

Analysis of Longitudinal Marginal Structural Models

Jennifer F. Bryan* Zhuo Yu[†]

Mark J. van der Laan[‡]

*Dept. of Statistics and Biotechnology Lab., University of British Columbia

[†]Department of Statistics, University of California, Berkeley

[‡]Division of Biostatistics, School of Public Health, University of California, Berkeley

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

<http://biostats.bepress.com/ucbbiostat/paper120>

Copyright ©2002 by the authors.

Analysis of Longitudinal Marginal Structural Models

Jennifer F. Bryan, Zhuo Yu, and Mark J. van der Laan

Abstract

In this article we construct and study estimators of the causal effect of a time-dependent treatment on survival in longitudinal studies. We employ a particular marginal structural model (MSM), and follow a general methodology for constructing estimating functions in censored data models. The inverse probability of treatment weighted (IPTW) estimator is used as an initial estimator and the corresponding treatment-orthogonalized, one-step estimator is consistent and asymptotically linear when the treatment mechanism is consistently estimated. We extend these methods to handle informative censoring. A simulation study demonstrates that the the treatment-orthogonalized, one-step estimator is superior to the IPTW estimator in terms of efficiency. The proposed methodology is employed to estimate the causal effect of exercise on mortality in a longitudinal study of seniors in Sonoma County.

1 Introduction

The methods developed in this article make it possible to estimate the causal effect of a time-varying treatment on an outcome process, even in the presence time-dependent confounders and censoring. We rely on the abstract concept of counterfactual outcomes [Rubin, 1976] to define this causal effect, which we present briefly here and more formally in an explicitly longitudinal setting in section 1.1.

Let the set of possible treatments be denoted by \mathcal{A} . When subjected to the same treatment $a \in \mathcal{A}$, the outcome of interest for a population of exchangeable subjects, denoted Y_a , has some particular marginal distribution. Such an outcome Y_a is a 'counterfactual' in the sense that, for any given subject, we assume its existence, even if the subject actually receives a different treatment. Differences in the marginal distributions of counterfactual outcomes, e.g. F_{Y_a} versus $F_{Y_{a'}}$ for $a, a' \in \mathcal{A}$, are precisely what we mean by the causal effect of one treatment versus another. For example, if such a difference implies that treatment a induces a median survival time of 15 months, whereas treatment a' induces a median survival time of 21 months, this describes a positive causal effect of treatment a' on median survival, relative to treatment a . We have just described the use of *marginal structural models* [Robins, 2000] for the purposes of causal inference. That is, the causal effect of treatment is defined to be the structural differences in marginal counterfactual distributions induced by different treatments.

In real life, of course, each subject can only receive one treatment. On a randomly drawn subject, we observe the treatment $A \in \mathcal{A}$, the outcome $Y = Y_A$, and possibly baseline covariates W . Therefore, we regard the observed data as an incomplete observation of a subject's 'full data' ($Y_a : a \in \mathcal{A}, W$), i.e. the data we would observe under every possible treatment. Consequently, we approach the causal inference problem as if it were a censored or missing data problem. The causal effect of treatment $A = a$ is defined as the effect of a on the distribution of Y_a , possibly adjusted for selected baseline covariates $V \subset W$. One says that the treatment A is randomized w.r.t W if the probability of $A = a$ given the full data ($Y_a : a \in \mathcal{A}, W$), the so-called propensity score, only depends on W . In other words, the treatment assignment is only based on the observed covariates W . This randomization assumption is necessary to be able to identify the causal effect of A on the outcome Y . Under this assumption, we can exploit the theory of estimating functions in censored data models in order to estimate the marginal counterfactual distributions and, thereby, the causal effect of treatment. A locally optimal doubly robust method for estimation of the causal effect of a time-independent treatment in a marginal structural semiparametric model is given in Robins [1999]. Rosenbaum and Rubin [1983, 1985] and Rosenbaum [1987, 1988, 1995, 1996] propose to estimate causal effects by stratification on an estimate of the propensity score. A general presentation of the use of censored data methods to perform causal inference is given in van der Laan and Robins [To appear in 2002].

1.1 Longitudinal marginal structural models

The theory of counterfactual causal inference in longitudinal studies is first laid out in Robins [1986, 1987, 1989, 1997]. Let time t take values in $\tau = [0, T]$, where T is possibly infinite. For a time-dependent process $t \rightarrow Z(t)$, we denote its sample path up to time t^* with $\bar{Z}(t^*) = \{Z(t) : t \leq t^*\}$ and its complete sample path by $\bar{Z} = \bar{Z}(T) = \{Z(t) : t \in \tau\}$. In this paper, we will work with finite T and divide $\tau = [0, T]$ into K intervals of equal length. Therefore, time is discrete and takes values in $\tau = \{t_0, t_1, \dots, t_{K-1}, t_K = T\}$. For any time-dependent process, we may use the abbreviated notation Z_k and \bar{Z}_k in place of $Z(t_k)$ and $\bar{Z}(t_k)$.

Let $t \rightarrow A(t)$ be a time-dependent treatment process and let \mathcal{A} be the set of possible sample paths of \bar{A} , where we assume that \mathcal{A} is finite. Let $Y_{\bar{a}}(t_k)$ be the counterfactual outcome process under treatment \bar{a} and let $L_{\bar{a}}(t_k)$ be the corresponding covariate process. For each possible treatment regime \bar{a} , we define $\bar{X}_{\bar{a}}(t_k)$ as the data one would observe on the subject up to time t_k , if the subject were to follow treatment regime \bar{a} . Note that this only depends on the treatment history *prior* to time t_k , i.e. $\bar{X}_{\bar{a}}(t_k) = \bar{X}_{\bar{a}_{k-1}}(t_k)$. The complete sample path $\bar{X}_{\bar{a}}$ is a counterfactual and is comprised of the paths of the outcome process $\bar{Y}_{\bar{a}}$, the covariate process $\bar{L}_{\bar{a}}$, and baseline covariates W . The full data for a subject is the collection of counterfactuals generated by allowing treatment to range over the entire space \mathcal{A} , e.g.:

$$X = (\bar{Y}_{\bar{a}}, \bar{L}_{\bar{a}}, W : \bar{a} \in \mathcal{A})$$

The observed data is given by

$$O = (\bar{A}, \bar{X}_{\bar{A}}) = (\bar{A}, \bar{Y}_{\bar{A}}, \bar{L}_{\bar{A}}, W).$$

We see that the observed data are only those elements of the full data corresponding to the actual observed treatment \bar{A} .

Consider a study designed to determine the causal effect of a time-dependent treatment on survival or, alternatively, the hazard of mortality. Let $S_{\bar{a}}$ be a counterfactual survival time and let $Y_{\bar{a}}(t_k) = I(S_{\bar{a}} \leq t_k)$ be a counting process that indicates for death prior to time t_k , where $Y_{\bar{a}}(t_0)$ is 0 for all $\bar{a} \in \mathcal{A}$ by definition. The survival time can be recovered from the path $\bar{Y}_{\bar{a}}$, up to the resolution permitted by discrete time. A subject's record continues until failure. Recall that $V \subset W$ is a selected subset of the baseline covariates for which want to adjust. We can assume a marginal structural intensity model for the process $Y_{\bar{a}}(\cdot)$:

$$E(dY_{\bar{a}}(t_k) | \bar{Y}_{\bar{a}}(t_{k-1}), V) = \lambda(t_k, \bar{a}_{k-1}, \bar{Y}_{\bar{a}}(t_{k-1}), V | \alpha),$$

where $\lambda(t_k, \bar{a}_{k-1}, \bar{Y}_{\bar{a}}(t_{k-1}), V | \alpha)$ is a known function, up to a p -dimensional parameter α . Typically, $\lambda(\cdot | \alpha)$ is the product of an indicator that the subject is at risk for the event of interest and a logistic function:

$$\lambda(t_k, \bar{a}_{k-1}, \bar{Y}_{\bar{a}}(t_{k-1}), V | \alpha) = I(Y_{\bar{a}}(t_{k-1}) = 0) \times \pi(t_k, \bar{a}_{k-1}, V | \alpha), \quad (1)$$

where $\pi(t_k, \bar{a}_{k-1}, V | \alpha)$ might be something like $\text{logit}^{-1}(\alpha_0 + \alpha_1 t_k + \alpha_2 a_{k-1} + \alpha_3 V)$. Our goal is to estimate α based on n i.i.d. copies (O_1, O_2, \dots, O_n) of O .

All timepoints t_k are not created equal, in terms of the actions taken and the information collected. Continuing the above example, the potential failure times t_k will be called monitoring times. The interval length $t_k - t_{k-1}$ will generally be quite small and corresponds to the resolution with which we record survival time, for example, up to the month of death. At a given subset of the monitoring times, which we call the measuring times, we observe the covariate process L_k . These measuring times generally coincide with a regular assessment such as a medical check-up. Typically, the treatment A_k can change at these measuring times, but one can also imagine situations in which the treatment changes at even fewer time points. We call these treatment times, which are a subset of the measuring times, which are a subset of the monitoring times. Schematically, at a time t_k which is a monitoring, measuring, and treatment time, here is what happens for a subject:

1. Determine the outcome Y_k , i.e. confirm that subject survived the interval (t_{k-1}, t_k) .
2. Measure the covariate L_k .
3. Fix (or at least record) the treatment A_k .

In order to identify causal effects, we must assume that the probability of a particular treatment decision at a treatment time t_k only depends on the observed history $(\bar{A}_{k-1}, Y_k, L_k, W) = (\bar{A}_{k-1}, \bar{X}_{\bar{A}}(t_k))$ of the subject. This assumption is called the sequential randomization assumption (SRA). To formally define the SRA, we recall the full data for a subject: $X = (\bar{X}_{\bar{a}} : \bar{a} \in \mathcal{A})$. The treatment mechanism satisfies SRA if

$$\begin{aligned} g(\bar{A} | X) &= g(A_0 | X) \prod_{k=1}^{K-1} g(A_k | \bar{A}_{k-1}, X) \\ &= g(A_0 | \bar{X}_{\bar{A}}(t_0)) \prod_{k=1}^{K-1} g(A_k | \bar{A}_{k-1}, \bar{X}_{\bar{A}}(t_k)) \end{aligned} \quad (2)$$

In other words, conditional on the observed past, the treatment decision at time t_k is independent of the full set of counterfactual data X [Robins, 1997]. This assumption is also referred to as the assumption of no unmeasured confounders.

It is important to point out why conventional approaches will not produce valid results from this data structure. The most common method for handling confounders is to adjust for them or, in other words, to include all confounders in the regression model. In a point-treatment study, the

resulting regression coefficient for the treatment does indeed have a causal interpretation. However, in a longitudinal study where the treatment changes over time – possibly in response to observed confounders, which are also affected by past treatment – such a regression will produce a biased estimate of the causal treatment effect [Robins et al., 2000], even if the assumption of no unmeasured confounders holds and the regression model is correctly specified. Another reason to look beyond the usual approach is the need to describe treatment effects for large diverse populations, for example, in policy-making. We may be truly interested the **marginal** effect of treatment on a population, as opposed to the treatment effect conditional on the values of certain covariates.

In section 2 we provide two ways to estimate α from the observed data $O = (\bar{A}, \bar{X}_{\bar{A}})$. First, the inverse-probability-of-treatment (IPTW) estimator of Robins et al. [2000] is presented. Second, a more efficient and more robust estimator, the 'treatment-orthogonalized' IPTW estimator (TO), is presented. We implement this as a one-step estimator, using the IPTW estimator as the initial estimator. For both the IPTW and TO estimators, we provide confidence intervals for α . In section 3 we extend the approach to handle informative censoring.

In order to compare the practical performance of the estimators under consideration, we present the results of a simulation study in section 4. We note that the one-step TO estimator can be far more efficient than the IPTW estimator. Finally, in section 5, we apply the extended methodology to estimate the causal effect of exercise on mortality in a longitudinal study of seniors in Sonoma County.

2 Estimation and inference

2.1 IPTW estimator

In this section we describe the first of two estimators of α : the IPTW estimator. This estimator, denoted $\hat{\alpha}_n^{iptw}$, forms the basis of the treatment-orthogonalized estimator described in section 2.2. An IPTW estimator is obtained as the solution of an estimating equation and the relevant estimating function results from the mapping of a full data estimating function into an observed data estimating function. We refer to this mapping as the IPTW mapping from full data estimating functions to observed data estimating functions. For details on this mapping and the fact that the corresponding treatment-orthogonalized IPTW estimating functions comprise the class of all observed data estimating functions, and thus includes, in particular, the optimal estimating function, we refer the reader to van der Laan and Robins [To appear in 2002].

Let $h(\cdot)$ be any function of time, the selected baseline covariates V , and the observed history of the treatment and outcome processes; a typical choice of h would be the score function from the marginal structural model. For every h , we can define an IPTW estimating function $IC_{iptw}(O|g, \alpha, h)$:

$$\begin{aligned} IC_{iptw}(O|g, \alpha, h) &= \sum_k sw(t_k) \times h(t_k, \bar{A}_k, \bar{Y}_k, V) \times \epsilon_{\bar{A}}(\alpha) \\ &= \sum_k \left(\prod_{j=0}^k \frac{g(A_j|\bar{A}_{j-1}, V)}{g(A_j|\bar{A}_{j-1}, \bar{X}_{\bar{A}}(t_j))} \right) \times h(t_k, \bar{A}_k, \bar{Y}_k, V) \times \epsilon_{\bar{A}}(\alpha). \end{aligned} \tag{3}$$

This has the familiar form of an estimating function, namely, a product of a residual and some function of the data. However in this case, we additionally have stabilized weights $sw(t_k)$ that capture the probability of the observed treatment given the past. In practical terms, any IPTW-type estimator works by upweighting (downweighting) subjects that, given their covariate values, have received an unusual (typical) treatment. This is achieved through the use of weights inversely proportional to the probability of the observed treatment, given the covariate. The stabilized weights $sw(t_k)$ in (3) [Robins, 1998] include a numerator term that, in the absence of time-dependent confounding, will equal the denominator and will produce an unweighted estimating function. However, in the presence of confounding, stabilized weights increase the efficiency and robustness of the IPTW estimator $\hat{\alpha}_n^{iptw}$. It is important that the denominator of $sw(t_k)$ be non-negligible for all treatment actions and for all possible histories. In practice, this often implies that the dimension of the treatment space must be modest. Using g_n to denote an estimator of the treatment mechanism, we obtain the following

estimating equation

$$\sum_{i=1}^n IC_{iptw}(O_i|g_n, \hat{\alpha}_n^{iptw}, h) = 0.$$

The solution of this equation, $\hat{\alpha}_n^{iptw}$, is the IPTW estimator.

Let us consider the marginal structural logistic model (1). We will make the usual choice for h , namely:

$$h(t_k, \bar{A}_k, \bar{Y}_k, V) = \frac{d}{d\alpha} \text{logit } \pi(t_k, \bar{A}_k, V|\alpha) \quad (4)$$

Standard software can be employed to solve the weighted estimating equation implied by (4). Practical details for implementing this are provided in section 4.2.

2.2 Treatment-orthogonalized IPTW estimator

In this section we provide a second estimator of α , namely, the treatment-orthogonalized (TO) estimator. This estimator, denoted $\hat{\alpha}_n^{to}$, builds upon the IPTW estimator, but is more efficient. This is the benefit of constructing an estimating function that is orthogonal to the nuisance tangent space T_{SRA} . Recall the treatment mechanism given in equation (2), which satisfies the SRA. Under SRA, the observed data likelihood factorizes into terms arising from the distribution of the full data and terms arising from the treatment mechanism. The nuisance tangent space T_{SRA} is defined by the nuisance scores obtained by varying the treatment mechanism g .

It will be convenient in this and later sections to define $\mathcal{F}_k = (\bar{A}_{k-1}, \bar{X}_{\bar{A}}(t_k))$; in words, \mathcal{F}_k is the observed past just prior to the treatment assignment A_k . The tangent space for the treatment mechanism nuisance parameter at time t_k , denoted $T_{SRA,k}$, is the space of scores obtained by varying $g(A_k|\mathcal{F}_k)$. This is the space of all functions of A_k and \mathcal{F}_k that have conditional mean zero, given the observed past \mathcal{F}_k , i.e. for d_k ranging over all functions of A_k and \mathcal{F}_k we have that

$$T_{SRA,k} = \{d_k(A_k, \mathcal{F}_k) - E(d_k(A_k, \mathcal{F}_k)|\mathcal{F}_k)\}.$$

The factorization of $g(\bar{A}|X)$ into time-specific terms implies that

$$\begin{aligned} T_{SRA} &= T_{SRA,0} \oplus T_{SRA,1} \oplus \dots \oplus T_{SRA,K-1} \\ &= \left\{ \sum_k d_k(A_k, \mathcal{F}_k) - E(d_k(A_k, \mathcal{F}_k)|\mathcal{F}_k) \right\}. \end{aligned}$$

If the time-specific treatment actions A_k take values 0 or 1, it can be shown that the difference $d_k(A_k, \mathcal{F}_k) - E(d_k(A_k, \mathcal{F}_k)|\mathcal{F}_k)$ can be expressed as

$$\begin{aligned} d_k(A_k, \mathcal{F}_k) - E(d_k(A_k, \mathcal{F}_k)|\mathcal{F}_k) &= (d_k(1, \mathcal{F}_k) - d_k(0, \mathcal{F}_k)) (A_k - E(A_k|\mathcal{F}_k)) \\ &= H_k(\mathcal{F}_k) dM_G(t_k), \end{aligned}$$

where $H_k(\mathcal{F}_k) = d_k(1, \mathcal{F}_k) - d_k(0, \mathcal{F}_k)$ and $dM_G(k) = A_k - E(A_k|\mathcal{F}_k)$. Therefore, T_{SRA} can be expressed as

$$T_{SRA} = \left\{ \sum_k H_k(\mathcal{F}_k) dM_G(t_k) \right\}.$$

We obtain a TO estimating function from the IPTW estimating function $IC_{iptw}(O|g, \alpha, h)$ by subtracting its projection onto T_{SRA} . The projection of $IC_{iptw}(O|g, \alpha, h)$ onto T_{SRA} is given by

$$\begin{aligned} IC_{SRA}(O|g, Q) &= IC_{SRA}(O|g, Q_{\alpha, h}) \\ &= \sum_k E(IC_{iptw}(O|g, \alpha, h)|A_k, \mathcal{F}_k) - E(IC_{iptw}(O|g, \alpha, h)|\mathcal{F}_k) \\ &= \sum_k \left[Q(A_k, \mathcal{F}_k) - \sum_{a_k} Q(a_k, \mathcal{F}_k) g(a_k|\mathcal{F}_k) \right], \end{aligned}$$

where $Q(\cdot)$ is the conditional expectation of the IPTW estimating function, given the past A_k and \mathcal{F}_k , and $Q_{\alpha,h}(\cdot)$ includes the causal parameter α and the function h . The TO estimating function can now be written as

$$IC_{to}(O|g, \alpha, h, Q) = IC_{iptw}(O|g, \alpha, h) - IC_{SRA}(O|g, Q).$$

We see the estimating function depends on the causal parameter α , the choice of h and two nuisance parameters: the treatment mechanism g and the conditional expectation Q . As described above, h can be quite general, but a typical choice would be the score function, such as in equation (4). The estimation of g has been discussed earlier and will be demonstrated in later sections. The estimation of Q is generally carried out via regression and is demonstrated later. Assuming that Q_n and g_n are estimators of Q and g , respectively, we define the TO estimator $\hat{\alpha}_n^{to}$ as the solution to the following estimating equation

$$\sum_{i=1}^n IC_{to}(O_i|g_n, \hat{\alpha}_n^{to}, h, Q_n) = 0. \quad (5)$$

If a \sqrt{n} -consistent initial estimator of α is available, there is an attractive alternative to solving equation (5): one can construct a one-step estimator. The one-step estimator is asymptotically equivalent to the actual solution of the estimating equation in the sense that both are asymptotically linear with the same influence curve. A one-step estimator is obtained from a \sqrt{n} -consistent initial estimator and an adjustment that comes from executing the first step of the Newton-Raphson algorithm for solving (5). Concretely, the adjustment is the inverse of the derivative w.r.t. α of the estimating equation (5), evaluated at the initial estimator. We use the IPTW estimator $\hat{\alpha}_n^{iptw}$ of section 2.1 as the initial estimator. The one-step TO estimator, denoted $\hat{\alpha}_n^{to}$, is defined as

$$\hat{\alpha}_n^{to} = \hat{\alpha}_n^{iptw} + \frac{1}{n} \sum_i -c_n^{-1} IC_{to}(O_i|g_n, \hat{\alpha}_n^{iptw}, h, Q_n), \quad (6)$$

where

$$c_n = \frac{1}{n} \sum_i \frac{d}{d\alpha} IC_{to}(O_i|g_n, \alpha, h, Q_n)|_{\alpha=\hat{\alpha}_n^{iptw}}$$

or

$$c_n = \frac{1}{n} \sum_i \frac{d}{d\alpha} IC_{iptw}(O_i|g_n, \alpha, h)|_{\alpha=\hat{\alpha}_n^{iptw}}.$$

In the case where the treatment A_k only takes on values of 0 and 1, the projection can be written as

$$IC_{SRA}(O|g, Q) = \sum_k (Q(1, \mathcal{F}_k) - Q(0, \mathcal{F}_k)) dM_g(k), \quad (7)$$

where $dM_g(k) = A_k - E(A_k|\mathcal{F}_k)$. Therefore, to actually compute the TO estimator, we will have to estimate the expected value of $IC_{iptw}(O|g, \alpha, h)$ given the observed past \mathcal{F}_k and all possible current treatment options: either $A_k = 0$ or $A_k = 1$. This is generally done through regression and, given estimated parameters, we can estimate IC_{SRA} for each subject using (7). In this manner we obtain \widehat{IC}_{to} and can compute the one-step TO estimator given in (6). Practical details are provided later and an Splus example is given in the appendix.

In fact, we can actually use the estimating functions to evaluate the projection fit. The success of the projection of the IPTW estimating function onto the nuisance tangent space determines the performance gain of the TO estimator. By the definition of IC_{SRA} , the following inner product must be zero:

$$\langle IC_{SRA}, IC_{iptw} - IC_{SRA} \rangle = \langle IC_{SRA}, IC_{to} \rangle = 0.$$

This suggests an empirical way to evaluate the quality of the projection. Let \widehat{IC}_{SRA} and \widehat{IC}_{to} be estimates of the respective quantities. We can calculate \widehat{IC}_{SRA} and \widehat{IC}_{to} for each subject; both will

be vectors of the same dimension as α . The empirical correlation of each component of \widehat{IC}_{SRA} with each component of \widehat{IC}_{to} should be close to zero, if the projection was successful. We collect these correlations in a square, nonsymmetric matrix $\widehat{\rho}$ in which the i, j -th element provides the correlation of the i -th component of \widehat{IC}_{SRA} and the j -th component of \widehat{IC}_{to} .

2.3 Asymptotic performance

The TO estimator will remain unbiased even if the conditional expectation Q has been incorrectly specified, as long as the treatment mechanism g has been correctly specified. That is, even if $Q_n \rightarrow Q'$ and $Q' \neq Q$, the TO estimator is consistent when $g_n \rightarrow g$.

Under regularity conditions, assuming that $Q_n \rightarrow Q'$, and $g_n \rightarrow g$, the TO estimator $\widehat{\alpha}_n^{to}$, defined in equation (5), is a regular asymptotically linear estimator with influence curve

$$IC(\cdot) = -c^{-1} [IC_{to}(\cdot|g, \alpha, h, Q') - \Pi(IC_{to}(\cdot|g, \alpha, h, Q')|T_g)] \quad (8)$$

where $T_g \subset T_{SRA}$ is the tangent space for the treatment mechanism g under the assumed model for g . For more details see van der Laan and Robins [To appear in 2002].

Since the TO estimator $\widehat{\alpha}_n^{to}$ is asymptotically linear with influence curve given above, the asymptotic covariance matrix of $\widehat{\alpha}_n^{to}$ can be estimated by

$$\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \widehat{IC}(O_i)^{\otimes 2}.$$

Since it is impractical to actually take the projection in equation (8), we can estimate the asymptotic covariance matrix of $\widehat{\alpha}_n^{to}$ conservatively by

$$\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \widehat{IC}_{to}(O_i|g_n, \widehat{\alpha}_n^{to}, h_n, Q_n)^{\otimes 2}.$$

If $Q' = Q$, then the two influence curves used in variance estimation agree.

The above variance estimates can be used to construct a 95% confidence interval for the j -th component of α ,

$$\widehat{\alpha}_j^{to} \pm 1.96 \frac{\widehat{\Sigma}_{jj}}{\sqrt{n}}.$$

3 Causal inference from censored data

Suppose that the survival time S is subject to right censoring, possibly informative, and that C is the censoring time; that is, we observe

$$O = (R = S \wedge C, \bar{A}^t(R), \bar{Y}(R), \bar{L}(R), W).$$

Here we use $A^t(\cdot)$ to denote the usual treatment process; the superscript t has been added to emphasize the treatment action. We define a censoring process $A^c(t_k) = A_k^c = I(C \leq t_k)$ and the *action process* $A = (A^t, A^c)$ will now be defined more generally and refer to both treatment and censoring. That is, instead of just making a treatment decision at time t_k , we take an action that consists of (1) choosing a treatment and (2) determining censorship status. If the censoring time is unobserved, i.e. if $R = S$, then $C = \infty$ by definition. The observed data can be represented by

$$O = (R = S \wedge C, \bar{A}(R), \bar{X}_{\bar{A}^t}(R)),$$

where $\bar{X}_{\bar{A}^t}(R) = (\bar{L}_{\bar{A}^t}(R), \bar{Y}_{\bar{A}^t}(R), W)$, $S = S_{\bar{A}^t}$, and $\bar{A}(R) = (\bar{A}^t(R), \bar{A}^c(R))$. We assume that the counterfactual response and covariate processes evolve independently of the censoring process. This can be viewed as a general censored data structure where A is the censoring variable and the full data is $X = (X_{\bar{a}^t} : \bar{a}^t \in \mathcal{A}^t)$. The sequential randomization assumption in this case is

$$g(A_k | \bar{A}_{k-1}, X) = g(A_k | \bar{X}_{\bar{A}^t}(t_k), \bar{A}_{k-1})$$

Thus, we have

$$\begin{aligned} g(\bar{A}|X) &= \Pi_k g(A_k | \bar{X}_{\bar{A}^t}(t_k), \bar{A}_{k-1}) = \Pi_k g(A_k^c, A_k^t | \bar{X}_{\bar{A}^t}(t_k), \bar{A}_{k-1}) \\ &= \Pi_k g(A_k^c | \bar{X}_{\bar{A}^t}(t_k), \bar{A}_{k-1}, A_k^t) \times \Pi_k g(A_k^t | \bar{X}_{\bar{A}^t}(t_k), \bar{A}_{k-1}). \end{aligned} \quad (9)$$

We construct an inverse probability of action (IPAW) estimator as an extension of the IPTW estimator constructed in section 2.1. We simply must extend the stabilized weights to include information on the censoring process as well as the treatment process. The IPAW estimating function is given by

$$IC_{ipaw}(O|g, \alpha, h) = \sum_k sw(t_k) \times I(A_k^c = 0) \times h(t_k, \bar{A}_k^t, \bar{Y}_k, V) \times \epsilon_{\bar{A}^t}(\alpha),$$

where the stabilized weight $sw(t_k)$ is now given by

$$\begin{aligned} sw(t_k) &= \frac{g(\bar{A}_k^t, \bar{A}_k^c = 0 | V)}{g(\bar{A}_k^t, \bar{A}_k^c = 0 | X)} \\ &= \frac{\prod_{j=0}^k g(A_j^c = 0 | \bar{A}_j^t, \bar{A}_{j-1}^c = 0, V)}{\prod_{j=0}^k g(A_j^c = 0 | \bar{X}_{\bar{A}^t}(t_j), \bar{A}_j^t, \bar{A}_{j-1}^c = 0)} \\ &\quad \times \frac{\prod_{j=0}^k g(A_j^t | \bar{A}_{j-1}^t, \bar{A}_{j-1}^c = 0, V)}{\prod_{j=0}^k g(A_j^t | \bar{X}_{\bar{A}^t}(t_j), \bar{A}_{j-1}^t, \bar{A}_{j-1}^c = 0)}. \end{aligned}$$

Once again, we can use standard models, such as logistic regression, for the censoring process A^c . Given a choice of h , such as the usual choice, we can then solve the corresponding estimating equation for α as in section 2. For example, we can use the S-plus function `glm()` with weights $I(A_k^c = 0) sw(t_k)$ at time t_k .

The action orthogonalized estimating function is constructed by subtracting from IC_{ipaw} its projection on the tangent space T_{SRA} of the nonparametric model for $g(\bar{A} | X)$ defined by (9). Let $\mathcal{F}_k^c = (\bar{X}_{\bar{A}^t}(t_k), \bar{A}_{k-1}, A_k^t)$ and $\mathcal{F}_k^t = (\bar{X}_{\bar{A}^t}(t_k), \bar{A}_{k-1})$. We have

$$\begin{aligned} IC_{SRA} &= \Pi(IC_{ipaw} | T_{SRA}) \\ &= \sum_k (E(IC_{ipaw} | \mathcal{F}_k^c, A_k^c) - E(IC_{ipaw} | \mathcal{F}_k^c)) \\ &\quad + \sum_k (E(IC_{ipaw} | \mathcal{F}_k^t, A_k^t) - E(IC_{ipaw} | \mathcal{F}_k^t)) \end{aligned}$$

Note that the differences inside the above sums are zero when $t_k > C \wedge S$ and thus the summation ranges from t_0 to the last time t_k that precedes both censoring C and death S . In the same manner as we did for the TO estimating function, we can consider $IC_{SRA}(O|g, Q)$ as a function of the data O , the action mechanism g and the conditional expectations $Q = (Q^c, Q^t)$ of $IC_{ipaw}(\cdot | g, \alpha, h, Q)$, given \mathcal{F}_k^c and \mathcal{F}_k^t , respectively. The action orthogonalized estimating function is now defined by $IC_{ao}(O|g, \alpha, h, Q) = IC_{ipaw}(O|g, \alpha, h) - IC_{SRA}(O|g, Q)$.

4 Simulation Studies

In this section, we carry out a simulation study to compare various estimators of the causal parameter α :

- The naive estimator, denoted $\tilde{\alpha}_n$, obtained by a conventional regression approach of modeling the intensity conditional on the observed predictors, in this case the treatment history.
- The IPTW estimator $\hat{\alpha}_n^{iptw}$ of section 2.1, in which the naive estimating equation is weighted according to estimated propensity scores, i.e. the probability of the observed treatment given the past.

- The one-step treatment-orthogonalized IPTW estimator $\hat{\alpha}_n^{to}$ of section 2.2.

The mean squared error (MSE) of the naive estimator $\tilde{\alpha}_n$ is shown to be much greater than that of the IPTW and TO estimators; this is not surprising given the known bias of the naive estimator. Furthermore, the MSE of the IPTW estimator is shown to be much greater than that of the TO estimator in the presence of severe confounding.

Section 4.1 describes the causal marginal structural model and the treatment mechanism we use in the simulation study. Section 4.2 describes how we analyze the simulated data and thus includes concrete details on implementing these estimators. Section 4.3 gives the data-generating parameter values and all of the simulation results. Section 4.4 discusses the influence of censoring on the relative efficiency of the IPTW and TO estimators.

4.1 Data generating model

We continue to work in a discrete time setting, with a finite number of monitoring times. We also use the same survival counting process introduced earlier. The real-world example that loosely motivates the data generating model is a study of survival in AIDS patients. At measuring times, which are a subset of the monitoring times, we record a covariate process. The time-dependent covariate L_k is real-valued and always positive; a good example would be viral load. We specify a covariate process that grows linearly with time, with subject-specific slopes and intercepts that may be affected by treatment.

At measuring times, we may change the treatment and this decision may be a function of the covariate and treatment history of the subject. The treatment A_k takes on the values 0 and 1, corresponding to 'off' and 'on' treatment, respectively; an appropriate example might be the initiation of highly active anti-retroviral therapy (HAART). Once a subject is on treatment, s/he will remain so until failure. And, of course, a subject will not be considered for treatment after failure. Given the above considerations, we are left with a relatively small set of times at which treatment can change. At a treatment time t_k , the probability of initiating treatment is specified by the following logistic intensity model:


$$E(dA_k | \bar{L}_k, \bar{A}_{k-1}) = I(\bar{A}_{k-1} = 0) \times \text{logit}^{-1}(\theta_0 + \theta_1 t_k + \theta_2 L_k). \quad (10)$$

Therefore partial likelihood estimation can be used to estimate θ , as was mentioned in section 2. If a subject goes on treatment, the treatment initiation time is denoted t^* .

We now state the MSM and explain its motivation. In reality, the hazard obviously depends on how sick a patient is at a given time and an effective treatment mitigates this. In our simple story, this is captured by the covariate, e.g. viral load. The treatment effect might take the form of a one-time reduction in L_k followed possibly by a period of no change, followed by resumed linear growth, perhaps with a persistent change in the slope. One could even imagine that the one-time shift is proportional to the current viral load. Therefore, the optimal treatment time is not apparent, in that it may be possible to treat too early (implies a small one-time reduction in viral load) or too late (too much cumulative time with a high viral load). In this situation, the following MSM arises. The probability of failure in the upcoming interval is given by the intensity model

$$\lambda(t_k, \bar{a}, \bar{Y}_{\bar{a}}(t_{k-1}) | \alpha) = I(\bar{Y}_{\bar{a}}(t_{k-1}) = 0) \times \text{logit}^{-1}(\alpha_0 + d_1(t_k)\alpha_1 + a_k\alpha_2 + d_3(t_k)\alpha_3) \quad (11)$$

where $d_1(t_k) = (1 - a_k)t_k + a_k t^*$ and $d_3(t_k) = a_k(t_k - t^*)$. We see that the probability of failure depends on the current treatment status a_k and on either the study time elapsed t_k (for subjects off treatment) or the time since treatment initiation $t_k - t^*$ (for subjects on treatment). If there is no treatment effect, $\alpha_2 = 0$ and $\alpha_1 = \alpha_3$. All other things held equal, $\alpha_2 < 0$ corresponds to a positive treatment effect, i.e. treatment causes a persistent decrease in the hazard. The case $\alpha_3 < \alpha_1$ also corresponds to a positive treatment effect, i.e. treatment causes the hazard to grow more slowly as a function of time. In a situation where $\alpha_2 < 0$, but $\alpha_3 > \alpha_1$, the effect of treatment is ambiguous. At certain times, it is less hazardous to be on treatment, while at others, it is less hazardous to be off treatment. In reality, the outcome of interest is survival time $S_{\bar{a}}$ and, therefore, the optimal treatment regime is not immediately apparent. We note that a 'treatment regime' in this setting is completely



/home/jenny/research/msm/paper/images/true.surv.eps

Figure 1: True survival curves for different treatment regimes.

specified by the treatment initiation time. Figure 1 depicts the counterfactual survival curves for all possible treatments in the simulations described below and we see that, w.r.t median survival, the optimal treatment initiation time is t_5 .

The only remaining detail is exactly how to induce confounding. There must be a variable that affects both the treatment and outcome processes, but the distribution of the counterfactual survival process is only specified as a function of treatment. Conceptually, for each subject, we would like to have counterfactual outcomes for every possible path of the treatment process. We could then sequentially generate a particular outcome of the covariate and treatment processes, extract the appropriate counterfactual, and record it as the observed outcome. In practice, however, it is not computationally feasible to generate the full data on each subject.

The solution is to realize that the conditional hazard from the MSM in equation (11) completely determines the distribution function corresponding to any treatment path through the identity

$$F_{\bar{a}}(t) = 1 - S_{\bar{a}}(t) = 1 - \prod_{k:t_k \leq t} (1 - \lambda_{\bar{a},k}).$$

We can generate survival times with this distribution by applying the inverse of the CDF to random uniform $[0,1]$ deviates. Now imagine we draw an underlying health state $U \sim Unif[0,1]$ for each subject. By applying the appropriate inverse CDF to these health states, we can obtain that subject's counterfactual survival time under any treatment path. We accumulate a treatment history by sequentially generating covariate and treatment processes, constantly updating the cumulative hazard

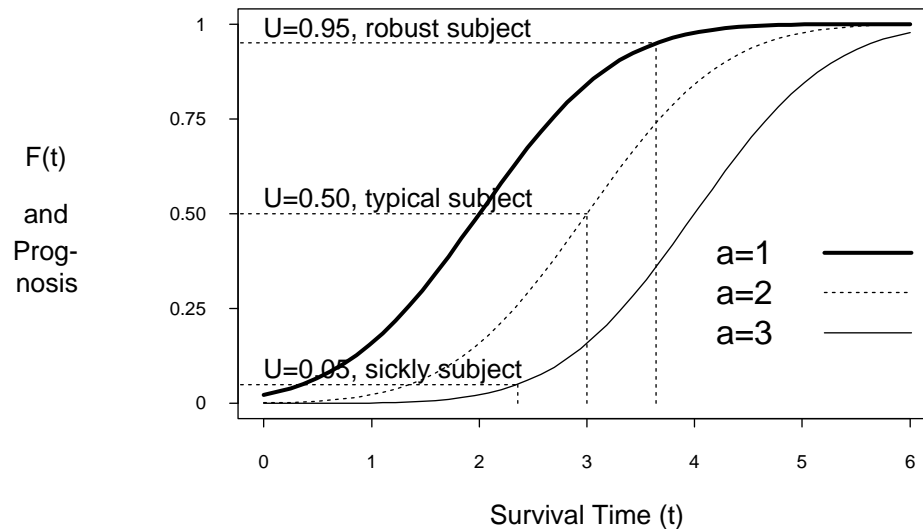


Figure 2: Illustration of confounding in a hazard MSM.

through the product of conditional hazard terms implied by the observed treatment. As soon as the cumulative hazard exceeds U , the subject fails and we obtain the observed survival time.

Figure 2 illustrates this basic idea in a simpler point-treatment example. It is apparent that the distribution of failure times for treatment 3 lies fully to the right of that for treatment 2 (and likewise for treatment 1). However, the most robust subjects, e.g. those with values of U closer to 1 than 0, are preferentially assigned to the worst treatment, treatment 3. And so, from observed data, it appears that treatment 1 induces survival times longer than those associated with treatment 2 (and likewise for treatment 3).

Here is pseudo-code for simulating data for one subject in our longitudinal setting:

- Draw U from a uniform $[0,1]$ distribution.
- Map U into a slope and intercept that describe the covariate process $L(\cdot)$.
- Loop through the t_k while subject still alive:
 - If t_k is a potential treatment time, measure the covariate and make a treatment decision. If t_k is not a potential treatment time, copy the covariate and treatment process state at t_{k-1} into t_k .
 - Using the MSM in equation (11), calculate the probability of surviving the interval (t_k, t_{k+1}) and record it. Take the product of all such probabilities for $t_j : t_j \leq t_k$. If this product is less than or equal to $1 - U$, then $Y(t_k) = 1$ and the subject has failed. Otherwise, $Y(t_k) = 0$ and the subject lives on.

4.2 Analysis of simulated data

First we describe the observed data structure. Each subject will exhibit an outcome of the survival process Y that is a vector of zeros followed by exactly one one, i.e. something of the form $(0, 0, \dots, 1)$. Of the same length as this vector, we will have the outcome of the covariate and treatment processes. By concatenating this subject-specific data, we create a dataset from which to estimate the treatment mechanism and the causal parameter α .

The IPTW estimator relies on solving weighted estimating equations, where the weights arise from the treatment mechanism. Here we describe how to compute the estimated stabilized weights $\widehat{sw}(t_k)$. Due to the curse of dimensionality, typically one cannot estimate g using the SRA alone;

it is usually necessary to assume a lower dimensional parametric or semiparametric model for the treatment mechanism. For example, we could assume a logistic model for the probability of treatment given the observed past (or perhaps some portion thereof).

Since we have used a logistic model to generate the data in this simulation, we know that this model will be correct. By using the treatment and covariate portion of the dataset described above, we estimate the parameter θ of the treatment mechanism in equation (10) with ordinary logistic regression. We treat each subject at each time t_k as an observation. Denote the resulting estimator $\hat{\theta}$. Let \hat{p}_{ik} be the implied estimated treatment probability at time t_k given the history of subject i . We also fit a second logistic regression model, but this time omitting L_k as a predictor. Denote the resulting estimator $\tilde{\theta} = (\tilde{\theta}_0, \tilde{\theta}_1, 0)$ and let \tilde{p}_{ik} denote the associated estimated treatment probabilities. The estimated stabilized weight $\widehat{sw}_i(t_k)$ at time t_k for subject i is given by the product

$$\widehat{sw}_i(t_k) = \prod_{j=0}^k \frac{\tilde{p}_{ij}^{A_{ij}} (1 - \tilde{p}_{ij})^{1-A_{ij}}}{\hat{p}_{ij}^{A_{ij}} (1 - \hat{p}_{ij})^{1-A_{ij}}}$$

Note that in the absence of confounding, i.e. if the covariate $L(\cdot)$ does not affect treatment decisions, the true stabilized weights are always one.

At this point we can calculate the IPTW estimator. We fit a logistic regression, using the outcomes from all subjects, of Y on the predictors used in the MSM given in equation (11). For example, one can use the Splus function `glm` and provide estimated weights $\widehat{sw}(t_k)$ through the `weights` argument; Splus code for an example can be found in the appendix. We carry this out using no weights and stabilized weights, which results in the naive estimator $\tilde{\alpha}_n$ and the IPTW estimator $\hat{\alpha}_n^{iptw}$ respectively. Both are estimators of the causal parameter α .

To calculate the TO estimator recall the special form of the projection (see equation (7)) when the treatment process takes on the values zero or one:

$$IC_{SRA}(O|g, Q) = \sum_k (Q(1, \mathcal{F}_k) - Q(0, \mathcal{F}_k)) dM_g(k), \quad (12)$$

where $dM_g(k) = A_k - E(A_k|\mathcal{F}_k)$ and $Q(a, \mathcal{F}_k) = E(IC_{iptw}(O|g, \alpha, h)|A_k = a, \mathcal{F}_k)$. We estimate the expectation $Q(\cdot)$ by regressing the estimated IPTW estimating function $IC_{iptw}(O_i|g_n, \hat{\alpha}_n^{iptw}, h)$ on the treatment and covariate histories \bar{A}_k and \bar{L}_k . From the fitted regression and an estimated treatment mechanism, we can compute the projection \widehat{IC}_{SRA} given in (12) for each subject. We subtract this from \widehat{IC}_{iptw} to obtain \widehat{IC}_{to} and use equation (6) to calculate the one-step TO estimator. Some Splus code for an example is in the appendix and we discuss this in more detail for individual simulations.

4.3 Simulation results

In all of the simulations described below, the study time τ is the interval $[0, 40]$. The monitoring times are $\{t_0 = 0, t_1 = 1, \dots, t_{39} = 39\}$. We measure the covariate once in every five intervals, therefore the measuring times are $\{t_0, t_5, \dots, t_{35}\}$. The treatment times are t_0 and t_5 , therefore the possible treatment paths can be listed easily by considering the treatment actions a_0 and a_5 :

$$\mathcal{A} = \{(a_0, a_5) : a_0, a_5 \in \{0, 1\}, a_0 \leq a_5\} = \{(0, 0), (0, 1), (1, 1)\}.$$

Here is the value of α , the parameter of the logistic MSM given in equation (11):

α_0	α_1	α_2	α_3
-3.04	0.175	-1.5	0.388

Recall that the counterfactual survival curves implied by α are given in figure 1.

4.3.1 Simulation 1

The main point of this simulation is to demonstrate the relative performance of the naive and IPTW estimators under conditions ranging from no confounding to quite severe confounding. The presence

and degree of confounding is completely determined by the parameter θ in model (10). No confounding implies that the coefficient for L_k equals zero. The degree of confounding intensifies as that coefficient is increased. In the case of severe confounding, the coefficient for time t_k equals zero and that for the covariate L_k does not; thus, all differences in the probability of initiating treatment are attributable to differences in the observed covariate. Here are the four values of θ that we used, listed in increasing order w.r.t confounding severity:

	θ_0 (Intercept)	θ_1 (Time)	θ_2 (Covariate)
$\theta^{(1)}$	-0.58	0.08	0.00
$\theta^{(2)}$	-0.71	0.06	0.01
$\theta^{(3)}$	-0.85	0.03	0.03
$\theta^{(4)}$	-0.99	0.00	0.04

The following table presents the average naive and IPTW estimates of α , under varying degrees of confounding. As a reference, the first row contains the true data-generating value of α . Recall that $\theta^{(1)}$ corresponds to no confounding and that $\theta^{(4)}$ corresponds to severe confounding.

	Naive estimator				IPTW estimator			
	α_0	α_1	α_2	α_3	α_0	α_1	α_2	α_3
Truth	-3.04	0.175	-1.5	0.388	-3.04	0.175	-1.5	0.388
$\theta^{(1)}$	-3.037	0.177	-1.494	0.388	-3.031	0.176	-1.498	0.388
$\theta^{(2)}$	-3.003	0.179	-1.602	0.391	-3.039	0.176	-1.487	0.386
$\theta^{(3)}$	-2.990	0.183	-1.722	0.401	-3.059	0.179	-1.511	0.391
$\theta^{(4)}$	-2.943	0.185	-1.918	0.413	-3.063	0.177	-1.543	0.397

We see that the average IPTW estimate is close to the corresponding true value, even in the presence of confounding, whereas the average naive estimate drifts away from the truth as the confounding increases. This is especially pronounced for α_2 , which is arguably the coefficient most directly responsible for the treatment effect. In figure 3 we display the relative efficiency of the estimators for α_2 ; qualitatively similar results are obtained for the remaining components of α . The total bar length reflects mean squared error and we see that MSE for the naive estimator is always greater than for the IPTW. Notice this is true even in the case of no confounding, e.g. θ^1 , which is an illustration of the efficiency gain from weighted estimators even in the absence of confounding. The difference in MSE becomes more dramatic as the confounding becomes more severe and this is driven by the increasing bias of the naive estimator.

4.3.2 Simulation 2

In this simulation, we introduce the TO estimator and compare the performance of all three estimators discussed in this paper. We consider only one value of the treatment mechanism parameter θ : namely, $\theta = (\theta_0, \theta_1, \theta_2) = (-0.8, 0, 0.04)$. Since θ_2 does not equal zero, we see that confounding is present.

At this point, we provide more details on the regression used to estimate $Q(A_k, \mathcal{F}_k)$, which is the conditional expectation of the IPTW estimating function $IC_{iptw}(O|g, \alpha, h)$, given the past A_k and \mathcal{F}_k . This is the greatest challenge in calculating the TO estimator. In this simulation $IC_{iptw}(O|g, \alpha, h)$ is a four dimensional vector, since α is four dimensional, and it is a subject-specific variable. Whenever we require $IC_{iptw}(O|g, \alpha, h)$ for each subject observation, such as in the regressions described below, the subject-specific vector is replicated. For concreteness, we concentrate on the first component of α the approach to the other three components is identical. In panel (a) of figure 4 we use one of the simulated datasets to present a typical scatterplot of the first component of \widehat{IC}_{iptw} versus L_k for all observations contributed by subjects when they are off treatment. Panel (b) of figure 4 presents observations from time 0 among treated subjects and panel (c) presents the same from time 1. The plot in panel (c) exhibits two distinct curves, corresponding to the subjects assigned to treatment at time t_0 , i.e. $t^* = t_0$, and those assigned at time t_5 , i.e. $t^* = t_5$. Using data from off-treatment subjects, we use linear regression to model the expectation of $IC_{iptw}(O|g, \alpha, h)$ as a function of L_k

Relative Efficiency for α_2

		Mean Squared Error			
		Squared Bias		Variance	
θ^1	Naive	0	+	0.077	= 0.077
	IPTW	0	+	0.051	= 0.051
θ^2	Naive	0.01	+	0.06	= 0.07
	IPTW	0	+	0.046	= 0.046
θ^3	Naive	0.049	+	0.083	= 0.132
	IPTW	0	+	0.052	= 0.052
θ^4	Naive	0.174	+	0.078	= 0.252
	IPTW	0	+	0.048	= 0.048

/home/jenny/research/msm/paper/rev_jenny4/sim1.rel.eff.eps

Figure 3: Relative efficiency of naive and IPTW estimators in Simulation 1.

and L_k^2 , for times 0 and 5 separately. Among on-treatment observations – that is, those depicted in panels (b) and (c) of figure 4 – it is clear that our model must include the treatment initiation time t^* . In this case, we use linear regression to model the expectation of $IC_{iptw}(O|g, \alpha, h)$ as a function of L_k, L_k^2, t^* , and $t^* \times L_k$ for times 0 and 5 separately.

Figure 5 depicts the relative performance of the three estimators for α_2 with respect to both bias and variance. Note that α_2 is our main parameter of interest since its value contributes heavily to the causal effect of the treatment regime. As expected, we see the inferior performance of the naive estimator, with respect to both bias and variance. When compared to the naive estimator, the IPTW estimator has a relative efficiency of four. Furthermore, the TO estimator has a relative efficiency of four, compared to the IPTW estimator. The advantages of the IPTW and TO estimators are substantial, in terms of bias and variance.

Finally, we evaluate the projection fit through inspection of the empirical inner products between elements of \widehat{IC}_{SRA} and \widehat{IC}_{to} , as described in section 2.2. The $\widehat{\rho}$ matrix is given below and we confirm that all elements are indeed close to zero:

$$\widehat{\rho} = \begin{pmatrix} 0.0459 & 0.0552 & 0.0453 & 0.0502 \\ 0.0447 & 0.0463 & 0.0328 & 0.0419 \\ 0.0673 & 0.0773 & 0.0543 & 0.0696 \\ 0.0787 & 0.0840 & 0.0586 & 0.0736 \end{pmatrix}$$

4.3.3 Simulation 3

The only difference between this simulation and the previous, described in section 4.3.2, is an additional treatment time at $t_{10} = 10$. This has the effect of increasing the size of the treatment space \mathcal{A} :

$$\begin{aligned} \mathcal{A} &= \{(a_0, a_5, a_{10}) : a_0, a_5, a_{10} \in \{0, 1\}, a_0 \leq a_5 \leq a_{10}\} \\ &= \{(0, 0, 0), (0, 0, 1), (0, 1, 1), (1, 1, 1)\}. \end{aligned}$$

All of the estimators are implemented in the manner detailed in previous sections, except we perform a pooled regression when estimating the projection, i.e. we do not fit separate models at different times. The relative performance of the estimators is summarized in figure 6; the results are essentially equivalent to those seen in the previous simulation. Likewise, the inner product matrix $\widehat{\rho}$ indicates an acceptable projection fit.

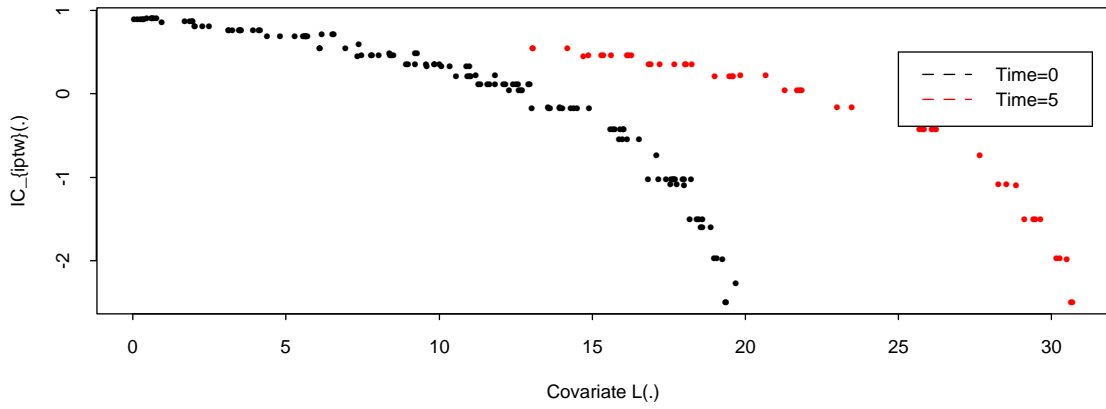
$$\widehat{\rho} = \begin{pmatrix} 0.0493 & 0.0586 & 0.0370 & 0.0528 \\ 0.0780 & 0.0703 & 0.0550 & 0.0632 \\ 0.0731 & 0.0768 & 0.0585 & 0.0721 \\ 0.0733 & 0.0822 & 0.0550 & 0.0769 \end{pmatrix} \tag{13}$$

4.4 Influence of end-of-study censoring


In the above simulations, the end-of-study time is $T = 40$ and, as was our intention, all subjects fail before this time in our simulated datasets. Since it is interesting to determine the effect of noninformative censoring on the relative performance of the estimators, we re-examine the simulated datasets and impose earlier end-of-study times, thereby creating censored observations. That is, we redefine the outcome process as $Y_{\bar{a}}(t_k) = I(S_{\bar{a}} \leq t_k, t_k < T)$. The following table gives the relative efficiency w.r.t MSE of the IPTW estimator versus the naive estimator (columns 1 - 3) and that of the TO estimator versus the IPTW estimator (columns 4 - 6) for different end-of-study times.

End-of-study time	Naive vs IPTW			IPTW vs TO		
	7	10	40	7	10	40
α_0	0.807	0.827	0.707	0.853	0.957	0.986
α_1	0.903	0.738	0.657	0.904	0.935	0.825
α_2	0.523	0.331	0.255	0.700	0.509	0.274
α_3	0.959	0.952	0.737	0.953	0.738	0.633

(a) 'No Treatment' Observations (A=0) at times 0 and 5



(b) 'Treatment' Observations (A=1) at time 0



(c) 'Treatment' Observations (A=1) at time 5




Figure 4: First component of $IC_{iptw}(O|g, \alpha, h)$.

Relative Efficiency for α

		Mean Squared Error		MSE	Rel Eff of MSE
		Squared Bias	Variance		
α_3	TO	0.0000 +	0.0005 =	0.0005	0.48
	IPTW	0.0000 +	0.0010 =	0.0010	
	Naive	0.0004 +	0.0010 =	0.0014	
α_2	TO	0.0000 +	0.0134 =	0.0134	0.26
	IPTW	0.0014 +	0.0492 =	0.0506	
	Naive	0.1292 +	0.0766 =	0.2058	
α_1	TO	0.0000 +	0.0003 =	0.0003	0.77
	IPTW	0.0000 +	0.0004 =	0.0004	
	Naive	0.0001 +	0.0004 =	0.0005	
α_0	TO	0.0000 +	0.0183 =	0.0183	0.89
	IPTW	0.0001 +	0.0206 =	0.0207	
	Naive	0.0152 +	0.0245 =	0.0397	

/home/jenny/research/msm/paper/rev_jenny4/sim2.rel.eff.eps

Figure 5: Relative efficiency of all three estimators in Simulation 2.

Relative Efficiency for α

		Mean Squared Error			Rel Eff of MSE
		Squared Bias	Variance	MSE	
α_3	TO	0.0000 +	0.0008 =	0.0008	0.88 0.69
	IPTW	0.0000 +	0.0009 =	0.0009	
	Naive	0.0003 +	0.0010 =	0.0013	
α_2	TO	0.0000 +	0.0181 =	0.0181	0.56 0.18
	IPTW	0.0002 +	0.0323 =	0.0325	
	Naive	0.1165 +	0.0662 =	0.1827	
α_1	TO	0.0000 +	0.0004 =	0.0004	0.81 0.43
	IPTW	0.0000 +	0.0005 =	0.0005	
	Naive	0.0005 +	0.0007 =	0.0012	
α_0	TO	0.0000 +	0.0202 =	0.0202	0.85 0.75
	IPTW	0.0000 +	0.0238 =	0.0238	
	Naive	0.0071 +	0.0247 =	0.0318	

/home/jenny/research/msm/paper/rev_jenny4/sim3.rel.eff.eps

Figure 6: Relative efficiency of all three estimators in Simulation 3.




Figure 7: Overview of SPARCS timeline.

We see that, as the proportion of censored subjects decreases or, equivalently, as end-of-study time grows later, the relative efficiencies become more dramatic. This can be seen in the general downward trend of the three relative efficiencies reported for any component of α , for either estimator comparison. The overall conclusion is that the performance gains for the IPTW and TO estimators are largest in the absence of censoring, although they persist even in the presence of censoring.

5 Analysis of SPARCS data

Here we apply the methodology developed in previous sections to analyze data from a project entitled “Study of Physical Performance and Age Related Changes in Sonomans” (SPARCS) [Tager et al., 2000]. SPARCS is a community-based longitudinal study of physical activity and fitness in people at least 55 years of age who live in Sonoma, California. One of the goals of SPARCS and the primary goal of the current analysis is to estimate the causal effect of increased physical activity on survival.

5.1 Data structure

The subset of the data that we examine here was collected in the first three home evaluations of female SPARCS participants ($n = 1197$), over the time period 5/1993 - 10/1999. We treat time as discrete and divide the study time into 6 month intervals. This implies that the monitoring times are given by $\tau = \{t_0, \dots, T = t_{10}\}$; at all time points t_k , we determine if the subject is alive. See figure 7 for a helpful picture. To simplify the analysis we proceed as if subjects were evaluated every 2.5 years; therefore, the measuring times are $t_0 = 0$ years, $t_5 = 2.5$ years, and $t_{10} = 5$ years (the actual time spacing of evaluations ranged from 2 to 3 years). Also, recall that the treatment and covariate processes are subject to change only at the measuring/treatment times t_0, t_5 , and t_{10} .

Our measure of physical activity is based on an *activity score* that is recorded for each subject at each evaluation. The activity score takes values in the set $\{1, 2, 3, 4\}$, where 4 corresponds to the highest level of activity. We define a time-dependent treatment process A_k^t that is an indicator for an activity score of 3 or 4 during the interval (t_k, t_{k+1}) , which implies that the subject is engaging in moderately vigorous activity. Note that, although subjects are not being actively treated in any way, we can simply handle a subject’s self-chosen activity level as an intervention whose efficacy we wish to measure. Following the approach introduced in section 1.1, we also define a counting process Y_k that is an indicator for death in the interval (t_{k-1}, t_k) . An uncensored subject’s survival time, denoted S to distinguish it from the end-of-study time T , can be recovered from the history of this process $\bar{Y}(T)$ by noting the jump time, i.e. $S = \min_k \{t_k : Y_k > 0\}$.

At the initial evaluation, information on the the following baseline covariates is obtained: age in years (*age*), indicator of activity decline in past 5 - 10 years, indicator of past habitual vigorous activity, indicator of participation in high school sports. This collection of variables is referred to collectively as W . At each evaluation, including the baseline evaluation, information on the following

time-dependent covariates is obtained: indicator of a cardiovascular condition (*card*), indicator of other health conditions, self-perception of health, NRB – a laboratory measure of performance in a variety of everyday tasks, BMI – body mass index, indicator of current smoking status, indicator of past smoking status. These variables will be referred to collectively as L_k .

It is reasonable to assume that the variables W and L_k can influence both the activity level and the survival time. Therefore, we must regard W and L_k as potential confounders in our study of the causal relationship between activity level and survival.

Two types of censoring are present in the data. The first type arises if a subject drops out of the study, which affects 58 participants. We refer to the drop out time as D . The second type of censoring arises if the subject survives the entire study time $[0, T]$. Of the 1197 participants, 958 were alive at their third and final evaluation. The remaining 181 subjects died during the study time. Each subject accumulates a history until the earliest of these events: death S , end-of-study T , or drop-out D .

We use the following imputation procedure to address missing values: missing continuous covariates are replaced by the mean of the nonmissing covariates for that subject at other time points. Missing categorical covariates are set to the nonmissing value measured at a nearby timepoint. Subjects with missing values for a covariate at all time points were excluded from the analysis.

5.2 Causal models

The full data for a subject includes the uncensored survival process (equivalent to the survival time $S_{\bar{a}^t}$) for every possible treatment history and the corresponding covariate process:

$$X = (S_{\bar{a}^t}, \bar{a}^t(S), \bar{L}_{\bar{a}^t}(S), W : \bar{a}^t \in \mathcal{A}^t)$$

In contrast, the observed data includes only the survival and covariate process corresponding to the subject's actual treatment history, possibly subject to censoring:

$$O = (R = S \wedge T \wedge D, \bar{A}^t(R), \bar{L}(R), W).$$

We consider two marginal structural models for the hazard of death as a function of time, the activity level $A^t(\cdot)$, and, perhaps, two baseline covariates: cardiovascular condition (*card*) and age (*age*). Since we are modeling hazard in a discrete time setting, we focus on the probability of a jump in the survival process $Y_{\bar{a}^t}(\cdot)$. Model 1 is given by

$$\begin{aligned} E(dY_{\bar{a}^t}(t_k) | \bar{Y}_{\bar{a}^t}(t_{k-1})) &= I(\bar{Y}_{\bar{a}^t}(t_{k-1}) = 0) \times \\ &\text{logit}^{-1}(\alpha_0 + \alpha_1 a^t(t_k) + \alpha_2 t_k + \alpha_3 \text{age} + \\ &\alpha_4 \text{card} + \alpha_5 (\text{age} \times a^t(t_k)) + \\ &\alpha_6 (\text{age} \times t_k) + \alpha_7 (\text{card} \times a^t(t_k)) + \\ &\alpha_8 (\text{card} \times t_k) + \alpha_9 (\text{age} \times \text{card})). \end{aligned}$$

A simpler alternative, Model 2, is given by

$$\begin{aligned} E(dY_{\bar{a}^t}(t_k) | \bar{Y}_{\bar{a}^t}(t_{k-1})) &= I(\bar{Y}_{\bar{a}^t}(t_{k-1}) = 0) \times \\ &\text{logit}^{-1}(\alpha_0 + \alpha_1 a^t(t_k) + \alpha_2 t_k). \end{aligned}$$

In order to form the weights necessary to use the observed data to estimate the causal parameter α , from either Model 1 or 2, we must model the treatment mechanism. We assume the following logistic regression model:

$$P(A^t(t_k) = 1 | \bar{A}^t(t_{k-1}), \bar{L}(t_k), W) = \lambda_{t_k}(\gamma), \quad (14)$$

where $\lambda_{t_k}(\gamma) = \text{logit}^{-1}(\gamma_0 + \gamma_1 A^t(t_{k-1}) + \gamma_2 t_k + \gamma_3^T L_k + \gamma_4^T W)$. The usual partial mle γ can be computed using standard software and we compute the treatment contribution to the estimated weights as described before in section 4.2.

We do not assume that the drop-out time D is independent of survival, but we do assume independence conditional on the observed covariate history. Therefore, our weights will include the

Table 1: Estimates of the causal parameter α in Model 1.

	Naive	IPAW	TO			
$\hat{\alpha}_0$	-12.5435	-13.2424	(1.42)	-12.9429	(1.40)	Intercept
$\hat{\alpha}_1$	0.6861	1.1873	(1.47)	0.8018	(1.43)	Activity level
$\hat{\alpha}_2$	-0.1163	-0.0618	(0.16)	-0.0825	(0.15)	Time
$\hat{\alpha}_3$	0.1048	0.1132	(0.02)	0.1091	(0.02)	Age at baseline
$\hat{\alpha}_4$	4.8012	4.2989	(1.66)	4.4791	(1.69)	Cardiovascular
$\hat{\alpha}_5$	-0.0202	-0.0263	(0.02)	-0.0204	(0.02)	Age \times activity
$\hat{\alpha}_6$	0.0025	0.0017	(0.00)	0.0019	(0.00)	Age \times Time
$\hat{\alpha}_7$	0.4386	0.6898	(0.37)	0.6885	(0.35)	Cardio \times activity
$\hat{\alpha}_8$	0.0101	0.0191	(0.04)	0.0179	(0.04)	Cardio \times time
$\hat{\alpha}_9$	-0.0535	-0.0488	(0.02)	-0.0509	(0.02)	Age \times cardio

probability of not being censored due to dropout, in addition to the above probability on the treatment mechanism. We define a drop-out censoring process $A^c(t_k) = I(D \in (t_k, t_{k+1}))$. If the subject drops out before death or end of the study, $\bar{A}^c = (0, \dots, 0, 1)$; otherwise $\bar{A}^c = (0, \dots, 0, 0)$. We assume the following intensity model for the drop-out process:

$$E(dA_k^c | \bar{L}_k, W, \bar{A}_k^t, \bar{A}_{k-1}^c) = I(\bar{A}_{k-1}^c = 0) \times \pi_{t_k}(\beta), \quad (15)$$

where $\pi_k(\beta) = \text{logit}^{-1}(\beta_0 + \beta_1 A_k^t + \beta_2 t_k + \beta_3^T L_k + \beta_4^T W)$. Just as with the treatment process, we fit the regression implied by (15). Denote the estimator by $\hat{\beta}$ and the implied probability of dropout at time t_k for subject i by \hat{q}_{ik} . The same quantities, but based on a regression in which \bar{L}_k is omitted as a predictor, are denoted by $\tilde{\beta}$ and \tilde{q}_{ik} . The censoring contribution to the estimated stabilized weight is then

$$\prod_{j=0}^k \frac{(1 - \tilde{q}_{ij})}{(1 - \hat{q}_{ij})}$$

The treatment contribution is calculated as it was earlier in the simulations and the estimated stabilized weights are an element-wise product of treatment and censoring terms. Splus code for a related example is in the appendix.

5.3 Results

Table 1 reports the Model 1 estimates of α based on the naive (unweighted), IPAW, and TO estimators. Estimated standard errors for the IPAW and TO estimators are reported in parentheses. In the context of the MSM, the effect of physical activity on the hazard can be obtained from the coefficients of all terms containing $a^t(\cdot)$. Specifically, let α_{act} reflect the effect of activity on the linear predictor (i.e. the logit of the probability of death); that is, in Model 1, $\alpha_{act}(age, card) = \alpha_1 + \alpha_5 age + \alpha_7 card$. Due to the interaction terms, α_{act} is a function of the baseline covariates age and $card$. We estimate α_{act} by substituting the relevant estimated coefficients. The standard error of $\hat{\alpha}_{act}$ can be recovered from the estimated covariance matrices. Table 2 reports the estimated treatment effect of activity $\hat{\alpha}_{act}$ for subjects of various ages and cardiovascular conditions based on the unweighted, IPAW, and TO estimators.

Among patients without a history of cardiovascular problems, physical activity decreases the hazard, with the effect size steadily increasing with the subject's age at baseline. This trend is seen in all three estimators, with statistical significance at the conventional 0.05 level for ages 70 and higher. For patients that have had cardiovascular problems, the results are more ambiguous. The unweighted estimator indicates that physical activity decreases the hazard, but the effects are much smaller than for the first population. The positive relationship between baseline age and treatment effect also remains. Both the IPAW and TO estimators weakly indicate that physical activity may increase the hazard for patients younger than 75 at baseline. The treatment effect at ages 75 and 80 is positive, but certainly does not approach statistical significance.

Table 2: Estimates of treatment effect α_{act} in Model 1

Age at baseline		60	65	70	75	80
$card = 0$	Naive	-0.5255	-0.6265	-0.7275	-0.8284	-0.9294
	IPAW	-0.3882	-0.5195	-0.6508	-0.7821	-0.9134
		(0.3794)	(0.3073)	(0.2514)	(0.2243)	(0.2362)
	TO	-0.4242	-0.5264	-0.6286	-0.7307	-0.8329
		(0.3691)	(0.2989)	(0.2442)	(0.2169)	(0.2273)
$card = 1$	Naive	-0.0870	-0.1879	-0.2889	-0.3899	-0.4908
	IPAW	0.3016	0.1703	0.0390	-0.0923	-0.2236
		(0.4007)	(0.3416)	(0.3021)	(0.2900)	(0.3088)
	TO	0.2642	0.1621	0.0599	-0.0423	-0.1445
		(0.3822)	(0.3246)	(0.2860)	(0.2746)	(0.2935)

Table 3: Estimates of α in Model 2.

	Naive	IPAW	TO
$\hat{\alpha}_0$	-4.4552	-4.8602 (0.1855)	-4.8532 (0.1782)
$\hat{\alpha}_1$	-1.0833	-0.4594 (0.1846)	-0.4448 (0.1684)
$\hat{\alpha}_2$	0.0755	0.0798 (0.0201)	0.0758 (0.0196)

Another way to summarize the treatment effect is to examine survival curves for different subpopulations. In figure 8, we present survival curves for subjects of different baseline ages and cardiovascular histories. All panels have the same color scheme, horizontal and vertical axes. Columns one and two contain survival curves for subjects without a cardiovascular history, both for the unweighted and the IPAW estimators. Columns three and four present the same information for patients with a cardiovascular history. In every case, survival curves corresponding to lack of activity (i.e. $\bar{a}^t = (0, 0, 0)$) and constant activity (i.e. $\bar{a}^t = (1, 1, 1)$) are drawn in red and green, respectively.

The main discordance between the two estimators is among younger patients with a history of cardiovascular trouble. For ages 60, 65, and 70, the unweighted estimator indicates there is a small benefit from increased activity, whereas the IPAW estimator indicates the effect is either absent or detrimental. Given the weak statistical significance, it is fair to say that the scientific conclusions would likely be very similar from either unweighted or weighted estimation of α , but that a causal interpretation is justified only in the case of the latter.

The estimates of α for Model 2 are reported in table 3. In this model, we choose not to condition on any baseline covariates; we want to state the effect of physical activity on survival for the entire population. For all estimators, the relevant estimated coefficient, $\hat{\alpha}_1$ is negative and, therefore, indicates that activity decreases the hazard. However, the magnitude of this effect is markedly different for the unweighted and weighted estimators. The unweighted estimator indicates a much larger benefit than the weighted estimators. This is quite understandable when one considers the risk of using naive estimators to draw causal inferences. In our setting, the risk is that those who opt for a more active lifestyle are different from those who do not *in a way that also affects survival*. It is quite likely that the people who are actually physically active are healthier across the board than those who do not and that people are more active during periods of health. Therefore, the observed survival experience for 'active intervals' is quite likely to be more positive than it would be if all subjects – regardless of current health status – engaged in activity. Through weighting, the IPAW and TO estimators partially compensate for this and, indeed in this case, present a less dramatic estimate of the universal benefit of physical activity. In fact, even in Model 1, the overall result is that the unweighted estimator indicates a larger, more positive treatment effect than that found by weighted estimation.




Figure 8: Survival curves for subjects aged 60 through 80 at baseline.

6 Conclusion

This paper presents two estimators for the causal parameter in a longitudinal marginal structural model: an IPTW or IPAW estimator and a treatment or action orthogonalized (one-step) estimator. Through weighting and projection, these estimators allow causal inferences to be drawn from data in which both treatment assignment and censoring may depend on past treatment, covariates, and response, i.e. when conventional randomization and censoring assumptions are violated. We have seen that, even in the absence of time-dependent confounding, the IPTW estimator is more efficient than the naive estimator. We note also the substantial efficiency gain of the treatment and action orthogonalized estimator, relative to the initial inverse-weighted estimator. Finally, the proposed estimators are easily implemented with standard software tools, such as Splus or R.

References

- J. M. Robins, M. A. Hernan, and B. Brumback. Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11(5):550–560, 2000.
- James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Math. Modelling*, 7(9-12): 1393–1512, 1986. ISSN 0270-0255. Mathematical models in medicine: diseases and epidemics, Part 2.
- James M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology, the environment, and clinical trials (Minneapolis, MN, 1997)*, pages 95–133. Springer, New York, 2000.
- J.M. Robins. Addendum to: “A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect” [Math. Modelling 7 (1986), no. 9-12, 1393–1512; MR 87m:92078]. *Comput. Math. Appl.*, 14(9-12):923–945, 1987. ISSN 0097-4943.
- J.M. Robins. The analysis of randomized and non-randomized AIDS treatment trials using a new approach to causal inference in longitudinal studies. In L. Sechrest, H. Freeman, and A. Mulley, editors, *Health Service Research Methodology: A Focus on AIDS*, pages 113–159. NCHSR, U.S. Public Health Service, Dordrecht, 1989.
- J.M. Robins. Causal inference from complex longitudinal data. In *Latent variable modeling and applications to causality (Los Angeles, CA, 1994)*, pages 69–117. Springer, New York, 1997.
- J.M. Robins. Marginal structural models. In *1997 Proceedings of the Section on Bayesian Statistical Science*, pages 1 – 10, Alexandria, VA, 1998. American Statistical Association.
- J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference. In *Proceedings of the American Statistical Association Section on Bayesian Statistical Science*, pages 6–10, 1999.
- P.R. Rosenbaum. Sensitivity analysis for certain permutation inferences in matched observational studies. *Biometrika*, 74(1):13–26, 1987. ISSN 0006-3444.
- P.R. Rosenbaum. Correction to: “Sensitivity analysis for certain permutation inferences in matched observational studies” [Biometrika 74 (1987), no. 1, 13–26; MR 88c:62114]. *Biometrika*, 75(2):396, 1988. ISSN 0006-3444.
- P.R. Rosenbaum. *Observational studies*. Springer-Verlag, New York, 1995. ISBN 0-387-94482-6.
- P.R. Rosenbaum. Observational studies and nonrandomized experiments. In *Design and analysis of experiments*, pages 181–197. North-Holland, Amsterdam, 1996.

- P.R. Rosenbaum and D.B. Rubin. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55, 1983. ISSN 0006-3444.
- P.R. Rosenbaum and D.B. Rubin. The bias due to incomplete matching. *Biometrics*, 41(1):103–116, 1985. ISSN 0006-341X.
- Donald B. Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976. With comments by R. J. A. Little and a reply by the author.
- I. Tager, T. Haight, Hollenberg, and Satariano. Physical functioning and mortality in elderly females. Unpublished manual, University of California, Berkeley, School of Public Health., 2000.
- M.J. van der Laan and J. M. Robins. *Unified methods for censored longitudinal data and causality*. Springer-Verlag, New York, To appear in 2002.

7 Appendix

Here we provide some Splus code to demonstrate the implementation of the proposed estimators with existing software tools.

7.1 IPAW estimators

Assume `mortality` is a dataframe with the following variables:

- ID: subject identifier
- Time: $k = 0, 1, \dots$, as in t_k
- BlCov: a baseline covariate, as in W
- LongCov: a time-dependent covariate, as in $L(t_k)$
- Trt: a time-dependent treatment process, as in A_k^t
- PastTrt: same as Trt, but with one unit of time lag
- TrtRand: indicator for treatment time
- CenPro: indicator for censoring
- Outcome: indicator for failure

Here is what the `mortality` dataframe might look like for two subjects:

ID	Time	BlCov	LongCov	Trt	PastTrt	TrtRand	CenPro	Outcome
1	0	5	0.5	0	0	1	0	0
1	1	5	0.5	0	0	0	0	0
1	2	5	0.5	0	0	0	0	0
1	3	5	0.5	0	0	0	0	0
1	4	5	0.5	0	0	0	0	0
1	5	5	0.6	1	0	1	0	0
1	6	5	0.6	1	1	0	1	0
2	0	12	1.0	0	0	1	0	0
2	1	12	1.0	0	0	0	0	0
2	2	12	1.0	0	0	0	0	0
2	3	12	1.0	0	0	0	0	0
2	4	12	1.0	0	0	0	0	0
2	5	12	1.5	0	0	1	0	0
2	6	12	1.5	0	0	0	0	0

2	7	12	1.5	0	0	0	0	0
2	8	12	1.5	0	0	0	0	0
2	9	12	1.5	0	0	0	0	1

Here is code to calculate the part of the stabilized weights arising from the treatment process:

```
treat.fit.big <-
  glm(Trt ~ PastTrt + LongCov + Time + BlCov,
      subset = (TrtRand == 1),
      family = binomial, data = mortality)

treat.fit.small <- update(treat.fit.big, . ~ . ~ - LongCov)

treat.prob.big <- rep(1, dim(mortality)[1])
treat.prob.small <- rep(1, dim(mortality)[1])

treat.prob.big[mortality$TrtRand == 1] <-
  ifelse(mortality$Trt[mortality$TrtRand == 1] == 0,
        1 - predict(treat.fit.big, type = "response"),
        predict(treat.fit.big, type = "response"))

treat.prob.small[mortality$TrtRand == 1] <-
  ifelse(mortality$Trt[mortality$TrtRand == 1] == 0,
        1 - predict(treat.fit.small, type = "response"),
        predict(treat.fit.small, type = "response"))

treat.weights <-
  unlist(tapply(treat.prob.small/treat.prob.big, mortality$ID,
              cumprod))
```

Here is the code to calculate the part of the stabilized weights arising from the censoring process:

```
sensor.fit <- glm(CenPro ~ LongCov + Time + BlCov + Trt,
                 family = binomial, data = mortality)

sensor.fit.small <- update(sensor.fit, . ~ . - LongCov)

sensor.weights.0 <- # censor weights at each subject time
  1-predict(sensor.fit, type = "response")

sensor.weights.0.small <-
  1-predict(sensor.fit.small, type = "response")

sensor.weights <- # time dependent weights
  unlist(tapply(sensor.weights.0.small/sensor.weights.0,
              mortality$ID, cumprod))
```

Here is the code to calculate the IPAW estimator:

```
iptw.fit <- glm(Outcome ~ Trt + BlCov + Time, family = binomial,
               weights = sensor.weights * treat.weights,
               data = mortality)
alpha.iptw <- iptw.fit$coef
```

7.2 TO estimators

To calculate a TO estimator, we must estimate the projection of IC_{iptw} on T_{SRA} . Therefore we will need \widehat{IC}_{iptw} for each subject; this is a vector of the same dimension as the causal parameter α , which

is 4 in this example. We store these as row vectors in the matrix $IcIptw$, in which \widehat{IC}_{iptw} is repeated as necessary within subject to facilitate the regression. Here is an example, where `design.matrix` is a `num.of.subj × num.of.par` matrix with one column for each covariate in the MSM:

```
# code to calculate IC_iptw
weights <- treat.weights * censor.weights
epsilon <- mortality$Outcome -
          1/(1 + exp(-1 * as.vector(design.matrix %*% alpha.iptw)))
ic.iptw <- weights * design.matrix * epsilon

IcIptw <- array(0, c(num.of.obs, num.of.par))
IcIptw.subj <- array(0, c(num.of.subj, num.of.par))
for(j in 1: num.of.par){
  IcIptw[, j] <- unlist(tapply(ic.iptw[, j], mortality$ID, rep.sum))
  IcIptw.subj[, j] <- unlist(tapply(ic.iptw[, j], mortality$ID, sum))
}

rep.sum <- function(x){
  rep(sum(x), length(x))
}
```

We also need to compute c_n and here is an example:

```
# code to calculate c_n

der.expit <- function(x){
  exp(x)/((1 + exp(x))^2)
}

temp1 <- der.expit(as.vector((design.matrix %*% alpha.iptw)) *
                  design.matrix)
temp2 <- weights * design.matrix

c.n <- array(0, c(num.of.par, num.of.par))
for(id in 1: num.of.subj){
  temp1.id <- temp1[mortality$ID = id, ]
  temp2.id <- temp2[mortality$ID = id, ]
  if(length(temp1.id) == num.of.par){ # if temp1.id is a vector
    c.n.id <- t(temp2.id) %*% temp1.id
  } else{
    c.n.id <- temp2.id %*% t(temp1.id)
  }
  c.n <- c.n + c.n.id
}
c.n <- c.n/num.of.subj
```

The estimated projection will be a 4-dimensional vector, with one vector per subject. We can now take this projection:

```
#Calculate the projection on Tangent space for censoring
proj.on.Tcar.1 <- array(0, c(num.of.subj, num.of.par)) # store the projection

covariates <- mortality[, c("Time", "Trt", "BlCov", "LongCov")]
glm.CenPro.fit <-
  glm(CenPro ~ Time + Trt + BlCov + LongCov,
      family = binomial, data = mortality)
```

```

epsilon <-
  mortality$CenPro-predict(glm.CenPro.fit, type = "response")

for(j in 1:num.of.par){
  # We use linear model to regress IC_{iptw} on the past. But reader
  # could do the regression more nonparametrically using gam
  # regress IC_{iptw} on past conditioning on $A^c(t)=1$.
  gam.fit.ic.iptw.1 <-
    gam(IcIptw[, j] ~ Time + Trt + BlCov + LongCov,
        subset = (CenPro == 1), data = mortality)
  # regress IC_{iptw} on past conditioning on $A^c(t)=0$.
  gam.fit.ic.iptw.2 <-
    gam(IcIptw[, j] ~ Time + Trt + BlCov + LongCov,
        subset = (CenPro == 0), data = mortality)
  temp <- predict.gam(gam.fit.ic.iptw.1, covariates) -
    predict.gam(gam.fit.ic.iptw.2, covariates)
  # The projection calculation is based on (8) type of representation
  proj.on.Tcar.1[, j] <- unlist(tapply(temp*epsilon, mortality$ID, sum))
}

# Calculate the projection on Tangent space for treatment
proj.on.Tcar.2 <- array(0, c(num.of.subj, num.of.par))
for(k in 1:num.of.par){
  temp <- rep(0, num.of.obs)
  gam.fit.ic.iptw.1 <- # Regress on (A_k, \mathcal{F}_k)
    gam(IcIptw[, k] ~ Time + Trt + PastTrt + BlCov + LongCov,
        subset = (TrtRand==1), data = mortality)
  gam.fit.ic.iptw.2 <- # Regress on \mathcal{F}_k
    gam(IcIptw[, k] ~ Time + PastTrt + BlCov + LongCov,
        subset = (TrtRand == 1), data = mortality)
  temp[mortality$TrtRand == 1] <-
    predict.gam(gam.fit.ic.iptw.1) - predict.gam(gam.fit.ic.iptw.2)
  proj.on.Tcar.2[, k] <- unlist(tapply(temp, mortality$ID, sum))
}

```

We can now calculate the TO estimator. Let $IcIptw$ now refer to the collection of 4-dimensional row vectors, each giving \widehat{IC}_{iptw} for one subject.

```

ic.to <- -(IcIptw.subj - proj.on.Tcar.1 -
  proj.on.Tcar.2) %*% solve(c.n) # c.n represents cn
mean.of.ic.to <- apply(ic.to, 2, mean)
alpha.to <- alpha.iptw + mean.of.ic.to # Calculate one step estimator
var.to <- var(ic.to)/num.of.subj # variance of TO estimator

```