)$ and $\hat{\beta}_{sem}(b, d)$.

Stage 1: compute the maximum likelihood estimator $\hat{\gamma}$ of γ^* solving $\mathbb{P}_n\{S_\gamma(\gamma)\} = 0$ where $S_\gamma(\gamma) \equiv \frac{\partial}{\partial \gamma} \mathcal{L}_A(\gamma)$, $\mathcal{L}_A(\gamma) \equiv \prod_{j=0}^K \lambda_j(A_j | \bar{O}_j, \bar{A}_{j-1}; \gamma)$.

Stage 2: if model DYP-Par-MSM-obs* was assumed, compute $\hat{\beta}_{par}(b, d)$ solving

$$\mathbb{P}_n\{S_{par}(\beta, \hat{\gamma}, b) - S_{aug}(\hat{\gamma}, d)\} = 0 \quad (17)$$

and if model DYP-Sem-MSM-obs* was assumed compute $\hat{\beta}_{sem}(b, d)$ solving

$$\mathbb{P}_n\{S_{sem}(\beta, \hat{\gamma}, b) - S_{aug}(\hat{\gamma}, d)\} = 0. \quad (18)$$

In the preceding algorithm we adopted the convention, that we will also adopt throughout, that for n i.i.d. copies $V_1, \dots, V_n$ of any random vector V , $\mathbb{P}_n(V)$ stands for $n^{-1} \sum_{i=1}^n V_i$.

The special case in which the functions d_k are all identically 0 yields $S_{aug}(\hat{\gamma}, d) = 0$ so the estimators $\hat{\beta}_{par}(b, d)$ and $\hat{\beta}_{sem}(b, d)$ solve respectively the reduced equations

$$\mathbb{P}_n\{S_{par}(\beta, \hat{\gamma}, b)\} = 0 \text{ and } \mathbb{P}_n\{S_{sem}(\beta, \hat{\gamma}, b)\} = 0. \quad (19)$$

To simplify notation, we will use $\hat{\beta}_{par}(b)$ and $\hat{\beta}_{sem}(b)$ to denote the estimators solving these reduced equations. Lemma 2 implies that when model (16)

is correctly specified $\widehat{\beta} \cdot (b, d)$ and $\widehat{\beta} \cdot (b)$ are, under regularity conditions and the corresponding model (10) or (12), consistent and asymptotically normal estimators of β^* .

The left hand sides of the equations in displays (17) and (18) are the difference of two terms, the term $\mathbb{P}_n \{S_{par}(\beta, \widehat{\gamma}, b)\}$ and the term $\mathbb{P}_n \{S_{aug}(\widehat{\gamma}, d)\}$. The first term is a sum (or integral) over $x \in \mathcal{X}_{pos}$ of x -specific estimating functions $\mathbb{P}_n \{\omega_K^x(\widehat{\gamma}) U_{par}(x; \beta, b)\}$ and $\mathbb{P}_n \{\omega_K^x(\widehat{\gamma}) U_{sem}(x; \beta, b)\}$. These x -specific estimating functions are sums only over subjects whose data are consistent with having followed regime g_x and their contribution to the sum is equal to the functions $U_{par}(x; \beta, b)$ and $U_{sem}(x; \beta, b)$, weighted by the inverse of the product of the occasion-specific estimated conditional probabilities that at each occasion they took the treatment they actually took given their recorded history of outcomes and treatments at that occasion. Subjects that did not follow regime g_x contribute indirectly to these x -specific estimating functions through the estimate $\widehat{\gamma}$ of γ in the model for the treatment probabilities. In contrast, the second term $\mathbb{P}_n \{S_{aug}(\widehat{\gamma}, d)\}$ is a sum over all subjects, regardless of whether or not they followed any given regime g_x . This term does not depend on β , but with clever choices of functions $\{d_k\}_{k=0, \dots, K}$, its inclusion can help improve the efficiency with which β^* is estimated. This point is elaborated in Section 5.3.

Under (10) and (12), the functions $U_{par}(x; \beta, b)$ and $U_{sem}(x; \beta, b)$ respectively would have mean zero at β^* in a study in which subjects are randomized to the sequence of treatments $A_k, k = 0, \dots, K$, at baseline (with randomization probabilities that may depend on Z) but do not generally have mean zero in an observational study. Weighting by $\omega_K^x(\overline{O}_K, \overline{A}_K)$ creates pseudo-samples from a pseudo-population in which all subjects followed regime g_x . Thus, weighting by an estimator of $\omega_K^x(\overline{O}_K, \overline{A}_K)$ and summing over all $x \in \mathcal{X}_{pos}$, implicitly removes bias by creating, for large samples, a pseudo-sample from a pseudo-population in which subjects were assigned with equal probability to each of the regimes $g_x, x \in \mathcal{X}_{pos}$ (Hernán, Brumback and Robins, 2000).

In accordance with similar terminology used in estimation of standard, i.e. non-dynamic, marginal structural models (Robins, 1999) we refer to the estimator $\widehat{\beta} \cdot (b)$ as an inverse probability of treatment (IPTW) estimator and to the estimator $\widehat{\beta} \cdot (b, d)$ (with non-zero d_k 's) as an augmented inverse probability of treatment (AIPTW) estimator.

A point of note is that a given subject may follow more than one regime g_x and thus contribute a term to more than one x -specific estimating function $\mathbb{P}_n \{\omega_K^x(\widehat{\gamma}) U \cdot (x; \beta, b)\}$. This is a special feature of AIPTW estimating equations in dynamic-regime marginal mean models not shared by the IPTW

estimating equations in standard marginal mean models (see continuation of example 1 below for an illustration of this point).

Given $\hat{\beta}_{\cdot}(b, d)$, we estimate $x_{opt}(z)$ with $x_{opt}\left(z; \hat{\beta}_{\cdot}(b, d)\right)$ where for each β ,

$$x_{opt}(z; \beta) \equiv \arg \max_{x \in \mathcal{X}} h.(z, x; \beta). \quad (20)$$

Remark 2: The choice of measure P_X is inconsequential in the sense that the class of estimators $\left\{\hat{\beta}_{\cdot}(b, d) : b \text{ and } d \text{ unrestricted}\right\}$ is the same for any pair of mutually absolutely continuous measures P_X and P'_X . This holds because the estimator $\hat{\beta}_{\cdot}(b, d)$ that uses P_X is algebraically identical to the estimator $\hat{\beta}_{\cdot}(b', d)$ where $b' = b \times dP'_X/dP_X$.

Example 1 (continuation): returning to example 1 of Section 2.2 suppose that our goal is to find the optimal CD4 level at which to start treating HIV+ subjects with HAART. Suppose that for each integer x in the interval $[200, 600]$ there exist subjects in the population which started HAART as soon as it was detected that their CD4 count fell below x . In such case $\mathcal{X}_{pos}$ is the set of integers in $[200, 600]$ while $\mathcal{X}$ is the entire interval $[200, 600]$. In this example, we take P_X to be the uniform probability over the integers in $[200, 600]$, so integrals over $\mathcal{X}_{pos}$ with respect to P_X are simply averages over the integers x in $[200, 600]$. Suppose that the endpoint of interest, i.e. the utility, is defined to be equal to the time since baseline to the first occurrence of either death from any cause or diagnosis of clinical AIDS, if at least one of these two events occurred during the 37 months of follow-up, and to be equal to 72 months otherwise. The value of 72 months is chosen to represent the expected time since baseline to the occurrence of death or first diagnosis of clinical AIDS for subjects that are still alive and AIDS free at month 37. Thus, we write

$$u(O, A) = 72R_{K+1} + T_{K+1}(1 - R_{K+1})$$

Furthermore, suppose that $h_{\text{par}}(z, x; \beta)$ and $h_{\text{sem}}(z, x; \beta)$ are like in (11) and (13). To compute IPTW and AIPTW estimators of β^* we need to postulate first a parametric model $\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1}; \gamma)$ for the treatment probabilities $\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1})$. For example, we may choose

$$\lambda_k(a_k|\bar{o}_k, \bar{a}_{k-1}; \gamma) = \begin{cases} \frac{\exp\{\gamma'_k e_k(\bar{o}_k, \bar{a}_{k-1})a_k\}}{1 + \exp\{\gamma'_k e_k(\bar{o}_k, \bar{a}_{k-1})\}} & \text{if } r_k = 1 \text{ and } a_{k-1} = 0 \\ I_{\{a_{k-1}\}}(a_k) & \text{if } r_k = 1 \text{ and } a_{k-1} = 1 \\ I_{\{a_{k-1}\}}(a_k) & \text{if } r_k = 0 \end{cases}$$

where

$$e_k(\bar{O}_k, \bar{A}_{k-1}) = (1, \text{school}, \text{hmo}, \text{druguse}, \text{year}, \text{age}, \text{age}^2, \\ k, k^2, k^3, CD4_k, CD4_k^2, CD4_k^3, V_k, V_k^2, V_k^3)$$

with *school*, *hmo* and *druguse* being baseline binary indicators of completion of high school, affiliation with an HMO and drug use respectively, *year* is the calendar year of start of follow-up, $CD4_k$ and V_k , recall, are the CD4 count and the log-viral load at time k respectively. The choice of λ_k when $r_k = 1$ and $a_{k-1} = 1$ reflects the fact that in our example the probability that $A_k = 1$ is 1 when $R_k = 1$ and $A_{k-1} = 1$ because we have assumed that once an assignment to take HAART was made, it remained in place for the rest of the follow-up period. As discussed earlier in this Section, the choice of $I_{\{a_{k-1}\}}(a_k)$ when $r_k = 0$ is made so that there is no contribution to the product in the denominator of $\omega_K(\bar{O}_K, \bar{A}_K)$ after the event of interest has occurred.

Consider next the two stages of the algorithm to compute the IPTW and AIPTW estimators. In stage 1, the estimator $\hat{\gamma}$ of γ^* is effectively computed by pooled logistic regression with multiple contributions from each subject. Subject i contributes to the pooled logistic regression with the outcome $A_{k,i}$ and the covariates $e_k(\bar{O}_{ki}, \bar{A}_{k-1,i})$ for all times k such that a) $R_{k,i} = 1$ and b) $A_{k-1,i} = \dots = A_{0,i} = 0$.

To carry out stage 2 for IPTW estimation we must input a vector function $b(z, x)$. A simple choice for b in the model DYR-Par-MSM is $b(z, x) = \frac{\partial}{\partial \beta} h_{\text{par}}(z, x; \beta)$ which under model (11) yields

$$b(z, x) = (x - x^*, (x - x^*)^2, (x - x^*)z, (x - x^*)^2z, 1, z).$$

The IPTW estimator of $\beta = (\beta_1, \dots, \beta_6)$ in model (11) that uses this choice of b , then solves

$$0 = \sum_{i=1}^n \left[\sum_{x=200}^{600} \omega_{K,i}^x(\hat{\gamma}) \begin{pmatrix} x - x^* \\ (x - x^*)^2 \\ (x - x^*)Z_i \\ (x - x^*)^2Z_i \\ 1 \\ Z_i \end{pmatrix} \right] \times \\ \{72R_{K+1,i} + T_{K+1,i}(1 - R_{K+1,i}) - \beta_1(x - x^*) - \beta_2(x - x^*)^2 \\ - \beta_3(x - x^*)Z_i - \beta_4(x - x^*)^2Z_i - \beta_5 - \beta_6Z_i\}$$

where x^* is, say, 350.

If, instead, we wish to conduct estimation under the more flexible model DYR-Sem-MSM, then the choice $b(z, x) = \frac{\partial}{\partial \beta} h_{\text{sem}}(z, x; \beta)$ yields under model (13), $b(x, z) = (x - x^*, (x - x^*)^2, (x - x^*)z, (x - x^*)^2 z)'$. The IPTW estimator of $\beta = (\beta_1, \dots, \beta_4)$ now solves,

$$\sum_{i=1}^n \left[\sum_{x=200}^{600} \omega_{K,i}^x(\hat{\gamma}) \begin{pmatrix} x - \bar{x} \\ (x - x^*)^2 - \overline{(x - x^*)^2} \\ (x - \bar{x}) Z_i \\ \left\{ (x - x^*)^2 - \overline{(x - x^*)^2} \right\} Z_i \end{pmatrix} \right. \\ \left. \times \{ 72R_{K+1,i} + T_{K+1,i} (1 - R_{K+1,i}) - \beta_1 (x - x^*) \right. \\ \left. - \beta_2 (x - x^*)^2 - \beta_3 (x - x^*) Z_i - \beta_4 (x - x^*)^2 Z_i \} \right] = 0$$

where an overbar $\bar{\cdot}$ stands for average over the integers x in $[200, 600]$.

To construct augmented IPTW estimators of β^* we must choose functions $\{d_k(\bar{o}_k, \bar{a}_k)\}_{k=0, \dots, K}$ for each k . We postpone the discussion and illustration of convenient choices of d_k to Section 5.3.1.

Finally to illustrate the point made earlier that a subject may contribute to several x -specific estimating functions, consider a subject i whose CD4 count decreased over time, who started HAART at a CD4 count of 250 in some visit and whose previous visit CD4 count was 300. Then this subject's observed data was consistent with having followed regime g_x for all $x \in [250, 300]$. For such subject, $\omega_{K,i}^x(\hat{\gamma}) \neq 0$ if $x \in [250, 300]$ and $\omega_{K,i}^x(\hat{\gamma}) = 0$ otherwise. Thus, subject i enters in the estimating equation for β repeatedly over the sum on x , once for every regime g_x his data was consistent with. Furthermore, suppose that the subject died in the middle of the third month. Then for this subject $R_{K+1,i} = 0$ and $T_{K+1,i} = 2.5$ so his outcome is 2.5. In addition, $I_{\{g_{x,k}(\bar{o}_{k,i})\}}(A_{k,i}) / \lambda_k(A_{k,i} | \bar{O}_{k,i}, \bar{A}_{k-1,i}) = 1$ for $k = 3, \dots, K$ and consequently, $\omega_{K,i}^x(\hat{\gamma}) = \omega_{2,i}^x(\hat{\gamma}) \cdot \diamond$

5.2 Asymptotic Distribution of the IPTW and AIPTW Estimators of β^*

In this section we will derive the asymptotic distribution of the IPTW and AIPTW estimators of β^* proposed in the preceding subsection under the assumptions of model DYR-Par-MSM-obs*. Our results will require that we make the following strengthening of the $\text{PO}(\mathcal{X}_{\text{pos}})$ assumption which states that treatment probabilities of treatments consistent with regimes g_x not only are positive but also stay bounded away from zero.

Strengthened PO($\mathcal{X}_{pos}$) assumption: there exists $\sigma > 0$ such that for all $x \in \mathcal{X}_{pos}$, and all $k = 0, \dots, K$, $\mathbb{P} [\lambda_k (A_k^g | \overline{O}_k^g, \overline{A}_{k-1}^g) > \sigma] = 1$.

Under regularity conditions and the strengthened PO($\mathcal{X}_{pos}$) assumption, standard linearization arguments imply that when model DYR---MSM-obs* holds,

$$\sqrt{n} \left\{ \widehat{\beta}_{\cdot} (b, d) - \beta^* \right\} = \sqrt{n} J_{\cdot} (b)^{-1'} \mathbb{P}_n \{ M_{\cdot}^* (b, d) \} + o_p(1) \quad (21)$$

where

$$\begin{aligned} M_{\cdot}^* (b, d) &\equiv M_{\cdot} (b, d) - E \{ M_{\cdot} (b, d) S_{\gamma}' \} E \{ S_{\gamma}^{\otimes 2} \}^{-1} S_{\gamma}, \\ M_{\cdot} (b, d) &\equiv S_{\cdot} (\beta^*, \gamma^*, b) - S_{aug} (\gamma^*, d), \quad S_{\gamma} \equiv S_{\gamma} (\gamma^*), \\ J_{\cdot} (b) &\equiv -E \left\{ \frac{\partial}{\partial \beta} \int_{\mathcal{X}_{pos}} b(x, Z) h_{\cdot} (x, Z; \beta) dP_X(x) \right\} \Big|_{\beta=\beta^*} \end{aligned} \quad (22)$$

and for any $V, V^{\otimes 2} \equiv VV'$. Consequently,

$$\sqrt{n} \left\{ \widehat{\beta}_{\cdot} (b, d) - \beta^* \right\} \rightarrow N(0, \Omega_{\cdot} (b, d)) \quad (23)$$

where

$$\Omega_{\cdot} (b, d) = J_{\cdot} (b)^{-1'} E \left(M_{\cdot}^* (b, d)^{\otimes 2} \right) J_{\cdot} (b)^{-1}$$

which can be consistently estimated with $\widehat{J}_{\cdot} (b)' \mathbb{P}_n \left(\widehat{M}_{\cdot}^* (b, d)^{\otimes 2} \right) \widehat{J}_{\cdot} (b)$ where $\widehat{J}_{\cdot} (b)$ and $\widehat{M}_{\cdot}^* (b, d)$ are computed like $J_{\cdot} (b)$ and $M_{\cdot}^* (b, d)$ but with the population mean E replaced by the empirical mean $\mathbb{P}_n$ and with all quantities evaluated at $(\widehat{\beta}_{\cdot} (b, d), \widehat{\gamma})$ instead of at (β^*, γ^*) .

In the following sections we will derive the optimal functions b_{opt} and $d_{\cdot, opt, k}^b$ for the functions $d_k, k = 0, \dots, K$, that yield $\widehat{\beta}_{\cdot} (b_{opt}, d_{\cdot, opt})$ with smallest asymptotic variance among all AIPTW estimators $\widehat{\beta}_{\cdot} (b, d)$ under model DYR---MSM-obs*. Our plan is as follows. In the next subsection we will derive, for each fixed b , the optimal choices $d_{\cdot, opt, k}^b$ for the functions $d_k, k = 0, \dots, K$, in the sense that

$$\Omega_{\cdot} (b, d) \geq \Omega_{\cdot} (b, d_{\cdot, opt}^b) \text{ for all } d \text{ (} b \text{ fixed)}$$

where for any squared matrices Ω_1 and $\Omega_2, \Omega_1 \geq \Omega_2$ stands for $\Omega_1 - \Omega_2$ is positive semidefinite. The optimal functions $d_{\cdot, opt}^b$ are not available for data analysis because they depend on the unknown law of the data, however, we will indicate how to use the knowledge of their functional form to derive locally efficient estimators of β^* in model DYR---MSM-obs* for the class of AIPTW

estimators that use a fixed b (a property that is defined at the end of Section 5.3.1). In a subsequent section, we will show that these locally efficient estimators are also double-robust, a property that we will define then.

In Section 7 we will derive the optimal function $b_{,opt}$ in the sense that

$$\Omega.(b, d_{,opt}^b) \geq \Omega.(b_{,opt}, d_{,opt}^{b_{,opt}}) \text{ for all } b. \quad (24)$$

Thus, under model DYR---MSM-obs*, the choice $b_{,opt}, d_{,opt,k}$ where $d_{,opt,k} = d_{,opt,k}^{b_{,opt}}$ is optimal in the sense that $\hat{\beta}.(b_{,opt}, d_{,opt})$ has the smallest asymptotic variance among the asymptotic variances of all $\hat{\beta}.(b, d)$. These optimal functions are also not available for data analysis as they depend on the unknown observed data law, but we will show how to exploit the knowledge of their form to compute estimators of β^* that (i) are locally efficient under model DYR---MSM-obs* for the class of all AIPTW estimators and (ii) are double-robust.

5.3 Efficient Estimation for a Fixed Choice of b

To derive the optimal functions $d_{,opt,k}^b, k = 0, \dots, K$, we reason as follows.

1. Since $J.(b)$ does not depend on d , $d_{,opt}^b$ must minimize $E(M^*(b, d)^{\otimes 2})$ over all d .
2. Define the set

$$\begin{aligned} \Lambda &\equiv \{S_{aug}(\gamma^*, d) : d_k, k = 0, \dots, K \text{ arbitrary scalar functions}\} \\ &= \left\{ \sum_{k=0}^K \{d_k(\bar{O}_k, \bar{A}_k) - E_{\gamma^*}[d_k(\bar{O}_k, \bar{A}_k) | \bar{O}_k, \bar{A}_{k-1}]\} : \right. \\ &\quad \left. d_k, k = 0, \dots, K, \text{ are arbitrary scalar functions} \right\}. \end{aligned}$$

The second expression for the set indicates that Λ is comprised by sums from $k = 0$ to K of functions of $(\bar{O}_k, \bar{A}_k)$ with mean zero given $(\bar{O}_k, \bar{A}_{k-1})$ under $\lambda_k(a_k | \bar{O}_k, \bar{A}_{k-1}; \gamma^*)$.

3. Suppose that model (16) is correctly specified. Consider the score vector S_γ for γ evaluated at the true parameter γ^* . The j^{th} entry of S_γ is equal to $\sum_{k=0}^K \partial \log \lambda_k(A_k | \bar{O}_k, \bar{A}_{k-1}; \gamma) / \partial \gamma_j^1|_{\gamma=\gamma^*}$. Since each term $\partial \log \lambda_k(A_k | \bar{O}_k, \bar{A}_{k-1}; \gamma) / \partial \gamma_j^1|_{\gamma=\gamma^*}$ is a function of $(\bar{O}_k, \bar{A}_k)$ with mean zero given $(\bar{O}_k, \bar{A}_{k-1})$, then every entry of the vector S_γ is an element of the set Λ .

4. It follows from item 3 that $E \{M. (b, d) S'_\gamma\} E \{S_\gamma^{\otimes 2}\}^{-1} S_\gamma$, being a vector whose entries are linear combinations of the entries of S_γ , can indeed be written as $S_{aug}(\gamma^*, d_{ML})$ for some vector valued functions $\{d_{ML,k}\}_{k=0,\dots,K}$ of the same dimension as β^* .
5. If we could find $S_{aug}(\gamma^*, d_{\cdot, opt}^b)$ such that each entry of the column vector $M^*(b, d_{\cdot, opt}^b) = S.(\beta^*, \gamma^*, b) - S_{aug}(\gamma^*, d_{\cdot, opt}^b)$ was uncorrelated with all the elements of Λ , then we would have that

$$E \left(M^*(b, d) M^*(b, d_{\cdot, opt}^b)' \right) = E \left(M^*(b, d_{\cdot, opt}^b)^{\otimes 2} \right) \quad (25)$$

and consequently we would obtain the desired inequality

$$\begin{aligned} E \left(M^*(b, d)^{\otimes 2} \right) - E \left(M^*(b, d_{\cdot, opt}^b)^{\otimes 2} \right) &= \\ E \left(\{ M^*(b, d) - M^*(b, d_{\cdot, opt}^b) \}^{\otimes 2} \right) &\geq 0. \end{aligned}$$

That (25) is a consequence of $M^*(b, d_{\cdot, opt}^b)$ and Λ being uncorrelated follows from

$$\begin{aligned} M^*(b, d) &= S.(\beta^*, \gamma^*, b) - S_{aug}(\gamma^*, d) - E \{M. (b, d) S'_\gamma\} E \{S_\gamma^{\otimes 2}\}^{-1} S_\gamma \\ &= S.(\beta^*, \gamma^*, b) - S_{aug}(\gamma^*, d) - S_{aug}(\gamma^*, d_{ML}) \\ &= M^*(b, d_{\cdot, opt}^b) + \\ &\quad + \{ S_{aug}(\gamma^*, d_{\cdot, opt}^b) - S_{aug}(\gamma^*, d) - S_{aug}(\gamma^*, d_{ML}) \} \end{aligned}$$

where the first equality follows by (22) and the second by item 4. In the third equality the term between curly brackets is in Λ and hence it is uncorrelated with $M^*(b, d_{\cdot, opt}^b)$, thus yielding (25).

6. A random vector $S_{aug}(\gamma^*, d_{\cdot, opt}^b)$ with the properties of item 5 indeed exists. To find it we reason as follows. The set Λ is a linear and closed subspace of the (Hilbert) space of all mean zero, finite variance, random functions of (O, A) with covariance inner product. By the projection theorem (Luenberger, Ch. 3, Theorem 2, 1969), for any $q \times 1$ random vector Q , there exists a (mean zero, finite variance) $q \times 1$ random vector denoted $\Pi[Q|\Lambda]$, such that each entry of $\Pi[Q|\Lambda]$ is in Λ and such that $Q - \Pi[Q|\Lambda]$ is uncorrelated with all the elements of Λ . The vector $\Pi[Q|\Lambda]$ is called the projection of Q onto Λ . Because of the special form of Λ (Λ is a direct sum of K subspaces, the k^{th} one being the set of functions of $(\overline{O}_k, \overline{A}_k)$ with mean zero given $(\overline{O}_k, \overline{A}_{k-1})$) it is a standard calculation

to compute $\Pi [Q|\Lambda]$ for any Q . For completeness, in Section 2.4.1 of the companion paper (Orellana, Rotnitzky and Robins, 2010) we provide this calculation and show that

$$\Pi [Q|\Lambda] = \sum_{k=0}^K \{E (Q|\bar{O}_k, \bar{A}_k) - E (Q|\bar{O}_k, \bar{A}_{k-1})\}. \quad (26)$$

Applying this result to $Q = S. (\beta^*, \gamma^*, b)$, another straightforward calculation, also derived in Section 2.4.3 of the companion paper, gives that

$$\Pi [S. (\beta^*, \gamma^*, b) |\Lambda] = S_{aug} (\gamma^*, d_{\cdot, opt}^b) \quad (27)$$

where $d_{\cdot, opt, k}^b (\bar{o}_k, \bar{a}_k) \equiv d_{\cdot, \beta^*, \gamma^*, opt, k}^b (\bar{o}_k, \bar{a}_k)$, and where for any β, γ ,

$$\begin{aligned} d_{\cdot, \beta, \gamma, opt, k}^b (\bar{o}_k, \bar{a}_k) &\equiv \int_{\mathcal{X}_{pos}} b. (x, Z) \omega_k^x (\bar{o}_k, \bar{a}_k; \gamma) \\ &\quad \{ \phi_{k+1}^x (\bar{o}_k) - h. (x, Z; \beta) \} dP_X (x), \end{aligned} \quad (28)$$

with ϕ_k^x defined like ϕ_k in (8) but with g_x instead of g .

The preceding argument shows that inclusion of the augmentation term $S_{aug} (\gamma, d)$, with adequately chosen functions d_k , helps improve the efficiency with which one can estimate β^* : estimating β^* with a non-augmented estimating equation is equivalent to solving equations (17) using $d_k = 0$ for all k , but these null functions are generally not the optimal choices for d_k .

5.3.1 Locally Efficient Estimation for a Fixed Choice of b

The functions ϕ_{k+1}^x are unknown and thus the functions $d_{\cdot, \beta, \gamma, opt, k}^b$ are not available for data analysis. To obtain estimators of β^* with good efficiency properties we may consider replacing the unknown functions ϕ_{k+1}^x with estimators of them under some working model. Specifically, we propose carrying out the following algorithm:

a. Postulate a model

$$\phi_{k+1}^x (\bar{o}_k) = \phi_{k+1}^x (\bar{o}_k; \tau^*), k = 0, \dots, K, \quad (29)$$

where for each x and $\bar{o}_k$, $\phi_{k+1}^x (\bar{o}_k; \cdot)$ is a smooth function and τ^* is an $r \times 1$ unknown parameter vector.

b. Estimate τ^* with $\hat{\tau}$ solving $\mathbb{P}_n S_\tau(\tau) = 0$ where

$$S_\tau(\tau) \equiv \int_{\mathcal{X}_{pos}} S_\tau(x, \tau) dP_X(x)$$

and

$$\begin{aligned} S_\tau(x, \tau) \equiv & \left[\prod_{j=0}^K I_{\{g_{x,j}(\bar{O}_j)\}}(A_j) \right] \{u(O, A) - \phi_{K+1}^x(\bar{O}_K; \tau)\} \frac{\partial \phi_{K+1}^x(\bar{O}_K; \tau)}{\partial \tau} + \\ & \sum_{k=0}^{K-1} \left[\prod_{j=0}^k I_{\{g_{x,j}(\bar{O}_j)\}}(A_j) \right] \{\phi_{k+2}^x(\bar{O}_{k+1}; \tau) - \phi_{k+1}^x(\bar{O}_k; \tau)\} \frac{\partial \phi_{k+1}^x(\bar{O}_k; \tau)}{\partial \tau}. \end{aligned}$$

c. Estimate the function $d_{\cdot, \beta, \gamma, opt, k}^b$ with $d_{\cdot, \beta, \gamma, \hat{\tau}, opt, k}^b$ where for each τ , $d_{\cdot, \beta, \gamma, \tau, opt, k}^b$ is defined like $d_{\cdot, \beta, \gamma, opt, k}^b$ but with $\phi_{k+1}^x(\bar{O}_k; \tau)$ instead of the unknown $\phi_{k+1}^x(\bar{O}_k)$.

d. Solve the estimating equation

$$\mathbb{P}_n [S(\beta, \hat{\gamma}, b) - S_{aug}(\hat{\gamma}, d_{\cdot, \beta, \hat{\gamma}, \hat{\tau}, opt}^b)] = 0$$

and call its solution (in a slight abuse of notation) $\hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b)$.

The form of the estimating function in step b) is motivated by the definition (8) of the function $\phi_{k+1}^x(\bar{O}_k)$ which implies that when model (29) is correctly specified,

$$E[\phi_{k+2}^x(\bar{O}_{k+1}; \tau^*) | \bar{O}_k, \bar{A}_k = \bar{g}_k^x(\bar{O}_k)] = \phi_{k+1}^x(\bar{O}_k; \tau^*)$$

and consequently,

$$\begin{aligned} E \left[\left\{ \left. \frac{\partial \phi_{k+1}^x(\bar{O}_k; \tau)}{\partial \tau} \right|_{\tau=\tau^*} \right\} \left\{ \prod_{j=0}^k I_{\{g_{x,j}(\bar{O}_j)\}}(A_j) \right\} \right. \\ \left. \times \{\phi_{k+2}^x(\bar{O}_{k+1}; \tau^*) - \phi_{k+1}^x(\bar{O}_k; \tau^*)\} \right] = 0. \end{aligned}$$

Then $S_\tau(\tau^*)$ has mean zero and hence $\mathbb{P}_n S_\tau(\tau) = 0$ is an unbiased estimating function for τ^* .

The implementation of the estimator $\widehat{\beta} \cdot (b, \widehat{d}_{\cdot, opt}^b)$ is facilitated by noting that the estimating function for $\widehat{\beta} \cdot (b, \widehat{d}_{\cdot, opt}^b)$ simplifies considerably. Specifically, in Section 2.4.3 of the companion paper (Orellana, Rotnitzky and Robins, 2010) we show that $S \cdot (\beta, \gamma, b) - S_{aug}(\gamma, d_{\cdot, \beta, \gamma, \tau, opt}^b)$ can be re-written as

$$\int_{\mathcal{X}_{pos}} b(x, Z) S_{\cdot, \beta}(x; \beta, \gamma, \tau) dP_X(x) \quad (30)$$

where

$$\begin{aligned} S_{\cdot, \beta}(x; \beta, \gamma, \tau) &\equiv \omega_K^x(\gamma) \{u(O, A) - h \cdot (x, Z; \beta)\} \\ &\quad - \sum_{k=0}^K \{\omega_k^x(\gamma) - \omega_{k-1}^x(\gamma)\} \{\phi_{k+1}^x(\overline{O}_k; \tau) - h \cdot (x, Z; \beta)\} \end{aligned} \quad (31)$$

and $\omega_{-1}^x(\gamma) \equiv 1$. Thus, $\widehat{\beta} \cdot (b, \widehat{d}_{\cdot, opt}^b)$ solves

$$\mathbb{P}_n \left[\int_{\mathcal{X}_{pos}} b(x, Z) S_{\cdot, \beta}(x; \beta, \widehat{\gamma}, \widehat{\tau}) dP_X(x) \right] = 0.$$

Standard linearization arguments imply that when model (16) is correctly specified, then under regularity conditions and the strengthened $\text{PO}(\mathcal{X}_{pos})$ assumption,

$$\sqrt{n} \left\{ \widehat{\beta} \cdot (b, \widehat{d}_{\cdot, opt}^b) - \beta^* \right\} \rightarrow N(0, \Omega \cdot (b, d_{\cdot, \lim}^b))$$

where $d_{\cdot, \lim}^b$ is defined like $d_{\cdot, opt}^b$ but with $\phi_{k+1}^x(\overline{O}_k)$ replaced by $\phi_{k+1}^x(\overline{O}_k; \tau_{\lim})$ with $\tau_{\lim}$ the probability limit of $\widehat{\tau}$. In particular, if in addition to model (16), model (29) is also correctly specified, then the asymptotic variance of $\widehat{\beta} \cdot (b, \widehat{d}_{\cdot, opt}^b)$ is equal to $\Omega \cdot (b, d_{\cdot, opt}^b)$.

In summary, when model (16) holds, $\widehat{\beta} \cdot (b, \widehat{d}_{\cdot, opt}^b)$ satisfies:

- a) it is a consistent and asymptotically normal estimator of β^* regardless of whether or not model (29) is correctly specified, and
- b) if model (29) is correctly specified, it has the smallest asymptotic variance among all estimators in the class

$$\mathcal{T}_b = \left\{ \widehat{\beta} \cdot (b, d) : d_k \text{ arbitrary, } k = 0, \dots, K \right\}.$$

In general, suppose that $\mathcal{E}$ is a class of consistent and asymptotically normal estimators of a parameter β under all laws P in $\mathcal{M}$. Suppose that for each $P \in \mathcal{M}$, $\Psi(P)$ is a lower bound for the variance of the limiting distributions (under P) of the estimators in $\mathcal{E}$. Suppose that there exists an estimator in the class $\mathcal{E}$ whose limiting distribution has variance equal to $\Psi(P)$ when $P \in \mathcal{M}_0 \subset \mathcal{M}$. Such estimator is called locally efficient in model $\mathcal{M}$ for the class $\mathcal{E}$ at the submodel $\mathcal{M}_0$. According to this definition, if DYR--MSM-obs** denotes the model that imposes the restrictions of model DYR--MSM-obs* and additionally imposes the strengthened $\text{PO}(\mathcal{X}_{pos})$ assumption, the preceding properties (a) and (b) imply that the estimator $\hat{\beta}_{\cdot}(b, \hat{d}_{\cdot, opt}^b)$ is locally efficient in model DYR--MSM-obs** for the class $\mathcal{T}_b$ at the submodel of DYR--MSM-obs** that imposes the additional restriction (29).

6 Double-Robust Estimation

In the preceding subsection we have seen that the estimator $\hat{\beta}_{\cdot}(b, \hat{d}_{\cdot, opt}^b)$ was consistent and asymptotically normal for β^* if model (16) was correctly specified regardless of whether or not model (29) was correct. In fact, the reverse is also true: $\hat{\beta}_{\cdot}(b, \hat{d}_{\cdot, opt}^b)$ is consistent and asymptotically normal for β^* if model (29) is correctly specified regardless of whether model (16) is correct or not. This remarkable property of $\hat{\beta}_{\cdot}(b, \hat{d}_{\cdot, opt}^b)$ is usually referred to as double-robustness. More precisely, we say that $\hat{\beta}_{\cdot}(b, \hat{d}_{\cdot, opt}^b)$ is double-robust in the union model $\text{DYR--MSM-obs}^{**} \cup \text{DYR--MSM-obs}^{\dagger}$ where $\text{DYR--MSM-obs}^{\dagger}$ is the model defined by the restrictions of model DYR--MSM-obs, the additional restriction (29) and regularity conditions required to ensure the convergence in law of $\hat{\beta}_{\cdot}(b, \hat{d}_{\cdot, opt}^b)$.

The double-robustness property of $\hat{\beta}_{\cdot}(b, \hat{d}_{\cdot, opt}^b)$ follows after applying the general result in Robins, Rotnitzky and van der Laan (Sec 7, Lemma 1, 2000) for deriving double-robust estimators of parameters of models with factorized likelihoods. The conditions of their result are satisfied in our problem and they specialize to the following two conditions,

1. the likelihood factorizes into the product of two parts, $\mathcal{L}_O$ and $\mathcal{L}_A$, and the parameter of interest β depends on the law of the observed data P^{marg} only through laws entering one component, the $\mathcal{L}_O$ -part in our problem, of the factorization (see display (9) and the comment thereafter), and

2. there exists a function, in our case $S.(\beta, \gamma, b)$, that
 - (a) depends on the laws for the $\mathcal{L}_A$ -part of the likelihood through a parameter γ indexing a model for them,
 - (b) depends on the laws for the $\mathcal{L}_O$ -part of the likelihood only through β and
 - (c) it has mean zero at β^* when it is evaluated at the true value γ^* of γ when the model it indexes is correct.

The result in Robins et al. (2000) provides a prescription on how to obtain, from $S.(\beta, \gamma, b)$, a new estimating function, say $T(\beta, \gamma, \tau, b)$, that depends on β, γ and on a parameter τ indexing a model for the laws in $\mathcal{L}_O$ with the following properties:

- I. $T(\beta^*, \gamma^*, \tau, b)$ has mean zero if the model for the $\mathcal{L}_A$ -part of the likelihood is correct regardless of the correct or not specification of the model for the $\mathcal{L}_O$ -part, and
- II. $T(\beta^*, \gamma, \tau^*, b)$ has mean zero if the model for the $\mathcal{L}_O$ -part of the likelihood is correct and τ^* is the true value of τ , regardless of the correct or not specification of the model for the $\mathcal{L}_A$ -part.

According to Robins et. al. (2000) result, the function $T(\beta, \gamma, \tau, b)$ is equal to the difference of $S.(\beta, \gamma, b)$ and its projection under the law $P_{\gamma, \tau}$ on the (non-parametric) tangent space for the laws in the $\mathcal{L}_A$ -part of the likelihood, i.e. the space of scores for all possible parametric submodels for these laws. But this tangent space is precisely the set Λ defined in point (2) of Section 5.3 and we have already seen in point (6) of that section (display (27)) how to compute its projection. Furthermore, in Section 5.3.1, we have shown that $S.(\beta, \gamma, b)$ minus its projection into Λ can be expressed as the function in (30). Thus in our problem, (30) is the function $T(\beta, \gamma, \tau, b)$ satisfying properties (I) and (II).

This result is the essential point driving the double-robustness of $\hat{\beta} \cdot \left(b, \hat{d}_{\cdot, opt}^b \right)$. Specifically, under regularity conditions, $\hat{\beta} \cdot \left(b, \hat{d}_{\cdot, opt}^b \right)$ converges in probability to a solution of $\int_{\mathcal{X}_{pos}} E \left[b(x, Z) S_{\cdot, \beta}(x; \beta, \gamma^\dagger, \tau^\dagger) \right] dP_X = 0$ where $\gamma^\dagger$ and $\tau^\dagger$ are the probability limits of $\hat{\gamma}$ and $\hat{\tau}$. When model (29) is correct, $\tau^\dagger = \tau^*$ and, as just argued, this population moment equation is solved at β^* . It follows that if the population moment equation has a unique solution then $\hat{\beta} \cdot \left(b, \hat{d}_{\cdot, opt}^b \right)$ converges in probability to β^* if model (29) holds even if model (16) is incorrectly

specified. Standard results for solutions of unbiased estimating equations imply that, under regularity conditions, the limiting distribution of $\widehat{\beta} \cdot (b, \widehat{d}_{\text{opt}}^b)$ is normal regardless of the validity of models (29) and (16) thus concluding our argument supporting the double-robustness of $\widehat{\beta} \cdot (b, \widehat{d}_{\text{opt}}^b)$.

In spite of the preceding general result, with the goal of making this article self-contained, we will now show from direct calculations, that $S_{\cdot, \beta}(x; \beta, \gamma, \tau^*)$, and consequently the function defined in (30) evaluated at τ^* , has mean zero if model (29) is correct.

Rearranging terms, we can re-express $S_{\cdot, \beta}(x; \beta, \gamma, \tau)$ as

$$\begin{aligned} S_{\cdot, \beta}(x; \beta, \gamma, \tau) &= \{\phi_1^x(\overline{O}_0; \tau) - h_{\cdot}(x, Z; \beta)\} \\ &\quad + \sum_{k=1}^K \omega_{k-1}^x(\gamma) \{\phi_{k+1}^x(\overline{O}_k; \tau) - \phi_k^x(\overline{O}_{k-1}; \tau)\} \\ &\quad + \omega_K^x(\gamma) \{u(O, A) - \phi_{K+1}^x(\overline{O}_K; \tau)\}. \end{aligned}$$

Now,

$$\begin{aligned} &E[\omega_{k-1}^x(\gamma) \{\phi_{k+1}^x(\overline{O}_k; \tau) - \phi_k^x(\overline{O}_{k-1}; \tau)\} | \overline{O}_{k-1}] \\ &= E \left[\frac{I_{\{\bar{g}_{x,k-1}(\overline{O}_{k-1})\}}(\overline{A}_{k-1}) \{\phi_{k+1}^x(\overline{O}_k; \tau) - \phi_k^x(\overline{O}_{k-1}; \tau)\}}{\prod_{j=0}^{k-1} \lambda_j(g_{x,j}(\overline{O}_j) | \overline{O}_j, \overline{A}_{j-1} = \bar{g}_{x,j-1}(\overline{O}_{j-1}); \gamma)} \middle| \overline{O}_{k-1} \right] \\ &= \frac{E[\{\phi_{k+1}^x(\overline{O}_k; \tau) - \phi_k^x(\overline{O}_{k-1}; \tau)\} | \overline{O}_{k-1}, \overline{A}_{k-1} = \bar{g}_{x,k-1}(\overline{O}_{k-1})]}{\prod_{j=0}^{k-1} \lambda_j(g_{x,j}(\overline{O}_j) | \overline{O}_j, \overline{A}_{j-1} = \bar{g}_{x,j-1}(\overline{O}_{j-1}); \gamma)} \\ &\quad \times \mathbb{P}[\overline{A}_{k-1} = \bar{g}_{x,k-1}(\overline{O}_{k-1}) | \overline{O}_{k-1}]. \end{aligned}$$

The numerator in the fraction of the last equality is equal to 0 when model (29) is correctly specified and τ is equal to the true value τ^* since $\phi_{k+1}^x(\overline{O}_k)$ is, by definition, equal to $E[\phi_{k+1}^x(\overline{O}_k) | \overline{O}_{k-1}, \overline{A}_{k-1} = \bar{g}_{x,k-1}(\overline{O}_{k-1})]$.

Likewise,

$$\begin{aligned} & E \left[\omega_K^x(\gamma) \left\{ u(O, A) - \phi_{K+1}^x(\overline{O}_K; \tau) \right\} \middle| \overline{O}_K \right] \\ &= \frac{E \left[\left\{ u(O, A) - \phi_{K+1}^x(\overline{O}_K; \tau) \right\} \middle| \overline{O}_K, \overline{A}_{K-1} = \overline{g}_{x,K-1}(\overline{O}_{K-1}) \right]}{\prod_{j=0}^{K-1} \lambda_j(g_{x,j}(\overline{O}_j) | \overline{O}_j, \overline{A}_{j-1} = \overline{g}_{x,j-1}(\overline{O}_{j-1}); \gamma)} \\ & \times \mathbb{P} \left[\overline{A}_{K-1} = \overline{g}_{x,K-1}(\overline{O}_{K-1}) | \overline{O}_{K-1} \right] \end{aligned}$$

and

$$E \left[\left\{ \phi_1^x(\overline{O}_0; \tau) - h.(x, Z; \beta^*) \right\} \middle| Z \right]$$

are equal to 0 when model (29) is correctly specified and $\tau = \tau^*$. The equality to 0 of the last display follows from the fact that $h.(x, Z; \beta^*) = E[u(O^{g_x}, A^{g_x}) | Z]$ by the assumption of the dynamic regime MSM model and $E[\phi_1^x(\overline{O}_0) | Z] = E[u(O^{g_x}, A^{g_x}) | Z]$ by (7).

Because the preceding equalities with 0 when model (29) is correct and $\tau = \tau^*$ hold even if $\lambda_j(\cdot | \cdot, \cdot; \gamma)$ are not the true conditional treatment probabilities, we conclude that if model (29) is correctly specified then

$$\int_{\mathcal{X}_{pos}} E[b(x, Z) S_{\cdot, \beta}(x; \beta, \gamma, \tau^*)] dP_X = 0$$

regardless of the value of γ .

In conclusion, the estimator $\widehat{\beta}_{\cdot} \left(b, \widehat{d}_{\cdot, opt}^b \right)$ not only has attractive efficiency properties when model (16) is correct but it also gives the analyst double protection against model misspecification. The following Lemma summarizes the asymptotic properties of $\widehat{\beta}_{\cdot} \left(b, \widehat{d}_{\cdot, opt}^b \right)$. Point i) states its local efficiency property and points ii) and iii) state its double-robustness.

Lemma 3. Suppose model DYR--MSM-obs holds.

- i) if the strengthened positivity assumption $\text{PO}(\mathcal{X}_{pos})$ and models (16) and (29) hold, then under regularity conditions,

$$\sqrt{n} \left\{ \widehat{\beta}_{\cdot} \left(b, \widehat{d}_{\cdot, opt}^b \right) - \beta^* \right\} \rightarrow N \left(0, \Omega_{\cdot} \left(b, d_{\cdot, opt}^b \right) \right);$$

- ii) if the strengthened positivity assumption $\text{PO}(\mathcal{X}_{pos})$ and model (16) hold, then under regularity conditions, $\sqrt{n} \left\{ \widehat{\beta}_{\cdot} \left(b, \widehat{d}_{\cdot, opt}^b \right) - \beta^* \right\}$ converges in law to a mean zero normal distribution;

- iii) if model (29) holds, then under regularity conditions $\sqrt{n} \times \left\{ \hat{\beta} \cdot \left(b, \hat{d}_{\text{opt}}^b \right) - \beta^* \right\}$ converges in law to a mean zero normal distribution.

In practice, successful modelling of the functions $\phi_{k+1}^x(\bar{o}_k)$ will be difficult because: *i*) for each fixed x , a model for $\phi_{k+1}^x(\bar{o}_k)$ generally imposes complicated restrictions for the functions $\phi_{j+1}^x(\bar{o}_j)$ for $j < k$ (for a discussion of this point see Murphy, van der Laan and Robins, 2001), *ii*) few interaction terms between x and some of the components of $\bar{o}_k$ in the model for $\phi_{k+1}^x(\bar{o}_k)$ may fail to capture the complicated ways in which the effect of regime g_x may be modified by the past history $\bar{o}_k$, and *iii*) a correctly specified model for $\phi_{k+1}^x(\bar{o}_k)$ should yield a model for $\phi_1^x(o_0)$ compatible with the MSM model $h.(z, x; \beta)$, i.e. our model for $\phi_{k+1}^x(\bar{o}_k)$ should be such that there exists τ^* that satisfies $E[\phi_1^x(o_0; \tau^*)] = h.(z, x; \beta)$.

Example 1 (continuation): we consider now the construction of a double-robust estimator of β^* in the parametric and semiparametric dynamic regime MSM models $h_{\text{par}}(z, x; \beta^*)$ and $h_{\text{sem}}(z, x; \beta)$ defined like in (11) and (13). To do so we need to specify two models, one for the treatment probabilities, $\lambda_k(a_k | \bar{o}_k, \bar{a}_{k-1})$ and another for the functions $\phi_{k+1}^x(\bar{o}_k)$. The model for $\lambda_k(a_k | \bar{o}_k, \bar{a}_{k-1})$ was discussed in Section 5.1. To guide our choice of the model for $\phi_{k+1}^x(\bar{o}_k)$, we can use the interpretation (a) of this function given at the end of Section 3.2. Specifically,

Case 1) if $r_k = 0$, then by definition, $u(o, a) = t_k$, the event time recorded for occasion k . Consequently, any reasonable model should postulate that $\phi_{k+1}^x(\bar{o}_k; \tau^*) = t_k$ is fixed and known (in particular, it is not a function of unknown parameters τ^*),

Case 2) if $r_k = 1$, then our model should distinguish two cases, depending on whether or not the subject's CD4 count has crossed the threshold x or not. Specifically,

Case 2.1) if $CD4_j > x$ for all $j \leq k$ then in the world in which g_x was implemented, a subject with such CD4 recorded history would have not yet started HAART. The expected utility of such subject would depend on the threshold x determining the future value of CD4 count at which the subject will be put on HAART. Furthermore, a patient whose $r_k = 1$ had not yet experienced the event point and, with the definition of utility given in the example in Section 5.1, his/her utility will be at least k . So, our model for the expected utility should account for this. In addition, it is reasonable to expect that the utility will be predicted by the CD4 status at time k and that present CD4 could possibly be a modifier of the effect of starting HAART in the future at different values of x . One possible model contemplating these

considerations could be

$$\begin{aligned}\phi_{k+1}^x(\bar{o}_k; \tau^*) &= \tau_1 + \tau_2 k + \tau_3 CD4_k + (\tau_4 + \tau_5 k + \tau_6 CD4_k) x \\ &\quad + (\tau_7 + \tau_8 k + \tau_9 CD4_k) x^2 + (\tau_{10} + \tau_{11} k + \tau_{12} CD4_k) x^3\end{aligned}$$

which allows for the utility function to be a cubic polynomial in x with coefficients that depend on present $CD4_k$ and the clinic visit k . Of course, more flexible models could be envisioned which incorporate dependence not only on current CD4 count but on past CD4 counts and on current and past viral loads. Our ability to fit more flexible models is limited by the sample size of the study.

Case 2.2) if $CD4_j > x$ for all $j < k$ and $CD4_k \leq x$ then in the world in which x was implemented, a subject with such CD4 recorded history would start HAART at clinic visit k . A reasonable model would allow for the possibility that the expected utility of such subject is equal to k plus some function that depends on the current CD4 values of the subject and the rate of decline in CD4 at the last interval. As such, we could contemplate postulating that

$$\phi_{k+1}^x(\bar{o}_k; \tau^*) = k + \tau_{13} + \tau_{14} CD4_k + \tau_{15} (CD4_k - CD4_{k-1})$$

Note that the case $CD4_j \leq x$ for $j < k$ is not needed because for such k , $\omega_{k-1}^x(\gamma) = \omega_k^x(\gamma)$ so in the estimating equation (31) the term where $\phi_{k+1}^x(\bar{o}_k; \tau^*)$ appears is 0 because $\omega_k^x(\gamma) - \omega_{k-1}^x(\gamma) = 0$.

So, finally, putting the three cases together, our model would be

$$\begin{aligned}\phi_{k+1}^x(\bar{o}_k; \tau^*) &= \\ &= (1 - r_k) t_k + r_k \prod_{j=0}^k I_{(x, \infty)}(CD4_j) \times \{ \tau_1 + \tau_2 k + \tau_3 CD4_k + \\ &\quad (\tau_4 + \tau_5 k + \tau_6 CD4_k) x + (\tau_7 + \tau_8 k + \tau_9 CD4_k) x^2 + \\ &\quad (\tau_{10} + \tau_{11} k + \tau_{12} CD4_k) x^3 \} + r_k \left[\left(1 - \prod_{j=0}^k I_{(x, \infty)}(CD4_j) \right) I_{[0, x]}(CD4_k) \right] \\ &\quad \times \{ k + \tau_{13} + \tau_{14} CD4_k + \tau_{15} (CD4_k - CD4_{k-1}) \}.\end{aligned}$$

6.1 Double-Robust Inference for β^*

We now derive an estimator of the asymptotic variance of $\hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b)$ which is consistent when either model (16) or model (29) holds. To do so, we simply derive the limiting normal distribution of $\hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b)$ without assuming that

either model is correct and estimate the variance of this limiting distribution with the usual sandwich variance estimator, replacing the unknown model parameters with their estimators. Standard Taylor expansion arguments yield

$$\sqrt{n} \left\{ \widehat{\beta} \cdot \left(b, \widehat{d}_{\cdot, opt}^b \right) - \beta^\dagger \right\} = J. (b)^{-1'} \sqrt{n} \times \mathbb{P}_n \left\{ \int_{\mathcal{X}_{pos}} b(x, Z) Q. (x; \beta^\dagger, \gamma^\dagger, \tau^\dagger) dP_X(x) \right\} + o_P(1) \quad (32)$$

where

$$Q. (x; \beta, \gamma, \tau) \equiv S_{\cdot, \beta}(x; \beta, \gamma, \tau) - J_{\cdot, \gamma}(\beta, \gamma, \tau) I_\gamma(\gamma)^{-1} S_\gamma(\gamma) - J_{\cdot, \tau}(\beta, \gamma, \tau) I_\tau(\tau)^{-1} S_\tau(\tau), \quad (33)$$

$\gamma^\dagger$ and $\tau^\dagger$ and $\beta^\dagger$ are the probability limits of $\widehat{\gamma}$, $\widehat{\tau}$ and $\widehat{\beta} \cdot \left(b, \widehat{d}_{\cdot, opt}^b \right)$ respectively, $J. (b)$ is defined as in Section 5.2 except that the derivative is evaluated at $\beta^\dagger$ instead of β^* ,

$$I_\gamma(\gamma) \equiv E \left[\frac{\partial}{\partial \gamma'} S_\gamma(\gamma) \right]; \quad I_\tau(\tau) \equiv E \left[\frac{\partial}{\partial \tau'} S_\tau(\tau) \right]$$

and

$$J_{\cdot, \gamma}(\beta, \gamma, \tau) \equiv E \left[\frac{\partial}{\partial \gamma} S_{\cdot, \beta}(x; \beta, \gamma, \tau) \right]; \quad J_{\cdot, \tau}(\beta, \gamma, \tau) \equiv E \left[\frac{\partial}{\partial \tau} S_{\cdot, \beta}(x; \beta, \gamma, \tau) \right].$$

The preceding expansion implies that under regularity conditions and regardless of whether model (16) or model (29) hold,

$$\sqrt{n} \left\{ \widehat{\beta} \cdot \left(b, \widehat{d}_{\cdot, opt}^b \right) - \beta^\dagger \right\} \rightarrow N(0, \Gamma. (b; \beta^\dagger, \gamma^\dagger, \tau^\dagger))$$

where

$$\Gamma. (b; \beta^\dagger, \gamma^\dagger, \tau^\dagger) = J. (b)^{-1'} E \left\{ \int_{\mathcal{X}_{pos}} b. (x, Z) Q. \{x; \beta^\dagger, \gamma^\dagger, \tau^\dagger\} dP_X(x) \right\}^{\otimes 2} J. (b)^{-1}$$

If model (16) holds, then $\beta^\dagger = \beta^*$, $\gamma^\dagger = \gamma^*$ and $\Gamma. (b; \beta^\dagger, \gamma^\dagger, \tau^\dagger)$ coincides with $\Omega. (b, d_{\cdot, lim}^b)$ the asymptotic variance derived in Section 5.3.1. If in addition, model (29) is also correctly specified, then $\Gamma. (b; \beta^\dagger, \gamma^\dagger, \tau^\dagger)$ coincides with $\Omega. (b, d_{\cdot, opt}^b)$, the asymptotic variance of the optimal AIPTW estimator for a given choice of b .

A consistent estimator of $\Gamma(b; \beta^\dagger, \gamma^\dagger, \tau^\dagger)$ is given by

$$\hat{\Gamma} \cdot (b) = \hat{J} \cdot (b)^{-1} \mathbb{P}_n \left\{ \int_{\mathcal{X}_{pos}} b \cdot (x, Z) \hat{Q} \cdot \left\{ x; \hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b), \hat{\gamma}, \hat{\tau} \right\} dP_X(x) \right\}^{\otimes 2} \hat{J} \cdot (b)^{-1}$$

where $\hat{Q} \cdot (x; \beta, \gamma, \tau)$ is defined like $Q \cdot (x; \beta, \gamma, \tau)$ but with the population mean E replaced by the empirical mean $\mathbb{P}_n$ and $\hat{J} \cdot (b)$ is defined as in Section 5.2 but evaluated at $\hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b)$ instead of $\hat{\beta} \cdot (b, d)$.

The variance estimator just described can be used to construct Wald confidence ellipsoids centered at $\hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b)$ with asymptotic coverage probability equal to the nominal level if either model (16) or model (29) is correctly specified, but not necessarily both. Specifically, consider the ellipsoid

$$C_b = \left\{ \beta : n \left(\hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b) - \beta \right)' \hat{\Gamma} \cdot (b)^{-1} \left(\hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b) - \beta \right) \leq \chi_{p, 1-\alpha}^2 \right\}$$

where $\chi_{p, 1-\alpha}^2$ is the $(1 - \alpha) \times 100$ percentile of the χ_p^2 distribution. When either the assumptions of Lemma 3, part (ii) or the assumptions of part (iii) of the Lemma hold, $\beta^\dagger$ is equal to the true value β^* . Then, standard arguments for the properties of Wald confidence sets imply that C_b covers the true β^* with probability $(1 - \alpha)$ as n goes to infinity if either model (16) or model (29) is correctly specified, but not necessarily both.

6.2 Double-Robust Inference for $x_{opt}(z)$

Construction of asymptotically valid Wald confidence regions for $x_{opt}(z)$ that are valid when either model (16) or model (29) are correctly specified is straightforward when $\mathcal{X}$ is an open subset of $\mathbb{R}^s$ and certain regularity conditions hold. Specifically, if $\mathcal{X}$ is an open subset of $\mathbb{R}^s$, and if for each z , $x_{opt}(z; \beta)$ belongs to $\mathcal{X}$ and is differentiable with respect to β in a neighborhood of β^* with non-singular derivative at β^* , then if the assumptions of Lemma 3 part (ii) or the assumptions of part (iii) of that Lemma hold, an application of the delta method gives that

$$\sqrt{n} \left\{ x_{opt} \left(z; \hat{\beta} \cdot (b, \hat{d}_{\cdot, opt}^b) \right) - x_{opt}(z) \right\} \rightarrow N(0, \Sigma \cdot (b; \beta^*, \gamma^\dagger, \tau^\dagger))$$

where

$$\Sigma \cdot (b; \beta^*, \gamma^\dagger, \tau^\dagger) = \left\{ \frac{\partial x_{opt}(z; \beta)}{\partial \beta'} \Big|_{\beta=\beta^*} \right\} \Gamma \cdot (b; \beta^*, \gamma^\dagger, \tau^\dagger) \left\{ \frac{\partial x_{opt}(z; \beta)}{\partial \beta} \Big|_{\beta=\beta^*} \right\}.$$

A consistent estimator of $\Sigma. (b; \beta^*, \gamma^\dagger, \tau^\dagger)$ is obtained by evaluating the derivative of $x_{opt}(z; \beta)$ at $\hat{\beta}. (b, \hat{d}_{opt}^b)$ instead of β^* and replacing $\Gamma. (b; \beta^*, \gamma^\dagger, \tau^\dagger)$ with its consistent estimator $\hat{\Gamma}. (b)$. This estimator can be used to construct a Wald confidence region for $x_{opt}(z)$ centered at $x_{opt}\left(z; \hat{\beta}. (b, \hat{d}_{opt}^b)\right)$ which has coverage probability equal to the nominal one as n goes to ∞ when either model (16) or model (29) hold. Bembom and van der Laan (Appendix C, 2008) also discuss the construction of Wald confidence regions for $x_{opt}(z; \beta)$ when $\mathcal{X}$ is an open subset except that in their application the treatment probabilities are known by design, and as such their intervals are centered at non-double robust estimators and have length computed based on non-double robust estimates of standard error.

The following lemma gives the set of sufficient conditions for the existence of $\partial x_{opt}(z; \beta) / \partial \beta_j$ and an expression for it.

Lemma 4. If $\mathcal{X}$ is an open subset of $\mathbb{R}^s$ and if for each z, i) $x_{opt}(z; \beta)$ is a continuous function of β on an open neighborhood of β^* , ii) $h. (z, x; \beta)$ is differentiable with respect to (x, β) on an open neighborhood of $(x_{opt}(z; \beta^*), \beta^*)$ and iii) $\frac{\partial^2 h. (z, x; \beta^*)}{\partial x \partial x'} \Big|_{x=x_{opt}(z, \beta^*)}$ is non-singular, then for each $z, x_{opt}(z; \beta)$ is differentiable with respect to β in a neighborhood of β^* and

$$\frac{\partial x_{opt}(z; \beta)}{\partial \beta_j} \Big|_{\beta=\beta^*} = \left[\frac{\partial^2 h. (z, x; \beta^*)}{\partial x \partial x'} \Big|_{x=x_{opt}(z, \beta^*)} \right]^{-1} \left[\frac{\partial h. (z, x_{opt}(z, \beta^*); \beta)}{\partial \beta_j} \Big|_{\beta=\beta^*} \right].$$

When $\mathcal{X}$ is countable, whether finite or infinite, the preceding construction is not feasible because the derivative of $x_{opt}(z; \beta)$ is undefined. In this setting, we can consider the set

$$B_b = \{x_{opt}(z; \beta) : \beta \in C_b\}$$

where $x_{opt}(z; \beta)$ is defined in (20) and C_b is the Wald confidence ellipsoid for β of the preceding subsection. This set has coverage probability at least $(1 - \alpha)$ as n goes to ∞ when either model (16) or model (29) hold. In Section 3 of the companion paper (Orellana, Rotnitzky and Robins, 2010) we show that in the special, but widely applicable, case in which $h. (z, x; \beta)$ is a linear function of the β_j 's, determining whether or not a given x belongs to B_b entails determining if the intersection of $\#(\mathcal{X}) - 1$ half-hyperspaces of $\mathbb{R}^p$ and a ball centered at the origin is non-empty, where $\#(\mathcal{X})$ is the cardinality of the set $\mathcal{X}$. Although efficient algorithms to implement this task may be known in computer science or allied fields, we are not aware that such algorithms exist. In Section 3 of

the companion paper (Orellana, Rotnitzky and Robins, 2010) we provide an algorithm to compute a set B_b^* that includes, but is not necessarily included in, B_b in the aforementioned special case. Such set B_b^* is therefore a more conservative confidence set for $x_{opt}(z)$ than B_b . An alternative procedure would be to compute a confidence band for the function that maps x with $m(z, g_x)$ with simultaneous coverage no smaller than $(1 - \alpha)$ and, then to construct the set C comprised by all the x 's for which the confidence interval for $m(z, g_x)$ overlaps with the confidence interval for $m\left(z, g_{x_{opt}(z; \hat{\beta}_\cdot(b, d_{\cdot, opt}^b))}\right)$ both based on the band. Such set C would cover the true $x_{opt}(z)$ with probability at least $(1 - \alpha)$. Bembom and van der Laan (Appendix B, 2008) describe the construction of a simultaneous confidence band for the map $x \mapsto m(z, g_x)$ when $\mathcal{X}$ is finite, by Monte Carlo simulation.

7 Optimal Function b

From expansion (32) we can derive the optimal choice $b_{\cdot, opt}$ for the function b in the sense of minimizing the asymptotic variance of $\hat{\beta}_\cdot(b, d_{\cdot, opt}^b)$, i.e. satisfying

$$\Gamma_\cdot(b; \beta^\dagger, \gamma^\dagger, \tau^\dagger) \geq \Gamma_\cdot(b_{\cdot, opt}; \beta^\dagger, \gamma^\dagger, \tau^\dagger) \quad \text{for all } b \quad (34)$$

Specifically, in Section 2.5 of the companion paper we show that $b_{\cdot, opt}$ is the solution to the Type I Fredholm integral equations

$$-\frac{\partial}{\partial \beta} h_\cdot(x, z; \beta) \Big|_{\beta=\beta^\dagger} = \int_{\mathcal{X}_{pos}} C_\cdot(x, \tilde{x}; z) b_{\cdot, opt}(\tilde{x}, z) dP_X(\tilde{x}), \quad (35)$$

where

$$C_\cdot(x, \tilde{x}; z) \equiv E\{Q_\cdot^\dagger(x)^2 | Z = z\}$$

with

$$Q_{par}^\dagger(x) \equiv Q_{par}(x; \beta^\dagger, \gamma^\dagger, \tau^\dagger)$$

and

$$Q_{sem}^\dagger(x) \equiv Q_{sem}(x; \beta^\dagger, \gamma^\dagger, \tau^\dagger) - \int_{\mathcal{X}_{pos}} Q_{sem}(x; \beta^\dagger, \gamma^\dagger, \tau^\dagger) dP_X(x)$$

where $Q_\cdot(x; \beta, \gamma, \tau)$ is defined in (33).

A point of note is that the inequality (34) implies that $\hat{\beta}_\cdot(b_{\cdot, opt}, d_{\cdot, opt}^{b_{\cdot, opt}})$ has the smallest asymptotic variance among the asymptotic variances of members of the class

$$\left\{ \hat{\beta}_\cdot(b, d_{\cdot, opt}^b) : b \text{ arbitrary} \right\} \quad (36)$$

when either model (16) or (29) holds.

Another point of note is that although the optimal choice $b_{,opt}$ depends on the choice of P_X , the value taken by $\widehat{\beta}_{\cdot}(b, d_{,opt}^b)$ does not because if $b_{,opt}$ solves (35) for a given P_X , then $b'_{,opt} = b_{,opt} dP_X / dP'_X$ solves the integral equation for P'_X , provided P'_X is absolutely continuous with respect to P_X .

When $\mathcal{X}_{pos} = \{x_j : 1 \leq j \leq J\}$ is finite and P_X is the uniform distribution on $\mathcal{X}_{pos}$, the solution of equation (35) can be written in closed form. Specifically, $b_{,opt}(x_j, z)'$ is equal to the j^{th} row of the $J \times p$ matrix

$$\mathbf{b}_{\cdot}(z) \equiv C_{\cdot}(z)^{-1} \mathbf{m}_{\cdot}(z)$$

where $\mathbf{m}_{\cdot}(z)$ is a $J \times p$ matrix with j^{th} row equal to $\partial h_{\cdot}(x_j, z; \beta) / \partial \beta'$ and $C_{\cdot}(z)$ is the $J \times J$ matrix with $(j, l)^{th}$ entry equal to $C_{\cdot}(x_j, x_l; z)$. Of course, $\mathbf{b}_{\cdot}(z)$ is not available because $C_{\cdot}(z)$ is unknown. We may nevertheless consider estimators $\widehat{\beta}_{\cdot}(\widehat{b}, \widehat{d}_{,opt}^b)$ where $\widehat{\mathbf{b}}_{\cdot}(z) = \widehat{C}_{\cdot}(z)^{-1} \mathbf{m}_{\cdot}(z)$ for some estimator $\widehat{C}_{\cdot}(z)$ of $C_{\cdot}(z)$. To estimate $C_{\cdot}(z)$ we may write $C_{\cdot}(z) = D_{\cdot}(z)^{1/2} R_{\cdot}(z) D_{\cdot}(z)^{1/2}$ where $D_{\cdot}(z)$ is the diagonal matrix with diagonal equal to that of $C_{\cdot}(z)$ and postulate working variance and correlation models,

$$D_{\cdot jj}(z) = D_{jj}(z; \eta) \text{ and } R_{\cdot jk}(z) = R_{jk}(z; \delta), \quad 1 \leq j < k \leq J \quad (37)$$

where for each z , $D_{\cdot jj}(z; \cdot)$ and $R_{\cdot jk}(z; \cdot)$ are unknown smooth functions and η and δ are unknown parameters. For example, we may consider a linear exponential model

$$D_{\cdot jj}(z; \eta) = \exp(\eta_{1,j} + \eta_{2,j} z)$$

and a first order autoregressive model

$$R_{\cdot jk}(z; \delta) = \exp(-|j - k| \delta) \text{ for } \delta > 0.$$

We may then estimate η and δ as follows. We compute $\widehat{e}_{\cdot jl} = \widehat{Q}^{\dagger}(x_j) \widehat{Q}^{\dagger}(x_l)$ where $\widehat{Q}^{\dagger}(x)$ is defined like $Q^{\dagger}(x)$ but with $\widehat{\beta}_{\cdot}(b, d)$ (a preliminary estimate of β for arbitrary choices of b and d), $\widehat{\gamma}$, and $\widehat{\tau}$ instead of β^* , $\gamma^{\dagger}$, $\tau^{\dagger}$ and then estimate η with $\widehat{\eta}$ solving

$$\mathbb{P}_n \left[\sum_{j=1}^J \left\{ \widehat{e}_{\cdot jj}^2 - D_{\cdot jj}(z; \eta) \right\} \frac{\partial D_{\cdot jj}(z; \eta)}{\partial \eta} \right] = 0$$

and estimate δ with $\widehat{\delta}$ solving

$$\mathbb{P}_n \left[\sum_{\substack{l,j=1 \\ j \neq l}}^J \left\{ \frac{\widehat{e}_{\cdot jl}}{\sqrt{D_{\cdot jj}(z; \widehat{\eta}) D_{\cdot ll}(z; \widehat{\eta})}} - R_{\cdot jk}(z; \delta) \right\} \frac{\partial R_{\cdot jk}(z; \delta)}{\partial \delta} \right] = 0.$$

Finally, we compute the estimator $\widehat{C} \cdot (z) = D \cdot (z; \widehat{\eta})^{1/2} R \cdot (z; \widehat{\delta}) D \cdot (z; \widehat{\eta})^{1/2}$.

If for each z , $\widehat{C} \cdot (z)$ converges in probability to some matrix $C^* \cdot (z)$, then $\widehat{\mathbf{b}} \cdot (z)$ converges in probability to $b^* \cdot (z) = C^* \cdot (z)^{-1} m \cdot (z)$. If $b^* \cdot (x, z)$ stands for the j^{th} row of the $J \times p$ matrix $b^* \cdot (z)$, then results in Newey and McFadden (1994) imply that $\sqrt{n} \left\{ \widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right) - \beta^\dagger \right\}$ and $\sqrt{n} \left\{ \widehat{\beta} \cdot \left(b^* \cdot, \widehat{d}^{b \cdot}_{opt} \right) - \beta^\dagger \right\}$ have the same asymptotic distribution.

The estimators $\widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right)$ are locally efficient in two different models and for two different classes of estimators. Specifically,

1. In the special case in which at least one of model (16) or model (29) is correctly specified, all estimators in the class

$$\left\{ \widehat{\beta} \cdot \left(b, \widehat{d}^{b \cdot}_{opt} \right) : b \text{ arbitrary} \right\} \quad (38)$$

converge in probability to β^* , so in particular $\widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right)$ converges to β^* if either of these models is correct. Because when model (37) holds, the estimator $\widehat{C} \cdot (z)$ is consistent for $C \cdot (z)$, it follows that when model (37) holds, $\widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right)$ has asymptotic variance equal to the smallest asymptotic variance of estimators in the class (38). We conclude that $\widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right)$ is double-robust and locally efficient in the union model $\text{DYR} \cup \text{MSM-obs}^* \cup \text{DYR} \cup \text{MSM-obs}^\dagger$ (defined in Section 6) for the class (38) of double-robust estimators at the submodel of the union model that additionally imposes the restrictions (37).

2. We have already argued in Section 6.1 that when models (29) and (16) are correctly specified, $\Gamma \cdot (b; \beta^*, \gamma^*, \tau^*)$ coincides with $\Omega \cdot (b; d^{b \cdot}_{opt})$. Consequently, when all three models (37), (29) and (16) are correct, $\widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right)$ has asymptotic variance $\Omega \cdot \left(b_{opt}; d^{b \cdot}_{opt} \right)$ satisfying (24). We thus conclude that $\widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right)$ is a consistent and asymptotically normal estimator of β^* under model $\text{DYR} \cup \text{MSM}^{**}$ with asymptotic variance that coincides with the smallest possible asymptotic variance of AIPTW estimators in the class

$$\mathcal{T} = \left\{ \widehat{\beta} \cdot (b, d) : b \text{ and } d_k, k = 0, \dots, K, \text{ arbitrary} \right\} \quad (39)$$

if, in addition, models (29) and (37) hold. Thus, $\widehat{\beta} \cdot \left(\widehat{b} \cdot, \widehat{d}^{b \cdot}_{opt} \right)$ is a locally efficient estimator of β^* in model $\text{DYR} \cup \text{MSM}^*$ for the class $\mathcal{T}$ of AIPTW

estimators at the submodel of model DYR--MSM* which imposes the additional restrictions (29) and (37).

In practice, the set $\mathcal{X}_{pos}$, even if finite, may be very large. For instance, in Example 1 of Section 4, $\mathcal{X}_{pos}$ has 401 elements. In such case the square matrix $\widehat{C} \cdot (z)$ will be unduly large dimension and its inversion will generally be computationally intractable. For such scenarios, we may consider the following dimension reduction strategy. Let $\mathcal{X}_{pos}^*$ be a subset of $\mathcal{X}_{pos}$ with small cardinality J^* . Let the $J^* \times p$ matrix $\widehat{\mathbf{b}} \cdot (z)$ be defined as in the previous paragraph but with $\mathcal{X}_{pos}^*$ instead of $\mathcal{X}_{pos}$. Define for each $x \in \mathcal{X}_{pos}^*$, $b \cdot (x, z)'$ equal to the corresponding row of $\widehat{\mathbf{b}} \cdot (z)$. Finally, for any x' in $\mathcal{X}_{pos} - \mathcal{X}_{pos}^*$ set $b \cdot (x', z)$ equal to one of the following possible choices: *i*) 0, *ii*) $b \cdot (x, z)$ for x the closest neighbor of x' in the Euclidean distance sense, or *iii*) the linear interpolation between $b \cdot (x_1, z)$ and $b \cdot (x_2, z)$ for x_1 and x_2 the two closest neighbors of x' , if $\mathcal{X}_{pos} \subseteq \mathbb{R}$. We hypothesize that when the function $h \cdot (z, x; \beta^*)$ is a smooth function of x , the loss in efficiency due to implementing this strategy will be small if the distribution of the points in $\mathcal{X}_{pos}^*$ is similar to that of the points in $\mathcal{X}_{pos}$, for example if $\mathcal{X}_{pos}^*$ is comprised of one in every, say, ten (sorted) points in $\mathcal{X}_{pos}$ when $\mathcal{X}_{pos} \subseteq \mathbb{R}$. A thorough examination of the properties of this strategy is beyond the scope of this paper.

When $\mathcal{X}_{pos}$ is an infinite but bounded set, we recommend conducting the strategy just described on a set $\mathcal{X}_{pos}^*$ comprised by the points of a grid of $\mathcal{X}_{pos}$ of not too large cardinality. Once again, we hypothesize this strategy will not result in substantial efficiency losses when $h \cdot (z, x; \beta^*)$ is a smooth function of x .

8 Discussion

The dynamic regime marginal structural mean models developed in this paper have two appealing properties: they are easy to understand and they are easy to fit with standard off-the-shelf software that allows for weights. Nevertheless, estimation of the parameters of dynamic regime MSMs has some limitations. First, if the number of time periods is large, the denominator of $\omega_K^x(\gamma)$ is the product of many $\lambda_k^x(A_k | \overline{O}_k, \overline{A}_{k-1}; \gamma)$. Thus, even if the true $\omega_K^x(\gamma^*)$'s are not highly variable, small fluctuations in $\widehat{\gamma}$ may propagate quickly and result in large fluctuations of the $\omega_K^x(\widehat{\gamma})$. This may result in very large weights for a few compliers with regime x compared to the rest of the compliers in the sample. Even in moderate-sized studies, for estimation of β^* the consequence of this phenomenon in finite samples is bias and variance that is substantially larger

than the one prescribed by the asymptotic theory. Second, dynamic regime MSMs do not allow for direct modelling of interactions between treatment and evolving time dependent covariates and so do not permit qualitative effect modification to be considered directly. Finally, it is unclear how to construct extensions of dynamic regime MSMs suitable for assessing the sensitivity to reasonable departures from the assumption of sequential randomization. In contrast, the optimal regime structural nested mean (SNM) model of Robins (2004) does not suffer from any of these difficulties. However, as indicated in the introduction, SNM's cannot be used to estimate the optimal regime in a prespecified parametrized class of regimes (such as the optimal CD4 cell count at which to start HAART) that includes just the set of logistically feasible regimes.

In this paper we have assumed that patients came to the clinic at the end of each of K prespecified intervals. Admittedly, this assumption is unrealistic in most applications. In most observational databases there will be variability across patients in the frequency of, and reasons for, physician visits. This raises many important questions regarding the characteristics of the population in which the optimal treatment regime that we are aiming to estimate will be implemented. Clearly, the optimal treatment strategy depends on the timing of the clinic visits. Methods aiming to estimate the optimal treatment regime will depend on whether *i*) the subjects comprising the target population in which the regime is to be implemented are biologically similar to those in the population from which the database was obtained, and *ii*) the health care systems of the target and study populations differ for cultural, logistical, and financial reasons in the frequency of, and reasons for physician visits. Orellana and Rotnitzky (2007) and Robins, Orellana and Rotnitzky (2008) consider methodology to address these issues.

References

- Cox, D.R. (1958). *The design and planning of experiments*. Chapman and Hall, London.
- Bembom, O. and van der Laan, M. (2008). Analyzing sequentially randomized trials based on causal effect models for realistic individualized treatment rules. *Statistics in Medicine*. 27, 3689–3716
- Gill, R.D, van der Laan, M.J. and Robins J.M. (1997). Coarsening at random: characterizations, conjectures and counterexamples. *Proceedings of the First Seattle Symposium on Survival Analysis*, pp. 255-294.

- Hernán, M.A., Brumback B. and Robins J.M. (2000). Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*. 11, 561-570.
- Luenberger, D.G. (1969). *Optimization by vector spaces*. Wiley, New York.
- Murphy, S.A. (2003). Optimal dynamic treatment regimes, *Journal of the Royal Statistical Society, Series B* (with discussion) 65(2), 331-366.
- Murphy, S.A., van der Laan, M.J., Robins J.M. and CPPRG (2001). Marginal mean models for dynamic regimes. *Journal of the American Statistical Association*, 96, 1410-1423.
- Neyman, J. (1923). On the application of probability theory to agricultural experiments. Translated in *Statistical Science*, bf 5, 465-480 (1990).
- Newey, W. and Mc Fadden D. (1994). Large sample estimation and hypothesis testing, in R. Engle and D. McFadden (eds.) *Handbook of Econometrics*, Vol. 4, Amsterdam, North Holland, 2113-2245.
- Orellana, L. and Rotnitzky, A. (2007). Estimation of the effect of dynamic treatment regimes under flexible dynamic visit regimes. Liliana Orellana's Harvard. University, Dep of Biostatistics, Ph.D. Thesis, Chapter 2.
- Orellana, L., Rotnitzky, A. and Robins J.M. (2010). Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes. Part II: Proofs of results. *The International Journal of Biostatistics*, Vol. 6: Iss. 2, Article 8.
- Robins, J.M. (1986). A new approach to causal inference in mortality studies with sustained exposure periods -application to control of the healthy worker survivor effect. *Computers and Mathematics with Applications* 14, 1393-1512.
- Robins, J.M. (1987). Addendum to "A new approach to causal inference in mortality studies with sustained exposure periods - Application to control of the healthy worker survivor effect." *Computers and Mathematics with Applications* 14, 923-945.
- Robins, J. M. (1989). The analysis of randomized and nonrandomized AIDS treatment trials using a new approach to causal inference in longitudinal studies. In *Health Service Research Methodology: a Focus on AIDS* (eds L. Sechrest, H. Freeman and A. Mulley), pp. 113-159. US Public Health Service.
- Robins, J.M. (1993). Analytic methods for estimating HIV treatment and cofactor effects. In: *Methodological issues of AIDS Mental Health Research*.

- Eds.: D. G. Ostrow & R. Kessler. Plenum Publishing, New York, pp. 213–290.
- Robins, J.M. (1997). Causal inference from complex longitudinal data. Latent Variable Modeling and Applications to Causality. *Lecture Notes in Statistics* (120), 69-117, (eds: M. Berkane). Springer-Verlag, Inc, New York.
- Robins, JM. (1998). Marginal structural models. *Proceedings of the American Statistical Association. Section on Bayesian Statistics*, pp. 1-10.
- Robins, J.M. (1999). Marginal structural models versus structural nested models as tools for causal inference. *Statistical Models in Epidemiology: The Environment and Clinical Trials*. M.E. Halloran and D. Berry, Editors, IMA Volume 116, NY: Springer-Verlag, pp. 95-134.
- Robins, J.M. (2004). Optimal structural nested models for optimal sequential decisions. In D. Y. Lin and P. Heagerty (Eds.), *Proceedings of the Second Seattle Symposium on Biostatistics*, New York. Springer.
- Robins, J.M. and Rotnitzky, A. (1992). Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS Epidemiology, Methodological issues*. Birkhäuser, editors.
- Robins, J.M., Rotnitzky, A and Zhao L.P. (1994). Estimation of regression-coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427): 846-866.
- Robins, J.M., Rotnitzky A and van der Laan M. (2000). Comment on "On Profile Likelihood" by S.A. Murphy and A.W.van der Vaart. *Journal of the American Statistical Association – Theory and Methods*, 95(450):431-435.
- Robins, J., Orellana, L. and Rotnitzky, A. (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine*. 27(23):4678-721
- Rubin, D.B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics*, 6, 34-58.
- Tsiatis, A. (2006). *Semiparametric Theory and Missing Data*. Springer.
- van der Laan, M.J. (2006). Causal effects models for intention to treat and realistic individualized treatment rules. <http://www.bepress.com/ucbbiostat/paper203>
- van der Laan, M.J. and Petersen, M.L. (2007). Causal effect models for realistic individualized treatment and intention to treat rules. *The International Journal of Biostatistics*. 3(1), Article 3.

van der Laan, M.J. and Robins J.M. (2003). *Unified Methods for Censored Longitudinal Data and Causality*. Springer Verlag: New York.