

Linh Tran*, Constantin T. Yiannoutsos, Beverly S. Musick, Kara K. Wools-Kaloustian, Abraham Siika, Sylvester Kimaiyo, Mark J. van der Laan and Maya Petersen

Evaluating the Impact of a HIV Low-Risk Express Care Task-Shifting Program: A Case Study of the Targeted Learning Roadmap

DOI 10.1515/em-2016-0004

Abstract: In conducting studies on an exposure of interest, a systematic roadmap should be applied for translating causal questions into statistical analyses and interpreting the results. In this paper we describe an application of one such roadmap applied to estimating the joint effect of both time to availability of a nurse-based triage system (low risk express care (LREC)) and individual enrollment in the program among HIV patients in East Africa. Our study population is comprised of 16,513 subjects found eligible for this task-shifting program within 15 clinics in Kenya between 2006 and 2009, with each clinic starting the LREC program between 2007 and 2008. After discretizing follow-up into 90-day time intervals, we targeted the population mean counterfactual outcome (i. e. counterfactual probability of either dying or being lost to follow up) at up to 450 days after initial LREC eligibility under three fixed treatment interventions. These were (i) under no program availability during the entire follow-up, (ii) under immediate program availability at initial eligibility, but non-enrollment during the entire follow-up, and (iii) under immediate program availability and enrollment at initial eligibility. We further estimated the controlled direct effect of immediate program availability compared to no program availability, under a hypothetical intervention to prevent individual enrollment in the program. Targeted minimum loss-based estimation was used to estimate the mean outcome, while Super Learning was implemented to estimate the required nuisance parameters. Analyses were conducted with the ltmle R package; analysis code is available at an online repository as an R package. Results showed that at 450 days, the probability of in-care survival for subjects with immediate availability and enrollment was 0.93 (95 % CI: 0.91, 0.95) and 0.87 (95 % CI: 0.86, 0.87) for subjects with immediate availability never enrolling. For subjects without LREC availability, it was 0.91 (95 % CI: 0.90, 0.92). Immediate program availability without individual enrollment, compared to no program availability, was estimated to slightly albeit significantly decrease survival by 4% (95 % CI 0.03, 0.06, $p < 0.01$). Immediately availability and enrollment resulted in a 7% higher in-care survival compared to immediate availability with non-enrollment after 450 days (95 % CI $-0.08, -0.05$, $p < 0.01$). The results are consistent with a fairly small impact of both availability and enrollment in the LREC program on in-care survival.

Keywords: causal inference, causal estimation, targeted learning, causal road map, semiparametric models

*Corresponding author: **Linh Tran**, Department of Biostatistics, UC Berkeley, 101 Haviland Hall, Berkeley, CA 94720, USA, E-mail: tranlm@berkeley.edu

Constantin T. Yiannoutsos, Department of Biostatistics, Richard M. Fairbanks School of Public Health, Indiana University, Indianapolis, IN, USA

Beverly S. Musick, Department of Biostatistics, Indiana University School of Medicine, Indianapolis, IN, USA

Kara K. Wools-Kaloustian, Division of Infectious Diseases, Indiana University School of Medicine, Indianapolis, IN, USA

Abraham Siika, Sylvester Kimaiyo, Moi University School of Medicine, Eldoret, Central Kenya

Mark J. van der Laan, Department of Biostatistics, School of Public Health, UC Berkeley, 108 Haviland Hall, Berkeley, CA 94720-7360, USA

Maya Petersen, Department of Biostatistics, UC Berkeley, 101 Haviland Hall, Berkeley, CA 94720, USA

1 Introduction

van der Laan and Rose (2011), Petersen and van der Laan (2014), Pearl (2009) and others have advocated for the use of a systematic road map for translating causal questions into statistical analyses and interpreting the results. This road map requires the analyst to learn as much as possible about how the data were generated, posit a realistic statistical model that encompasses these findings, and assign a corresponding estimand that can answer the scientific question of interest. While in some cases this approach can be straightforward, in practice it generally requires much consideration when implementing. This is especially true in observational data, where a common objective is to estimate the joint effect of one or more longitudinal exposures, or a series of sequential treatment decisions (e. g. Bodnar 2004; Bryan 2004; Petersen et al. 2014). For example, one may be interested in contrasting immediate enrollment into a particular program (a single treatment decision) with delayed enrollment (a series of treatment decisions, in which enrollment is sequentially deferred at multiple time points at which it could have been initiated).

It is well recognized that the cumulative effect of such longitudinal exposures is often subject to time dependent confounding (e. g. Robins et al. 2000a; Bodnar 2004). For example, the decision to continue to defer enrollment at post-baseline time points may be affected by covariates that affect the outcome (and are thus confounders), and that are themselves affected by the prior decision not to enroll. If not properly accounted for, the resulting estimates may be biased for the parameter of interest.

Similar challenges arise in analyses aiming to estimate the effect of one longitudinal exposure while holding a second longitudinal exposure constant. For example, how would patient outcomes have differed under immediate versus deferred exposure to a clinic level intervention, if individual level enrollment were prevented? Such a controlled direct effect of two longitudinal exposures can be used to investigate the mechanisms by which an exposure is mediated by individual level uptake. In such scenarios, both longitudinal exposures may be subject to time dependent confounding.

Using the counterfactual framework, Robins (1986) and Robins et al. (1999) showed that both longitudinal effects and controlled direct effects can be identified in the presence of time dependent confounding. A range of estimators of such effects (or more precisely, estimators of the statistical estimands to which the counterfactual effects correspond under the sequential randomization assumption), have been developed, implemented, and applied (e. g. Robins 1987; Robins et al. 1999; Robins and Herná 2009; Hernán and Robins 2006; Pearl 2009). Prominent examples include inverse probability weighted (IPW) (Horvitz and Thompson 1952; Koul et al. 1981; Robins and Rotnitzky 1992; Satten and Datta 2001, 2004; Rotnitzky and Robins 2005; Hernán and Robins 2006), parametric g-computation (Robins et al. 1999; Robins and Hernán 2009; Hernán and Robins 2006; Taubman et al. 2009), and double robust estimating equation-based estimators (Robins et al. 2000b; van der Laan and Robins 2003; Bang and Robins 2005; Tsiatis 2006). More recently, Bang and Robins (2005) also introduced iterated conditional expectation g-computation and double robust substitution estimators, which van der Laan and Gruber (van der Laan and Gruber 2011) placed within the general targeted minimum loss-based estimation (TMLE) framework (van der Laan and Rubin 2006).

The TMLE (Bang and Robins 2005; van der Laan and Gruber 2011) offers a number of advantages, as have been reviewed elsewhere (van der Laan and Gruber 2011; van der Laan and Rose 2011). In particular, the TMLE does not suffer to the same extent from positivity violations (Hernán and Robins 2006; Westreich and Cole 2010) as IPW (Petersen et al. 2012), and allows for full integration of machine learning while retaining valid inference (van der Laan and Rose 2011; Brooks et al. 2013; Schnitzer et al. 2014; Decker et al. 2014). In this paper we describe an application of this methodology to estimate the joint effect of both time to program availability and individual enrollment in a task shifting program among HIV patients in East Africa. We provide a detailed description of this analysis, with an emphasis on several practical challenges likely to arise in similar applications. In doing so, we perform a case study of the targeted learning road map in the longitudinal setting (van der Laan and Rose 2011; Petersen and van der Laan 2014; Petersen 2014). In particular, we emphasize two primary issues. The first is the translation of the scientific questions

into counterfactual causal parameters, including a total effect and controlled direct effect. The hypothetical interventions for comparison are carefully defined to avoid known and potential positivity violations. The second is the use of recently developed estimation methods, including a double robust and semiparametric efficient targeted minimum loss-based estimator that integrates Super Learning for nuisance parameter estimation (van der Laan et al. 2007), and the imposition of global restraints for conditional distributions implied by our statistical model. We discuss several decisions and practical challenges that arise as a result, and provide corresponding R code.

The paper is organized as follows. In Section 2, we provide brief background on the scientific question of interest as well as review the targeted learning road map. In Section 3, we present the data as measured and formalized into a statistical framework. Section 4 posits the causal model for the data, while Section 5 addresses the target parameter, positivity violations, and identifiability. Section 6 focuses on estimating both the target and nuisance parameters. Section 7 presents the results of our approach, while Section 8 closes with a discussion and areas of future research.

2 Background

The majority of individuals with HIV live in settings, such as East Africa, where noticeable resource and infrastructure constraints place a limitation on the care these patients can receive (Stringer et al. 2006; World Health Organization 2013a). Antiretroviral therapy (ART) medication has been shown to reduce both viral loads and mortality in these patients (Gulick et al. 1997; Palella et al. 1998; Zwahlen et al. 2009; Dieffenbach and Fauci 2011; Danel et al. 2015; Lundgren et al. 2015), as well as reduce rates of transmission to persons uninfected (Quinn et al. 2000; Attia et al. 2009; Das et al. 2010; Cohen et al. 2011). However, the shortage of resources and health care professionals limit the number of patients who can be placed on ART (Van Damme et al. 2006; Koole and Colebunders 2010; World Health Organization 2013b; UNAIDS 2013).

Due to these limitations, various approaches have been undertaken in an effort to ensure that the maximal number of patients who need care can receive it. One such approach shifts care provision tasks for patients considered to be at low risk from physicians and clinical officers to other care professionals, such as nurses. Consequently, the workload for the physicians and clinical officers is reduced, theoretically increasing the attention that higher risk patients can receive for HIV or other conditions (World Health Organization 2013a).

One such program was implemented between 2007 and 2009 among clinics around eastern Africa. These clinics were followed as part of the Academic Model Providing Access to Healthcare (AMPATH) program, and contributed data to the International epidemiologic Databases to Evaluate AIDS, East Africa region (IeDEA-EA). The purpose of this Low Risk Express Care (LREC) program is to shift care tasks for patients considered to be at low risk from physician-centered care models to those utilizing non-physician health workers trained in simplified and standardized approaches to care. Patients were considered to be at low risk if they met the following set of criteria.

1. They were older than 18 years of age.
2. They were stable on ART for at least 6-months.
3. They had no AIDS-defining or AIDS-associated events within the past 180-days.
4. During past 6 months, they reported no missed pills when asked about the 7 days prior to each visit.
5. They were not pregnant within the past 30 days.
6. Their most recent CD4 count > 200 cells/ μL within 274 days prior to the current visit.

Once eligible, at each clinical visit, the clinician decided whether or not to enroll the patient into the LREC program. Patients enrolled had part of their care, such as identifying and managing ART side effects and opportunistic infections, shifted to nurses. Table 1 shows the differences between the standard of care and the LREC model.

Table 1: Proportion of visits responsibility assigned between the (a) standard and (b) LREC models of care provision.

Clinical monitoring	Standard model		LREC model	
	P/CO	Nurse	P/CO	Nurse
Request CD4/viral load measures	All	None	All	None
Monitor/support ART adherence	All	All	1/3	All
Determine functional status	All	None	1/3	2/3
Identify/manage ART side effects	All	None	1/3	2/3
Identify/manage opportunistic infections	All	None	1/3	2/3

Notes: P, Physician; CO, Clinical Officer.

In implementing this task-shifting program, a primary question of interest among health care providers is whether (i) clinic level exposure to and (ii) individual level enrollment in the program results in either better or worse clinical outcomes. For example, it could be that enrollment in the program increases loss to follow up and mortality because care is received from individuals who have a lower level of qualification/certification. Alternatively, enrollment in the program might decrease mortality and loss to follow-up due to more attentive and personal care given to enrolled patients. It could also be that an equivalent level of care is provided and thus, no impact is observed. Furthermore, receiving care at a clinic that has already implemented the LREC program might itself have a direct beneficial or detrimental effect on patient outcomes, in addition to an indirect effect mediated by patient enrollment in the program. Such an effect could result, for example, from changes in clinic practices made possible by shifts in workload.

2.1 Causal inference road map

We follow the targeted learning road map as presented by van der Laan and Rose (2011) and Petersen and van der Laan (2014). We briefly summarize the road map for longitudinal static interventions and corresponding notation below, and refer readers interested in further details to the cited articles.

Any causal scientific question must first be framed as a statistical question that can be answered with the observed data. Defining a statistical question includes properly defining the data, statistical model, and estimand or statistical parameter of interest. Here we consider data that consist of n independent and identically distributed observations of a random (vector) variable O with some underlying probability distribution P_0 . The statistical model \mathcal{M} should be chosen to ensure that it contains the true distribution P_0 , and thus any model assumptions should be based only on real knowledge about the process that generated the data. In practice this generally implies a semiparametric or nonparametric statistical model.

Translating causal questions into such a statistical estimation problem requires translating the scientific question into a formal causal query. This is typically defined as a parameter of the counterfactual distribution. We consider counterfactual experiments indexed by static interventions to set a vector \bar{A} of treatment or exposure variables equal to a set value \bar{a} , and denote the resulting counterfactual distribution $P_{\bar{a}}^0$. The statistical model \mathcal{M} can then be modified to incorporate non-testable causal assumptions about the data generating process. These causal assumptions are representable using structural causal diagrams or structural equation models (Pearl 2009). The resulting causal model $\mathcal{M}^{\mathcal{F}}$ is a model on $P_{\bar{a}}^0$ under each \bar{a} of interest.

One must determine what additional causal assumptions (if any) are needed in order to identify the target causal parameter from the distribution of the observed data. In this paper, we focus on the estimand provided by Bang and Robins (2005) through the iterated conditional expectation representation of the longitudinal g-formula. This parameter identifies the causal effects of longitudinal interventions under the sequential randomization and positivity assumptions.

Once identifiability has been established, the statistical model μ and target parameter of the observed data distribution P_0 to the real line are used to select a corresponding estimator $\hat{\Psi}$ which is a function of the empirical distribution of the observed data P_n . This process may additionally require nuisance parameter estimation. Here, we describe implementation of a double robust efficient targeted maximum likelihood estimator (Bang and Robins 2005). The resulting $\hat{\Psi}$ is used in calculating the parameter estimate $\hat{\Psi}(P_n) = \hat{\psi}_n$. Corresponding standard errors for the estimates are calculated based on the estimated influence curve and results are then interpreted. In the following sections we illustrate each of these steps in greater detail using the LREC analysis as a case study.

3 Data

Our study population is comprised of subjects found eligible for the LREC program within each of 15 clinics in Kenya between 2006 and 2009, with each clinic starting the LREC program between 2007 and 2008. Table 2 shows the characteristics of the 15 clinics.

Table 2: LREC program clinic characteristics (n = 15).

Area No.(%)		
Urban	6	(40.0)
Rural	9	(60.0)
Clinic Type No.(%)		
Referral hospital	3	(20.0)
(Sub) district hospital	7	(46.7)
Rural health center	5	(33.3)
Patients No.(%)		
≤ 500	4	(26.7)
501 – 1,000	3	(20.0)
1,001 – 1,500	4	(26.7)
> 1,500	4	(26.7)

In an effort to ensure that other unmeasured (or unmeasurable) aspects of care remained roughly constant, the first point of patient eligibility was truncated at 1-year prior to LREC program initiation at each clinic. Our target population is therefore comprised of patients found eligible for LREC within 1-year before the clinic's start date up to the administrative censoring date (5 March 2009). Our baseline timepoint (t=0) was defined to be the first date at which a patient was eligible for the program within this window. Figure 1 shows the distribution of the time from baseline to start of the LREC program. Subjects becoming eligible post LREC program initiation were assigned a time of 0. A small number of subjects had more than one year of follow up from baseline to LREC program initiation as a result of transferring to a new clinic with a later LREC start date. As patients generally visited clinics every three months, we discretized follow-up into time intervals of 90 days.

Patients were followed from the baseline time point defined above until one of four possible end points:

1. Death
2. Loss to follow-up (LTFU), defined here to be 198 days with no clinical visits
3. Database closure, occurring on 5 March 2009
4. Transfer to a clinic with no LREC program

For the present study, we were interested in the effect of LREC exposure and enrollment on the probability of remaining alive and in clinical care. In settings such as these, no standard national death

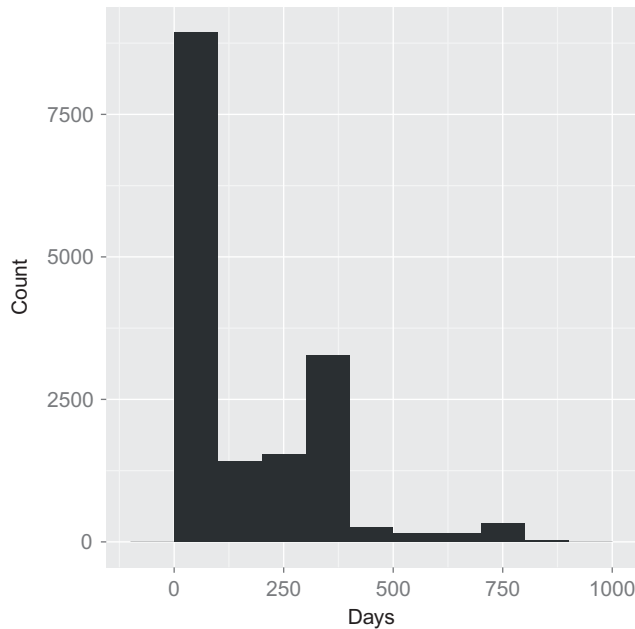


Figure 1: Histogram of time from eligibility to LREC program initiation.

registry exists. Consequently, relying solely on observed mortality would undercount the risk of death in our study population. Further, past research in this setting has found that a major reason (although certainly not the only reason) for apparent LTFU is unreported death (Geng et al. 2015, 2016). For example, a recent study by Geng et al. (2016) found that over 30% of subjects classified lost to follow up were unreported deaths. Without supplemental data, such unreported deaths preclude identification of causal parameters that treat death as the outcome and censor at LTFU, as well as parameters that treat LTFU as the outcome and censor at death (Geng et al. 2012). Furthermore, patients who do not return for continuing HIV care are subject to higher risk of complications and health decline (Kissinger et al. 1995; Samet et al. 2003; Giordano et al. 2005, 2007; Horstmann et al. 2010), placing them at unnecessarily higher mortality rates. We therefore define our outcome of interest as a composite of either the occurrence of death or LTFU. Patients were followed until this “failure” or until censoring due to either end of study or transfer. We aimed to evaluate the impact of (a) implementation of the LREC program at the clinic, and (b) enrollment into the LREC program after implementation on both retention “in-care” and survival.

3.1 Observed data

For notational convenience, we defined variables after the occurrence of any of these end points as equal to their last observed value. Following discretization of the data, we have a longitudinal data set with time-varying covariates, where the time points t correspond to 90-day intervals (e. g. 0, 90, 180,... days). Let W be the observed baseline time-independent covariates observed at the date the patient was first eligible for LREC (age at eligibility, CD4 count at start of ART, gender, indicator that ART regimen is protease-inhibitor based at eligibility, treated for tuberculosis at start of ART, indicator at urban clinic at eligibility, and WHO immunologic stage at both the start of ART and maximum stage prior to start). Let the time-varying variables from the observed data for each time point t be:

$$O(t) = (L1(t), Y(t), A1(t), A2(t), C1(t), C2(t)) : t = 0, 1, \dots, K, \quad [1]$$

where

- $L1(t)$ consists of the most recent measures of time-varying covariate values at the start of interval t , inclusive of covariates measured at the clinic level (i. e. calendar date, most recent, nadir, and zenith CD4 count, days since enrolling into the AMPATH program, an indicator of remaining on ART, pregnancy status, indicator in WHO stage III or IV, indicator of being treated for tuberculosis, clinic type (rural or urban), and an indicator of having at least one clinic visit in the previous interval).
- $Y(t)$ is an indicator that the patient is either (a) no longer in care (not seen in the clinic for 198 days) or (b) has died by the end of interval t . It jumps to and remains at 1 if either event occurs.
- $A1(t)$ is an indicator of the LREC program availability by the end of interval t . It jumps to and remains at 1 once the program has started.
- $A2(t)$ is an indicator of enrollment into the LREC program by the end of the interval. It jumps to and remains at 1 at time of enrollment and remains at 0 if $A1(t) = 0$.
- $C1(t)$ is an indicator that the patient transfers by the end of interval t to a clinic other than one of the 15 clinics that initiate the LREC program. It also jumps to and remains at 1 once the patient transfers.
- $C2(t)$ is an indicator of data base closure by the end of interval t . It jumps to and remains at 1 at the end of the study. Note that although database closure occurs at a fixed calendar date, censoring time due to database closure remains a random variable due to variability in time of eligibility for LREC.

We refer to the covariate and outcome nodes collectively as $L(t) = (L1(t), Y(t))$ and the treatment and censoring processes together as $A(t) = (A1(t), A2(t), C1(t), C2(t))$. Our observable data for each subject i can therefore be expressed as

$$O_i = (L_i(0), A_i(0), L_i(1), A_i(1), \dots, L_i(K+1)) \quad [2]$$

where $L(0)$ include our baseline variables W and $K+1$ is our final time point of interest, here equal to 4. We assume the observed data over all subjects consists of n copies of $O_i \stackrel{iid}{\sim} P_0$, where P_0 is the true underlying distribution.

4 Causal model

A causal model allows us to represent additional knowledge and assumptions associated with our scientific question that cannot be represented statistically. We present our causal model $\mathcal{M}^{\mathcal{F}}$ by making use of structural equation models to formally present how we assume each variable to be generated. We use a causal model that treats the 15 clinics as fixed, rather than sampled, and which describes an experiment in which individual subjects become eligible for LREC at random times. Specifically, we define the following non-parametric structural equation model (NPSEM) (Pearl 2009) to represent our knowledge about the causal process that generated the observed data.

$$\begin{aligned}
 L1(t) &= f_{L1(t)}(\bar{L}(t^-), \bar{A}(t^-), U_{L1(t)}) \\
 Y(t) &= f_{Y(t)}(\bar{L}(t^-), \bar{A}(t^-), U_{Y(t)}) \\
 &\quad \text{for } t = 0, 1, \dots, K, K+1 \\
 A1(t) &= f_{A1(t)}(\bar{L}(t), \bar{A}(t^-), U_{A1(t)}) \\
 A2(t) &= f_{A2(t)}(\bar{L}(t), \bar{A}(t^-), \bar{A}1(t), U_{A2(t)}) \\
 C1(t) &= f_{C1(t)}(\bar{L}(t), Y(t^-), \bar{A}(t^-), A1(t), A2(t), U_{C1(t)}) \\
 C2(t) &= f_{C2(t)}(L(0), Y(t^-), \bar{C}(t^-), C1(t), U_{C2(t)}) \\
 &\quad \text{for } t = 0, 1, \dots, K
 \end{aligned} \quad [3]$$

where $U \equiv (U_{L1(t)}, U_{Y(t)}, U_{A1(t)}, U_{A2(t)}, U_{C1(t)}, U_{C2(t)})$ are unmeasured exogenous variables from some underlying probability distribution P_U (assumed to satisfy the back door criterion (Pearl 2009) as detailed in Section 5.1.1) and $L(-1) = A(-1) = \emptyset$. For notational convenience, we define $t^- = t-1$ and define $\bar{X}(t) = (X(0), X(1), \dots, X(t))$ to denote the history of variable X up to time t .

This causal model specifies how we believe each of the variables in our data are deterministically generated, with randomness coming only from the unmeasured exogenous variables U . It tells us, for example, that individual enrollment immediately following LREC eligibility $A2(0)$ is generated as a deterministic function $f_{A2(0)}$ of $L(0)$, program availability $A1(0)$, and an error term $U_{A2(0)}$ drawn from some underlying population. Additionally, while not explicitly stated in eq. [3], we note that the deterministic function for enrollment $f_{A2(t)}$ sets $A2(t) = 0$ with probability 1 if $\bar{A1}(t) = 1$, i. e. if the program is not yet available.

Note that censoring due to end of study $C2(t)$ is not a function of time updated covariates $\bar{L}(t)$ beyond baseline covariates $L(0)$. This is because, while the values of $L(0)$ may vary due to the calendar date at which a subject's baseline eligibility occurs, once this date is set for a given subject, the censoring process due to database closure is deterministic.

Our outcome $Y(t)$ is assumed to be a function of treatment and censoring $A(t)$ only up to the previous time point, $t-1$. This restriction is imposed in order to preserve causal ordering, since death/LTFU ($Y(t)$) in an interval t clearly affects enrollment ($A(t)$) in the same interval. Consequently, the effects of availability and enrollment within an interval on the composite outcome are only captured beginning in the following interval.

5 Target parameter

In this framework, $t=1$ is the first time point at which treatment assignment $a(t)$ could affect survival. As the purpose of this study is to analyze the impact of different levels of treatment on the outcome, we conditioned on survival past $t=0$. Thus, we had that $Y(0) = 0$ for all subjects included in the study. That is, all subjects in the study survived past the first 90 days.

Conceptualizing an ideal hypothetical experiment can help in defining the target counterfactual parameter of interest. In order to evaluate the effect of exposure to and enrollment in the LREC program, we can conceive of an experiment in which we compare survival over time under alternative interventions to set time to program availability, time to enrollment following availability, and under an additional intervention to prevent censoring. As represented above in the causal model, these counterfactual outcome distributions are defined in terms of an intervention on the data-generating mechanism for $A(t) : t=0, 1, \dots, K$. In other words, we intervene to set values of program availability, enrollment, and censoring to some fixed values of interest at all time points.

Recall that our interest is evaluating both the total effect of exposure to the LREC program, and the direct effect of this exposure not mediated by individual level enrollment. Let $\bar{Y}_{\bar{a}}(t^*)$ denote the counterfactual outcome process up to time t^* under an intervention to set both time of program availability ($\bar{a}1$) and time of individual enrollment ($\bar{a}2$). Our target parameter for a given intervention of interest is $\mathbb{E}Y_{\bar{a}}(t^*) : t^* = 1, \dots, 4$, where $\mathbb{E}Y_{\bar{a}}(t^*)$ is the cumulative counterfactual probability of failure by time t^* under intervention \bar{a} . As we have conditioned on surviving past $t=0$ (the first 90-days of follow-up), our range represents the cumulative probability of failure between 91 and 450 days post-eligibility.

When contrasting counterfactual failure probabilities under distinct interventions, we focused on estimating the absolute risk difference (or average treatment effect). Specifically, we contrasted these counterfactual survival probabilities under the three following interventions: Our first intervention assigns all patients to have no program availability at all time points (set $A1(t) = A2(t) = 0 : t=1, 2, \dots, K$), and forces patients to remain uncensored (set $C1(t) = C2(t) = 0 : t=1, 2, \dots, K$). The corresponding 4 counterfactual survival probabilities $\mathbb{E}Y_{\bar{a}=0,0}(t^*) : t^* = 1, \dots, 4$ give us an understanding of survival patterns without the LREC program.

The second intervention of interest is to assign all patients to have immediate program availability (set $A1(t) = 1: t = 1, 2, \dots, K$), but not allow any subjects to enroll into the program (set $A2(t) = 0: t = 1, 2, \dots, K$). Patients would again be forced to remain uncensored and the counterfactual survival probability at each time point $\mathbb{E}Y_{\bar{a}=1,0}(t^*): t^* = 1, \dots, 4$ would be calculated. By evaluating $\Psi_{1,0}(P_0) = \mathbb{E}Y_{\bar{a}=1,0}(t^*) - \mathbb{E}Y_{\bar{a}=0,0}(t^*): t^* = 1, \dots, 4$, we target the controlled direct effect of exposure to the program if enrollment were prevented.

The third intervention is to assign all patients to have both immediate availability and enrollment (set $A1(t) = A2(t) = 1: t = 1, 2, \dots, K$). Again, censoring would be prevented and the counterfactual survival probability at each time point $\mathbb{E}Y_{\bar{a}=1,1}(t^*): t^* = 1, \dots, 4$ would be calculated. Evaluating $\Psi_{1,1}(P_0) = \mathbb{E}Y_{\bar{a}=1,1}(t^*) - \mathbb{E}Y_{\bar{a}=1,0}(t^*): t^* = 1, \dots, 4$ allows us to target the total effect of enrollment in a scenario where all subjects experienced immediate availability.

5.1 Identifiability

5.1.1 Sequential randomization

To establish identifiability of our causal parameter, we first make the assumption of sequential randomization. That is, we assume that

$$\bar{Y}_{\bar{a}}(t^*) \perp\!\!\!\perp A(t) | \bar{L}(t^-), \bar{A}(t^-) = \bar{a}(t^-): t^* = 1, \dots, 4; t = 0, \dots, t^* - 1. \quad [4]$$

Informally, we assume that the measured covariates are sufficient to control for confounding of our treatment effect. Consequently, at any time point t the observed treatment is independent of the full data given the data observed up to time point t . This assumption is met when the exogenous variables U are drawn from P_U such that the back door criterion is satisfied (Pearl 2009). Assuming that U are independent of each other and across time, for example, would satisfy the criterion. In this context, the major concern for violation of this assumption is that among patients classified as clinically stable, some patients are healthier or at lower risk of loss than others in ways not captured by the measured covariates, and these patients are differentially enrolled into the program.

5.1.2 Positivity

In addition, we assume that a subject had a positive probability of following each regime of interest (no availability, immediate availability and no enrollment, and immediate availability and enrollment) at each time point, regardless of their observed past:

$$P(A(t) = a(t) | \bar{L}(t^-), \bar{A}(t^-) = \bar{a}(t^-)) > 0 \text{ a.e.} \\ \text{for } t = 0, \dots, t^* - 1, \bar{a} \in \{0, 0 1, 0 1, 1\}. \quad [5]$$

Patients losing their eligibility for the LREC program posed a particular threat to this assumption. Our study population is comprised of patients initially deemed eligible for the LREC program due to their low risk. However, a noticeable proportion of the study population (32%) lost their eligibility at some point during follow-up. Once these patients were ineligible, they had a near 0 probability of subsequent program enrollment. To circumvent this potential positivity violation, we considered only treatment interventions which avoided enrollment at these time points. For example, consideration of patients who enroll immediately into the program would not encounter this issue, as all patients are eligible at the start of follow-up. Consideration of patients never enrolling into the program is also valid, as patients who are not eligible do not enroll. We further note that patients who lost their eligibility after enrollment into the LREC program were still considered to be enrolled. Similarly, patients who transferred to a new clinic without availability after receiving care at a clinic where LREC was available were considered exposed to the LREC program (in other words, we conducted an intent to treat type analysis of the effect of both availability and enrollment).

Under the assumptions of positivity and sequential randomization, the g-formula (Robins 1986, 1987) identifies the distribution that the observed data would have had under a counterfactual treatment, even when time-dependent confounders are present and are affected by prior components of the specified treatment. For our parameter, the standard g-computation representation for our target parameter is

$$\mathbb{E}Y_{\bar{a}}(t^*) = \sum_{\bar{l}(t^*-1)} \mathbb{E}(Y(t^*) \mid \bar{L}(t^*-1) = \bar{l}(t^*-1), \bar{A}(t^*-1) = \bar{a}(t^*-1)) \cdot \prod_{j=0}^{t^*-1} P(L(j) = l(j) \mid \bar{L}(j^-) = \bar{l}(j^-), \bar{A}(j^-) = \bar{a}(j^-)) \quad [6]$$

where the right hand side is a parameter of the observed data distribution P_0 . The g-computation recursive algorithm for estimating the cumulative probability of failure, introduced by Robins (2000) and expanded in Bang and Robins (2005), calls upon the tower rule of conditional expectations for this identity, suggesting an iterative conditional expectation (ICE) approach to estimating our parameters. Using their results, our parameter can therefore instead be expressed as

$$\mathbb{E}Y_{\bar{a}}(t^*) = \mathbb{E}[\mathbb{E}[\dots \mathbb{E}[\mathbb{E}[Y(t^*) \mid \bar{L}^{\bar{a}}(t^*-1)] \mid \bar{L}^{\bar{a}}(t^*-2)] \dots \mid \bar{L}^{\bar{a}}(0)]] \quad [7]$$

where $\bar{L}^{\bar{a}}(t)$ is the variable $L(t)$ from the post intervention distribution resulting from setting $\bar{A}(t^-) = \bar{a}(t^-)$ and the expectations are taken with respect to this distribution. Thus, given $\bar{L}^{\bar{a}}(t)$ is equivalent to conditioning on $\bar{L}(t), \bar{A}(t) = \bar{a}(t)$.

6 Estimation

The ICE approach has a number of advantages compared to the more simple standard g-computation procedure. The most noticeable among them is that we are only required to estimate the iteratively defined conditional expectations as opposed to the entire conditional densities of the covariate process $L(t)$. We therefore use the ICE approach here. A small disadvantage is that these set of regressions must be run separately for each desired treatment rule, \bar{a} , whereas the classical parametric g computation approach only has to estimate the conditional densities once. We note however that this is a very small price to pay when compared against the substantial gain achieved by not having to estimate the entire joint density, especially when dealing with high dimensional data.

While the ICE approach already provides a considerable advantage towards our estimation goals, using targeted minimum loss-based estimation (van der Laan and Gruber 2011; Bang and Robins 2005) provides a further gain. This approach provides a substitution estimator that solves the estimating equation corresponding to the efficient influence function $D(P)$ for our estimand. This removes the bias associated with the untargeted minimization of a global loss function for the density of the data. It is also known to be double robust, in that consistent estimation of the target parameter is obtained if either of the treatment mechanism or the iterated conditional expectations are estimated consistently. This approach is guaranteed to respect the parameter constraints implied by our statistical model. An R package titled `ltmle` has been developed which implements this estimator (Lendle et al. 2016; Petersen et al. 2013). This package takes longitudinal data supplied in wide format and estimates our target parameter for each specified treatment rule \bar{a} .

A number of options are available to users of the `ltmle` R package. For example, the probabilities of treatment and censoring $A(t)$ at each time point (for the specified treatment rule \bar{a}) can be separately estimated and subsequently fed into the estimation procedure, rather than being fit within. This allows the analyst the option of pooling the observations over all time points and estimating probabilities within this pooled sample, as opposed to fitting separate models for each time point. Doing so provides a lower variance in estimates of the treatment mechanisms at the cost of potentially higher bias.

An additional advantage of pooling over all observations in estimating our treatment and censoring mechanisms is that we can use additional data that is not included in modeling the ICEs. That is, data

observed beyond the final time point t^* can be used to aid in estimating the probabilities of treatment for $A(t): t=0, 1, \dots, t^*-1$. This can be advantageous and promotes stability in the estimates by borrowing information across all time points. Specifically, it helped in the present study in estimating our censoring from transfer mechanism, due to the extremely small number of transfers observed. For this study, we choose to pool across observations to fit the treatment mechanisms, though note that the decision between the two did not significantly affect the parameter estimates.

An additional package option is the ability to pool over all subjects when estimating the ICEs irrespective of observed treatment, as opposed to stratifying the data and using only subjects following the treatment rule specified. This choice implies an analogous bias-variance trade off. Pooling over subjects regardless of treatment history potentially allows for more stable estimates of the ICEs due to the lower variance. This option is helpful when the number of time points is large and the number of persons following a particular treatment regime over all time points is small.

Use of the `ltmle` R package requires that the data be provided in a wide format and with a time-ordering of the covariates. In doing so, two options are available for the $L(t)$ node at each time point t , i. e. $L1(t), Y(t)$ or $Y(t), L1(t)$. In our causal model in Section 4, the $L(t)$ node is not affected by the specified order. Therefore, use of either ordering will suffice in our study. We additionally note that even if the time-ordering did matter and our outcome of interest $Y(t^*)$ were, say, dependent upon the covariates $L1(t^*)$, there is still no need to condition on $L1(K+1)$ for the sake of estimating our target parameter as the efficient influence function $D(P)$ for our parameter based on the full data O is the same as the efficient influence function based on reduced data $O_r = O/L1(K+1)$. We provide a short proof for this in the Appendix.

6.1 Super learning the nuisance parameters

Consistent, asymptotically linear, and efficient estimation of our target parameter requires that the treatment and censoring $A(t)$ mechanisms as well as ICEs be estimated in a consistent manner and at a fast enough rate (van der Laan and Rose 2011). Using parametric models to do this requires correct specification of the functional form for the treatment mechanism and iterated conditional expectations. Given that we do not know a priori the form of the true probability distribution P_0 and the extreme unlikeliness that a simple parametric specification will result in a correctly specified model, use of these models will most likely result in overly biased estimates which will approach statistical significance with probability 1 as sample size increases regardless of whether a treatment effect exists. In other words, the use of misspecified parametric models elevates the risk of obtaining significant findings even if no true treatment effect is present.

An alternative approach is to use data-adaptive non-parametric methods which reside in a much larger statistical model or set of distributions. Examples include gradient boosting machines (Friedman 2001, 2002), neural networks (Mc-Culloch and Pitts 1943), and k-nearest neighbors (Altman 1992). In deciding which method to use, we recommend using the ensemble machine learning approach Super Learner, which is based on V-fold cross validation and implemented in the R package titled SuperLearner (Polley and van der Laan 2014). This algorithm takes a user-supplied loss function (chosen to measure performance) and a library of algorithms, which can include parametric models as well as non-parametric or machine learning algorithms such as those listed above. It uses V-fold cross validation to chose the convex combination of these algorithms that performs best on independent data (derived from internal data splits). If, as is likely, none of the algorithms in the library achieves the rate of convergence one would have with a correctly specified parametric model, the Super Learner will still perform asymptotically at least as well as the best algorithm in the library. Otherwise, it achieves an (almost) parametric rate of convergence. Furthermore, the derived oracle inequality showing the asymptotic performance also shows that the number of candidate algorithms considered in the cross-validation procedure can be polynomial in size proportional to the number of observations (van der Laan and Dudoit 2003; van der Vaart et al. 2006; van der Laan et al. 2007). Therefore, a large number of algorithms can be considered, which can grow with the number of observations, without fear of hampering the Super Learner's performance.

To use Super Learning, a loss function must be chosen and a user-specified library provided. We chose the non-negative log loss function for its desired steeper risk surface, as this function penalizes incorrect probabilities more severely than the more commonly used squared error loss. A number of default candidates are included in the SuperLearner package that were used here. These include the overall mean, main terms logistic model, step-wise regression with AIC (Hoerl and Kennard 1970), generalized additive model (Hastie and Tibshirani 1986), Bayesian generalized linear model (Gelman et al. 2008), k-nearest neighbors (Altman 1992), LASSO (Tibshirani 1996), ridge regression (Hoerl and Kennard 1970), neural net (McCulloch and Pitts 1943), multivariate adaptive polynomial spline (Friedman 1991), generalized boosted regression model (Friedman 2001, 2002), and support vector machine (Boser et al. 1992; Cortes and Vapnik 1995). Additionally, most of the algorithms have tuning parameters which can result in better candidate performance. To ensure that we were achieving satisfactory performance, we used different tuning parameters as additional candidates in the ensemble for the generalized additive models, k-nearest neighbors, neural nets, and generalized boosted regression models. We additionally used 4 user-specific parametric models as candidates in the library. The Super Learner fits were constructed using all potential confounders, listed above in Section 3.1 as W and $L1(t)$ for each time point t .

Of particular concern to the analyst when deciding on which candidates to include in the Super Learner library is the explicit condition that the candidates not be too data adaptive. This is because the empirical process conditions for the asymptotic linearity of our estimator require that we work within a Donsker class of estimators (van der Laan et al. 2009). Indeed, we have seen that algorithms that tend to overfit the empirical data, such as the machine learning algorithm random forest, will negatively impact our estimators. We therefore excluded these algorithms from our library, though note that such algorithms could still be used in a cross-validated parameter estimation approach such as cross-validated targeted minimum loss-based estimation (Zheng and van der Laan 2010).

Regarding the pooling of observations across time for the treatment mechanism, one further possible option is to use a Super Learner library that is doubled by including estimates from both the time stratified and pooled approach. Ensemble weights could then be calculated based on the best performing candidate in this larger library and subsequently fed into the ltmle package. Consequently, we continue benefiting from the borrowed information at different time points and simultaneously protect ourselves from the asymptotic bias of the previous approach. We opted not to additionally use the stratified approach, due to the computational intensity required of the approach.

6.1.1 Initial ICE fits

In using non-parametric estimators for the estimation of the ICEs, we may potentially disregard the global constraints implied by our statistical model. While all the estimators considered here are expected to work well at time $t = t^*$ where the outcome being modelled is binary, their use at $t < t^*$ where the outcome being modelled is known to fall within the interval $[0, 1]$ can present issues. For example, continuing to treat the outcome being modelled as binary may result in programmatic errors since many of the algorithms, such as Support Vector Machines or k-nearest neighbors, require that the outcome be supplied in classes or as factors. Alternatively, modelling the outcome continuously may result in extrapolations with estimates greater than 1 or less than 0. As our expectation is known to fall within $[0, 1]$, this would result in violations of the constraints of the outcome being modelled.

For each of the algorithms facing this potential issue, we implemented three approaches aimed at ensuring that estimates remained within the constraints. All three were used in the Super Learner library, allowing us to objectively compare the performance of each approach. We define each of the conditional expectations from eq. [7] to be $\bar{Q}_{0,L(t)}$ such that, for example, $\bar{Q}_{0,L(t^*)} \equiv \mathbb{E}[Y(t^*)|\bar{L}^a(t^* - 1)]$ and $\bar{Q}_{0,L(t^*-1)} \equiv \mathbb{E}[\bar{Q}_{0,L(t^*)}|\bar{L}^a(t^* - 2)]$. Let each estimator of the conditional expectation within the Super Learner

library at time t be denoted as $\bar{Q}_{n,L(t)}^j$ for $j \in 1, 2, \dots, J$ where J is the total number of candidates in the Super Learner library. We considered

1. Simply truncating $\bar{Q}_{n,L(t)}^j$ at both 0 and 1.
2. Taking the logit transformation of the outcome being modelled and truncating at a fixed threshold τ (set here to be 0.0001). We then modelled the transformed outcome on a continuous scale and took the inverse logit transformation on the fitted values.
3. Stratifying the observations by whether they were within the $(0, 1)$ open interval or within $\{0, 1\}$, i. e. whether they were continuous within the $(0, 1)$ interval or dichotomous with only values of 0 or 1. The former were fit on a continuous scale after taking the logit transformation, while the latter were modelled as a binary outcome.

We emphasize that use of Super Learner for estimation of the treatment mechanisms and ICEs provides two important primary benefits. Firstly, its use helps ensure the conditions for the asymptotic linearity of our estimator and the corresponding statistical inference are met by ensuring the consistent estimation of both the intervention mechanism and the iteratively defined conditional expectations. This allows us to establish robust confidence intervals for our estimator. Secondly, we gain efficiency in that we get an asymptotically efficient estimator if both the treatment mechanisms and ICEs are estimated consistently at fast enough rate. Thus, as long as at least one of the library candidates for each of the nuisance parameters achieve this optimal rate, our approach will have the lowest variance among all regular asymptotically linear estimators. Further, even if we fall short of this goal, the improved estimation of both nuisance parameters offered by Super Learner will generally improve finite sample efficiency.

7 Results

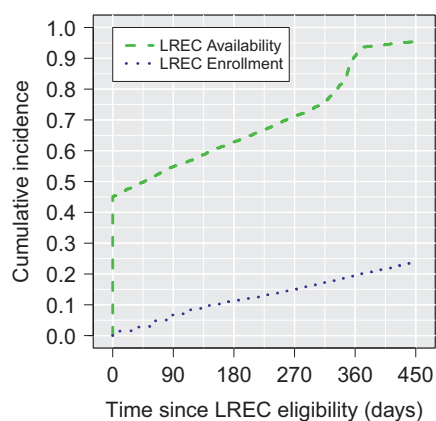
Among the 15 clinics implementing the LREC program, a total of 16,513 subjects (31% male) were found to be eligible for the program, of which 16,479 survived past $t=0$. As we only modelled survival up to 450 days, we report figures with follow-up truncated at that point. After discretizing the data, these patients contributed a total of 17,668 person-years of follow-up to the analyses. From these subjects 1,206 failure events were observed, of which 1,102 were losses to follow-up and 104 were deaths. All failure events observed at $t=1$ were deaths, since our definition of loss to follow-up required at least 198 days to pass before a subject could be lost to follow-up. A small number of subjects ($n=128$) were censored due to transfers to non-LREC clinics, while 3,889 subjects reached the end of study prior to $t=4$ and prior to experiencing a failure event and were administratively censored. Table 3 shows the characteristics of the patients conditioning on survival past $t=0$.

A small proportion of subjects died (42), were lost to follow-up (286), or were censored (60) before the LREC program became available. Of the 16,050 subjects who were at some point exposed to the program, most (15,294) experienced it by 1-year from baseline. Almost half of the study population began follow-up after the LREC program had already initiated, as indicated by the large spike in the cumulative incidence at time 0 (Figure 2). A noticeable spike was also seen at 1-year after baseline, representing the patients who had their first eligibility truncated at 1-year as stated in Section 3.

Patients who were not exposed to the LREC program could not enroll. Furthermore, once the LREC program was available, decisions on whether to enroll subjects rested upon the treating clinicians or clinical officers. Consequently, only 3,832 subjects were enrolled. As expected, subjects who were healthier were more likely to enroll into the LREC program. For example, univariate analyses showed that subjects who had higher CD4 counts, were receiving ARV, had a WHO stage I or II, were seen in clinics less often, and were not being treated for tuberculosis had higher probabilities of enrolling. Additionally, subjects from (sub) district hospitals and rural health centers (compared to referral hospitals), with fewer missed

Table 3: Characteristics of 16,479 patients at LREC eligibility (conditioning on survival past $t = 0$).

	Count	(%)
Age (years)		
< 30	2,793	(17 %)
30–39	7,017	(43 %)
≥ 40	6,669	(40 %)
Sex		
Female	11,441	(69 %)
Male	5,038	(31 %)
CD4 cell count (cells/μL) at ART start		
< 200	9,754	(59 %)
200–349	2,313	(14 %)
350–499	543	(3 %)
≥ 500	387	(2 %)
Unknown	3,482	(21 %)
CD4 cell count (cells/μL) at eligibility		
< 200	0	(0 %)
200–349	10,125	(61 %)
350–499	4,016	(24 %)
≥ 500	2,334	(14 %)
Unknown	4	(0 %)
PI-based ARV regimen		
No	15,600	(95 %)
Yes	879	(5 %)
Max WHO stage prior to ARV start		
I/II	7,696	(47 %)
III/IV	8,373	(51 %)
Unknown	410	(2 %)

**Figure 2:** Cumulative incidence of LREC availability and enrollment.

clinical visits, not on protease-inhibitor based regimen, and who were not pregnant also had higher probabilities of enrolling. We note that despite listing non-pregnancy as a criteria for being at low risk, a small number of subjects (18) enrolled while pregnant. Figure 2 shows the cumulative incidence of LREC availability and enrollment. Cumulative incidence of enrollment by 90 and 360 days after baseline was 7% (95% CI: 6.3%, 7.1%) and 19% (95% CI: 18.8%, 20.1%), respectively.

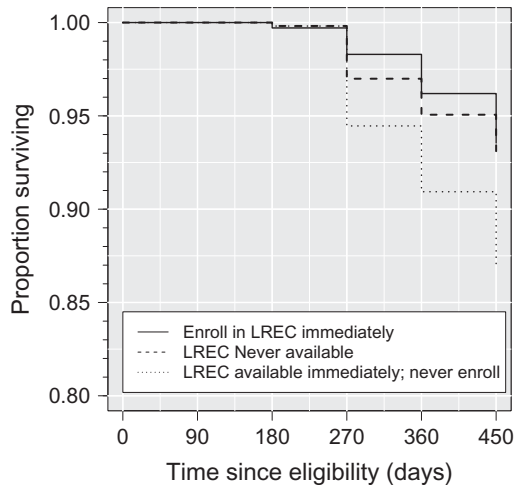


Figure 3: Unadjusted Kaplan-Meier survival curves.

As stated in Section 3, all patients in our study started follow-up eligible for the LREC program, leading to low or no variance in many of the confounders at early time points with a skewness towards the healthier values. During follow-up, however, many subjects who did not enroll subsequently became less healthy resulting in decreased probabilities of subsequent enrollment. Indeed, 3,920 subjects were found to have lost their eligibility prior to enrollment and prior to 1-year, precluding interventions to evaluate a range of different enrollment times, as discussed in Section 5.1.2.

Unadjusted analyses using the Kaplan-meier estimator showed overall high probabilities of in-care survival among all subjects (Figure 3). Those with immediate LREC availability who never enrolled had noticeably lower in-care survival than subjects never experiencing LREC availability. For example, at $t = 4$ the proportion of subjects still alive and in care was 0.77 for those not enrolling into LREC, compared to 0.93 for the group of subjects never experiencing LREC availability and 0.94 for subjects enrolling into LREC immediately. Conversely, those with immediate enrollment into the program had the highest survival probabilities. Differences in survival probabilities between treatment groups increased with time, with the largest differences seen between subjects with immediate enrollment and those with LREC availability never enrolling.

The cross-validated risks (using the non-negative log likelihood loss) for the treatment and censoring mechanisms are shown in Figure 4 under various models and algorithms. While observations at all time points were used in the mechanism fits, our interest is only in treatment interventions at $t \leq 3$. We therefore only calculated the risks using those time points. As expected, the Super Learner fit outperformed all of the candidate estimators in the supplied library, as well as the cross-validation selector (i. e. discrete Super Learner, which is equivalent to choosing the single algorithm in the library with the lowest cross validated risk). Compared to the mean model, which assumes no confounding and does not control for any confounders, the Super Learner fits for the LREC availability and end of study mechanisms showed an immense decrease in cross-validated risk. This gain was also noticeable when compared to the candidate model that only controls for time. The Super Learner fit for the enrollment mechanism also outperformed the mean model, though to a smaller degree. No noticeable gain was seen in the transfer mechanism, presumably due to the extremely low number of transfers observed (218). We did not present the cross-validated risks for the ICE fits as they are too numerous to describe in detail, though note that the gains from using Super Learner were similar to fits for the treatment and censoring mechanisms.

Adjustment for potential confounders using the Super Learner fits resulted in relatively small updates of the survival curves (Figure 5). Subjects enrolling immediately into the LREC program at eligibility

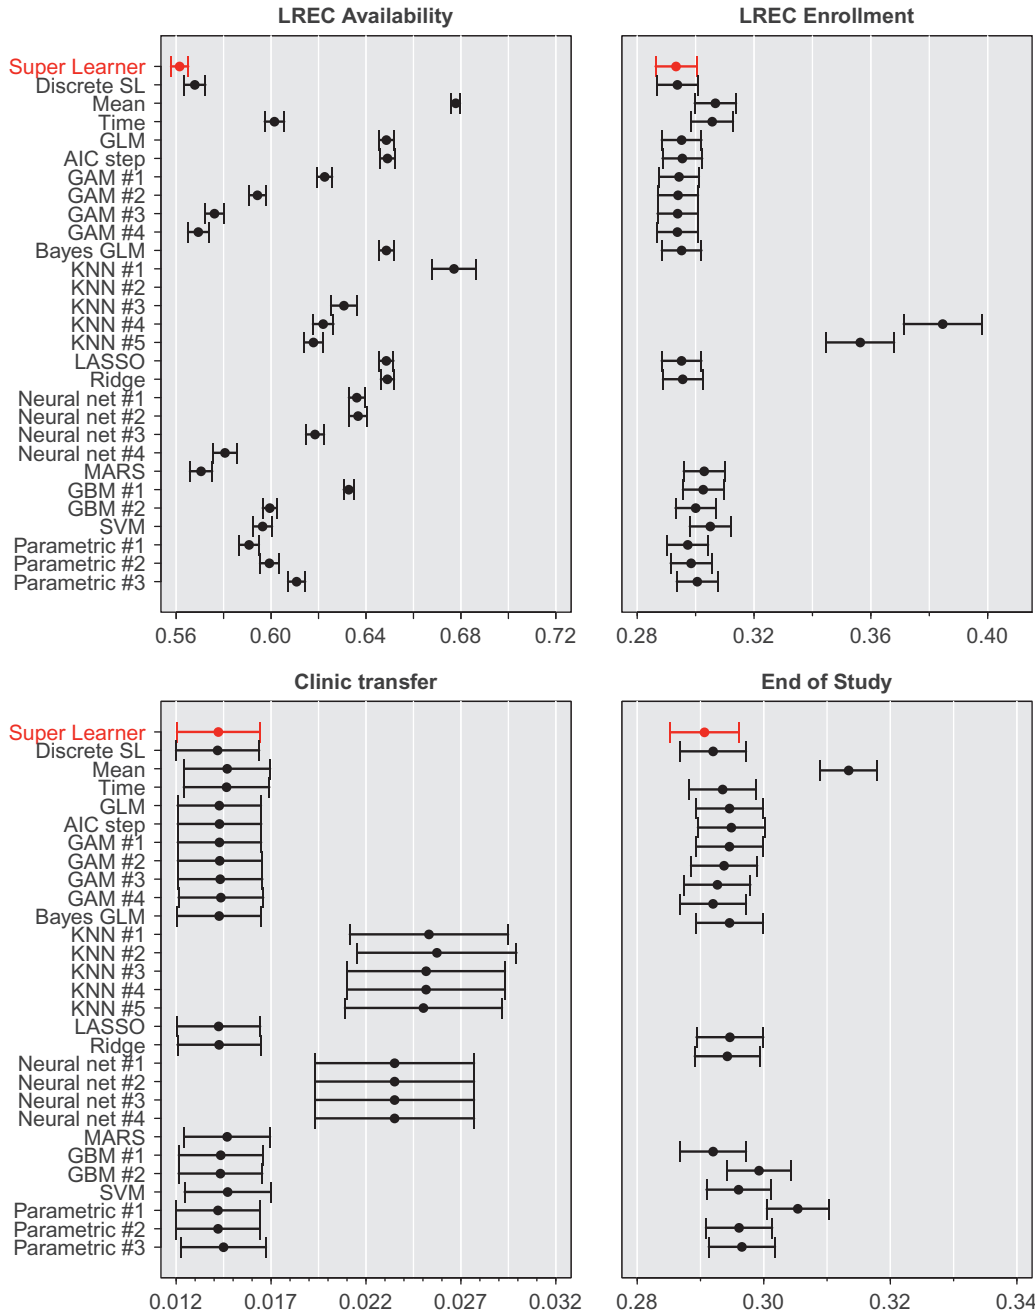


Figure 4: Cross-validated risk estimates (using non-negative log-likelihood loss) and 95 % confidence interval for the treatment and censoring mechanisms. A number of candidates had cross-validated risks too high to plot in the specified window and consequently are not shown. Mean: marginal probability, Time: logistic regression with time variable only, GLM: logistic regression with all confounders, AIC step: stepwise regression using the Akaike information criterion, GAM: generalized additive model, KNN: k-nearest neighbors, LASSO: least absolute shrinkage and selection operator, MARS: multivariate adaptive regression splines, GBM: generalized boosted regression models, SVM: support vector machines, Parametric: user specified logistic models using a subset of the confounders.

continued to have the highest survival probabilities, while those with immediate availability not enrolling had the lowest. Tables 4 and 5 show the calculated average treatment effects between the different interventions. Confidence intervals and p-values were calculated based on influence functions. As implied by the survival curves, immediate enrollment into the LREC program at eligibility had a beneficial effect

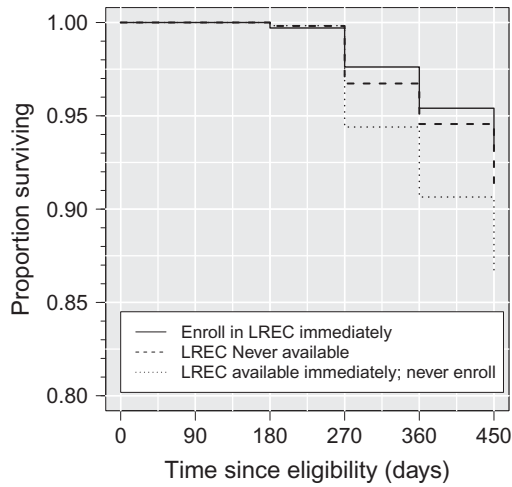


Figure 5: Survival curves adjusting for potential confounders.

Table 4: Unadjusted time-specific average treatment effects. (a) compares the intervention immediate LREC availability without enrollment to never having LREC available; (b) compares the intervention immediate LREC availability and enrollment to immediate LREC availability without enrollment.

Time (t^*)	(a) $EY_{\bar{a}=10}(t^*) - EY_{\bar{a}=00}(t^*)$			(b) $EY_{\bar{a}=11}(t^*) - EY_{\bar{a}=10}(t^*)$		
	Estimate	(95 % CI)	p-value	Estimate	(95 % CI)	p-value
1	0.00	(0.00,0.00)	0.93	0.00	(0.00,0.00)	0.56
	0.03	(0.02,0.03)	0.00	-0.04	(-0.05,-0.03)	0.00
	0.04	(0.03,0.05)	0.00	-0.05	(-0.07,-0.04)	0.00
	0.06	(0.05,0.08)	0.00	-0.07	(-0.09,-0.05)	0.00

Table 5: Time-specific average treatment effects adjusted for measured potential confounders. (a) compares the intervention immediate LREC availability without enrollment to never having LREC available; (b) compares the intervention immediate LREC availability and enrollment to immediate LREC availability without enrollment.

Time (t^*)	(a) $EY_{\bar{a}=10}(t^*) - EY_{\bar{a}=00}(t^*)$			(b) $EY_{\bar{a}=11}(t^*) - EY_{\bar{a}=10}(t^*)$		
	Estimate	(95 % CI)	p-value	Estimate	(95 % CI)	p-value
1	0.00	(0.00,0.00)	0.81	0.00	(0.00,0.01)	0.58
	0.02	(0.02,0.03)	0.00	-0.03	(-0.05,-0.02)	0.00
	0.04	(0.03,0.05)	0.00	-0.05	(-0.07,-0.03)	0.00
	0.04	(0.03,0.06)	0.00	-0.07	(-0.08,-0.05)	0.00

relative to never having LREC available, while having LREC immediately available and never enrolling was adverse. For example, at $t=4$ the probability of survival for subjects with immediate enrollment was 0.93 (95 % CI: 0.91, 0.95) and 0.87 (95 % CI: 0.86, 0.87) for subjects with immediate availability never enrolling. For subjects without LREC availability, it was 0.91 (95 % CI: 0.90, 0.92). Similar to the unadjusted estimates, the treatment effects increased with time. All estimates after $t=1$ showed statistical significance.

It is possible that near positivity violations can have large effects on estimates of our parameter. To test for this potential issue, we considered different truncation bounds for our treatment probabilities. Specifically, we considered using a bound of 0.001 and using untruncated probabilities. No differences were seen in the resulting mean outcome estimates.

8 Discussion

We have presented a comprehensive approach to applying longitudinal targeted minimum loss-based estimation to evaluate the impact of the LREC program. Corresponding code for the analyses have been uploaded to an online public repository as an R-package at www.github.com/tranlm/lrecImpact. The results support a somewhat negligible impact of implementation and enrollment, with the lowest survival among patients with immediate LREC availability never enrolling and similar survival among the other two interventions (Figure 5). Subjects enrolling immediately into the LREC program have almost identical survival to subjects never being exposed to the program. While the magnitude of difference in survival increased with time, this difference is modest.

We chose 90-day intervals for our time points in the current study, due to the understanding that patients would have visits approximately every 3-months. While smaller intervals could have been chosen, doing so can reduce the probability of following a given regime of interest given observed covariates (i. e. increase the extent of practical positivity violations), both by decreasing the probability of availability and enrollment occurring in the first interval, and because the probability of never enrolling given observed covariates involves taking a cumulative probability of not enrolling given the observed past over many more time points. Furthermore, the use of smaller intervals results in more time points, leading to higher computational costs. On the other hand, the use of larger intervals leads to discarding information in order to preserve time ordering, which can result in a less complete control for confounding as well as failure to capture the full causal effect of the intervention. In order to preserve time ordering, only covariate and outcome values measured at the end of the prior interval are considered possible causes of enrollment and availability in an interval. Longer intervals result in more problems with the assumption. We tested whether there was an effect in our study by re-running the analyses using 30-day intervals. The resulting survival estimates were similar to the ones reported here.

As with all studies, there are limitations that need to be considered. Firstly, it is possible that we did not sufficiently adjust for all the potential confounders. For example, the majority of subjects who had immediate availability and never enrolled had initial eligibility occur after the LREC program had already started. These subjects experiencing incidental eligibility (as opposed to prevalent eligibility from those eligible prior to the LREC program initiation) may have had factors placing them at higher risk. In addition, in defining our composite outcome “dead or lost to follow up” we implicitly assumed that not being seen in clinic for 198 days is an undesirable outcome reflecting out of care status. In practice, some of these patients might represent unreported transfers to care in an alternative clinic. Lastly, our analysis considered subjects from the same clinics to be causally independent of each other. In specifying our causal model we made a key decision to use an individual level NPSEM despite our interest in both an individual and a clinic level exposure variable. Such a formulation assumes that individuals within a given clinic are causally independent, and in particular, that the exposure received by one patient does not impact the outcome of another (the assumption of no causal interference) (Kline 2011; Tchetgen and VanderWeele 2012). A different formulation is possible that uses a hierarchical or clinic level NPSEM and corresponding hierarchical identification and analysis. We can think of the corresponding experiment as randomizing entire clinics to start the LREC program and within clinics with LREC available, randomizing patients to enroll. However, the sample size then becomes driven by the number of clinics and identification would require adequate variability in the introduction of LREC across clinics (Tchetgen and VanderWeele 2012). We therefore pursued an individual level formulation, while noting the limitations of this approach. Future research into improved approaches to interference effects in this setting should be undertaken.

We end by stating that, while not conducted here, this framework can be easily generalized to include dynamic interventions that are dependent upon other covariates. For example, there could be interest in intervening to enforce enrollment only on patients who retain eligibility during follow-up while exempting

patients who do not. Another option to consider is the use of marginal structural models to smooth treatment effects across time points, as well as availability and enrollment times (Robins 1999a, 1999b; Robins et al. 2000a; Petersen et al. 2013) though care should be taken when implementing as the number of regimes with available data would be limited. These models allow us to project the true underlying dose response curve onto a working parametric model, allowing us to conduct inference on a smaller set of parameters. The `ltmle` package includes a TMLE for causal parameters defined as the projection of the survival or failure curve onto user-specified parametric models.

In summary we applied the targeted learning roadmap to longitudinal data with a multilevel longitudinal treatment of interest to analyze a nurse-based triage system among HIV patients in East Africa. This included both definition and identification of our causal parameter. Issues with positivity were handled with careful selection of our target causal parameter. Nuisance parameters were estimated using Super Learner, a cross-validation ensemble algorithm using both parametric and machine learning algorithms. Observations for the estimation of the treatment mechanisms were pooled across time points, which aided us in estimating the censoring mechanism due to clinical transfers. Various approaches were implemented aimed at ensuring the machine learning estimates of the ICEs would respect the underlying statistical model. Estimates of survival at each time point were then contrasted by their differences and inference derived using the empirical influence functions. The results show a somewhat negligible impact of both availability and enrollment in the LREC program on in-care survival.

Appendix. Proof that O and O_r have equivalent influence functions

We provide a proof here showing that the influence function for the full data O is equivalent to the influence function we solve for in our analysis based on the reduced data $O_r = O/L1(K+1)$.

Firstly, we know the efficient influence function $D(P)$ for the reduced data O_r . Our goal is to compute the efficient influence function for the full data O for this same parameter (but now as a function on a model on O instead of O_r).

Recall that one can compute an efficient influence function $D^*(P)$ by first deriving an influence function of any estimator, and then projecting this influence function on the tangent space $T(P)$ of the model. $D(P)$ is one such influence function since it is the influence function of the TMLE ignoring $L1(K+1)$. Thus, the desired efficient influence function is $D^*(P) = \pi(D(P)|T(P))$ where π is the projection operator acting onto $T(P)$.

Now, we note that the likelihood of O can be factorized as $P(L1(K+1)|Y(K+1))P(O_r)$. Thus, the tangent space is the orthogonal sum of the tangent space $T_1(P)$ of the first factor $P(L1(K+1)|Y(K+1))$ and the tangent space $T_r(P)$ of $P(O_r)$. Consequently, $\pi(D(P)|T(P)) = \pi(D(P)|T_1(P)) + \pi(D(P)|T_r(P))$. However, the target parameter is only a function of $P(O_r)$, such that $P(L1(K+1)|Y(K+1))$ is actually a nuisance parameter. Because the efficient influence function is always orthogonal to a nuisance tangent space (i. e. the space of scores one gets by varying the nuisance parameters), it is also orthogonal to the tangent space of $P(L1(K+1)|Y(K+1))$. It follows that $D^*(P) = \pi(D(P)|T_r(P))$ (i. e. the component in $T_1(P)$ is zero).

However, $D(P)$ is the efficient influence function of the target parameter based on O_r and is therefore an element of the tangent space of $P(O_r)$, so that $D(P)$ is in $T_r(P)$. This results in $\pi(D(P)|T_r(P)) = D(P)$, which completes the proof that $D^*(P) = D(P)$. That is, the efficient influence function of our parameter based on O_r is the same as the efficient influence function of our parameter based on O .

Funding: United States Agency for International Development (Grant/Award Number: “AID-623-A-12-0001”); National Institute of Allergy and Infectious Diseases (Grant/Award Number: “U01AI069911”); Doris Duke Clinical Scientist Development Award.

References

- Altman, N. (1992). An introduction to kernel and nearest-neighbor nonparametric regression. *The American Statistician*, 46:175–185.
- Attia, S., Egger, M., Müller, M., Zwahlen, M., and Low, N. (2009). Sexual transmission of HIV according to viral load and antiretroviral therapy: Systematic review and meta-analysis. *AIDS (London, England)*, 23:1397–1404. <http://www.ncbi.nlm.nih.gov/pubmed/19381076>.
- Bang, H., and Robins, J. M. (2005). Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61:962–973. <http://www.ncbi.nlm.nih.gov/pubmed/16401269>.
- Bodnar, L. M. (2004). Marginal structural models for analyzing causal effects of time-dependent treatments: An application in perinatal epidemiology. *American Journal of Epidemiology*, 159:926–934. <http://aje.oupjournals.org/cgi/doi/10.1093/aje/kwh131>.
- Boser, B. E., Guyon, I. M., and Vapnik, V. N. (1992). A training algorithm for optimal margin classifiers. In: *Proceedings of the 5th annual ACM workshop on computational learning theory*, 144–152, <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.21.3818>.
- Brooks, J. C., van der Laan, M. J., Singer, D. E., and Go, A. S. (2013). Targeted minimum loss-based estimation of causal effects in right-censored survival data with time-dependent covariates: Warfarin, stroke, and death in atrial fibrillation. *Journal of Causal Inference*, 1:235–254. <http://www.degruyter.com/view/j/jci.2013.1.issue-2/jci-2013-0001/jci-2013-0001.xml>.
- Bryan, J. (2004). Analysis of longitudinal marginal structural models. *Biostatistics*, 5:361–380. <http://biostatistics.oupjournals.org/cgi/doi/10.1093/biostatistics/kxg041>.
- Cohen, M., Chen, Y., McCauley, M., Gamble, T., Hosseinipour, M. C., Kumarasamy, N., Hakim, J. G., Kumwenda, J., Grinsztejn, B., Pilotto, J. H., Godbole, S. V., Sanjay, M., Chariyalertsak, S., Santos, B. R., Mayer, K. H., Hoffman, I. F., Eshleman, S. H., Piwowar-Manning, E., Wang, L., Makhema, J., Mills, L. A., de Bruyn, G., Sanne, I., Eron, J., Gallant, J., Havlir, D., Swindells, S., Ribaud, H., Elharrar, V., Burns, D., Taha, T. E., Nielsen-Saines, K., Celentano, D., Essex, M., and Fleming, T. R. (2011). Prevention of HIV-1 infection with early antiretroviral therapy. *The New England Journal of Medicine*, 365:493–505. <http://www.nejm.org/doi/full/10.1056/nejmoa1105243>.
- Cortes, C., and Vapnik, V. N. (1995). Support-vector networks. *Machine Learning*, 20:273–297.
- Danel, C., Moh, R., Gabillard, D., Badje, A., Le Carrou, J., Kouame, G. M., Ntakpe, J. B., Ménan, H., Eholie, S., and Anglaret, X. (2015). Conference on retroviruses and opportunistic infections. In *Early ART and IPT in HIV-Infected African Adults With High CD4 Count (Temprano Trial)*, volume 17.
- Das, M., Chu, P. L., Santos, G.-M., Scheer, S., Vittinghoff, E., McFarland, W., and Colfax, G. N. (2010). Decreases in community viral load are accompanied by reductions in new HIV infections in San Francisco. *PloS One*, 5:e11068. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2883572&tool=pmcentrez&rendertype=abstract>.
- Decker, A. L., Hubbard, A., Crespi, C. M., Seto, E. Y., and Wang, M. C. (2014). Semiparametric estimation of the impacts of longitudinal interventions on adolescent obesity using targeted maximum-likelihood: Accessible estimation with the LTMLE package. *Journal of Causal Inference*, 2:95–108. <http://www.degruyter.com/view/j/jci.2014.2.issue-1/jci-2013-0025/jci-2013-0025.xml>.
- Dieffenbach, C., and Fauci, A. (2011). Thirty years of HIV and AIDS: Future challenges and opportunities. *Annals of Internal Medicine*, 154:766–771. <http://annals.org/article.aspx?articleid=746972>.
- Friedman, J. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29:1189–1232.
- Friedman, J. H. (1991). Multivariate adaptive regression splines. 19:1–67.
- Friedman, J. H. (2002). Stochastic gradient boosting. *Computational Statistics and Data Analysis*, 38:367–378.
- Gelman, A., Jakulin, A., Pittau, M. G., and Su, Y.-S. (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics*, 2:1360–1383.
- Geng, E. H., Glidden, D. V., Bangsberg, D. R., Bwana, M. B., Musinguzi, N., Nash, D., Metcalfe, J. Z., Yiannoutsos, C. T., Martin, J. N., and Petersen, M. L. (2012). A causal framework for understanding the effect of losses to follow-up on epidemiologic analyses in clinic-based cohorts: The case of HIV-infected patients on antiretroviral therapy in Africa. *American Journal of Epidemiology*, 175:1080–1087.
- Geng, E. H., Odeny, T. A., Lyamuya, R., Nakiwogga-muwanga, A., Diero, L., Bwana, M., Braitstein, P., Somi, G., Kambugu, A., Bukusi, E., Wenger, M., Neilands, T. B., Glidden, D. V., Wools-kaloustian, K., and Yiannoutsos, C. (2016). Retention in care and patient-reported reasons for undocumented transfer or stopping care among HIV-infected patients on antiretroviral therapy in eastern Africa: Application of a sampling-based approach. *Clinical Infectious Diseases*, 62:935–944.
- Geng, E. H., Odeny, T. A., Lyamuya, R. E., Nakiwogga-Muwanga, A., Diero, L., Bwana, M., Muyindike, W., Braitstein, P., Somi, G. R., Kambugu, A., Bukusi, E. A., Wenger, M., Wools-Kaloustian, K. K., Glidden, D. V., Yiannoutsos, C. T., and Martin, J. N. (2015). Estimation of mortality among HIV-infected people on antiretroviral treatment in east Africa: A sampling based approach in an observational, multisite, cohort study. *The Lancet HIV*, 2:107–116.

- Giordano, T. P., Gifford, A. L., White, A. C., Almazor, M. E. S., Rabeneck, L., Hartman, C., Backus, L. I., Mole, L. A. and Morgan, R. O. (2007). Retention in care: A challenge to survival with HIV infection. *Clinical Infectious Diseases*, 44:1493–1499. <http://cid.oxfordjournals.org/lookup/doi/10.1086/516778>.
- Giordano, T. P., Visnegarwala, F., White, A. C., Troisi, C. L., Frankowski, R. F., Hartman, C. M., and Grimes, R. M. (2005). Patients referred to an urban HIV clinic frequently fail to establish care: Factors predicting failure. *AIDS Care*, 17:773–783. <http://www.ncbi.nlm.nih.gov/pubmed/16036264>.
- Gulick, R., Mellors, J., Havlir, D., Eron, J., Gonzalez, C., McMahon, D., Richman, D., Valentine, F., Jonas, L., Meibohm, A., Emini, E., and Chodakewitz, J. (1997). Treatment with indinavir, zidovudine, and lamivudine in adults with human immunodeficiency virus infection and prior antiretroviral therapy. *The New England Journal of Medicine*, 337:734–739. <http://www.nejm.org/doi/full/10.1056/NEJM199709113371102>.
- Hastie, T., and Tibshirani, R. (1986). Generalized Additive Models. *Statistical Science*, 3:297–318.
- Hernán, M. A. and Robins, J. M. (2006). Estimating causal effects from epidemiological data. *Journal of Epidemiology and Community Health*, 60:578–586. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2652882&tool=pmcentrez&rendertype=abstract>.
- Hoerl, A. E., and Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12:55–67.
- Horstmann, E., Brown, J., Islam, F., Buck, J., and Agins, B. (2010). Retaining HIV Infected patients in care: Where are we? Where do we go from here? *Clinical Infectious Diseases*, 50:100201102709029–000. <http://cid.oxfordjournals.org/lookup/doi/10.1086/649933>.
- Horvitz, D., and Thompson, D. (1952). A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47:663–685.
- Kissinger, P., Cohen, D., Brandon, W., Rice, J., Morse, A., and Clark, R. (1995). Compliance with public sector HIV medical care. *Journal of the National Medical Association*, 87:19–24. <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2607741/>.
- Kline, R. B. (2011). *Principles and Practice of Structural Equation Modeling*. 3rd Edition. New York: The Guilford Press.
- Koole, O., and Colebunders, R. (2010). ART in low-resource settings: How to do more with less. *Lancet*, 376:396–398. <http://www.ncbi.nlm.nih.gov/pubmed/20638119>.
- Koul, H., Susarla, V., and Ryzin, J. V. (1981). Regression analysis with randomly right-censored data. *The Annals of Statistics*, 9:1276–1288. <http://www.jstor.org/stable/2240417>.
- Lendle, S., Schwab, J., Petersen, M. L., and van der Laan, M. J. (2016). Ltmle: An R package implementing targeted minimum loss-based estimation for longitudinal data. *Journal of Statistical Software*, In Press.
- Lundgren, J., Babiker, A., Gordin, F., Emery, S., Fätkenheuer, G., Molina, J.-M., Wood, R., and Neaton, J. D. (2015). Why START? Reflections that led to the conduct of this large long-term strategic HIV trial. *HIV Medicine*, 16 Suppl 1:1–9. <http://www.ncbi.nlm.nih.gov/pubmed/25711317>.
- McCulloch, W. S., and Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5:115–133.
- Palella, F. J., Delaney, K., Moorman, A. C., Loveless, M. O., Jack, F., Satten, G. A., Aschman, D. J., and Holmberg, S. D. (1998). Declining morbidity and mortality among patients with advanced human immunodeficiency virus infection. *The New England Journal of Medicine*, 338:853–860. <http://www.nejm.org/doi/full/10.1056/NEJM199803263381301>.
- Pearl, J. (2009). *Causality*. 2 Edition. New York: Cambridge University Press.
- Petersen, M., and van der Laan, M. J. (2014). Causal models and learning from data: Integrating causal modeling and statistical estimation. *Epidemiology*, 25:418–426.
- Petersen, M. L. (2014). Commentary: Applying a causal road map in settings with time-dependent confounding. *Epidemiology (Cambridge, Mass)*, 25:898–901. <http://www.ncbi.nlm.nih.gov/pubmed/25265135>.
- Petersen, M. L., Porter, K. E., Gruber, S., Wang, Y., and van der Laan, M. J. (2012). Diagnosing and responding to violations in the positivity assumption. *Statistical Methods in Medical Research*, 21:31–54. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4107929&tool=pmcentrez&rendertype=abstract>.
- Petersen, M. L., Schwab, J., Gruber, S., and Blaser, N. (2013). Targeted maximum likelihood estimation for dynamic and static longitudinal marginal structural working models. *Journal of Causal Inference*, 18;2(2):147–185.
- Petersen, M. L., Tran, L., Geng, E. H., Reynolds, S. J., Kambugu, A., Wood, R., Bangsberg, D. R., Yiannoutsos, C. T., Deeks, S. G., and Martin, J. N. (2014). Delayed switch of antiretroviral therapy after virologic failure associated with elevated mortality among HIV-infected adults in Africa. *AIDS*, 28:2097–2107.
- Polley, E., and van der Laan, M. J. (2014). SuperLearner: Super learner prediction. <https://github.com/ecpolley/SuperLearner>.
- Quinn, T. C., Wawer, M. J., Sewankambo, N., Serwadda, D., Li, C., Wabwire-Mangen, F., Meehan, M. O., Lutalo, T., and Gray, R. H. (2000). Viral load and heterosexual transmission of human immunodeficiency virus type 1. *The New England Journal of Medicine*, 342:921–929. <http://www.nejm.org/doi/full/10.1056/NEJM200003303421303>.
- Robins, J. (1999a). Marginal structural models versus structural nested models as tools for causal inference. *Statistical Models in Epidemiology, the Environment*, 1–30. http://link.springer.com/chapter/10.1007/978-1-4612-1284-3_2.

- Robins, J., Greenland, S., and Hu, F. (1999). Estimation of the causal effect of a time-varying exposure on the marginal mean of a repeated binary outcome. *Journal of the American*, 94:687–700. <http://amstat.tandfonline.com/doi/abs/10.1080/01621459.1999.10474168>.
- Robins, J. M. (1986). A new approach to causal inference in mortality studies with a sustained exposure period – application to control of the healthy worker survivor effect. *Mathematical Modelling*, 7:1393–1512.
- Robins, J. M. (1987). A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *Journal of Chronic Disease*, 40:139S–161S.
- Robins, J. M. (1999b). Association, causation, and marginal structural models. *Synthese*, 121:151–179.
- Robins, J. M. (2000). Robust estimation in sequentially ignorable missing data and causal inference models. *Proceedings of the American Statistical Association Section on Bayesian Statistical Science*, 6–10.
- Robins, J. M., and Hernán, M. A. (2009). Estimation of the causal effects of time-varying exposures. In: *Longitudinal Data Analysis*, G. M. Fitzmaurice, M. Davidian, G. Verbeke, and G. Molenberghs (Eds.), CRC Press, chapter 1.i.
- Robins, J. M., Hernán, M. A., and Brumback, B. (2000a). Marginal structural models and causal inference in epidemiology. *Epidemiology*, 11:550–560. <http://content.wkhealth.com/linkback/openurl?sid=WKPTLP:landingpage&an=00001648-200009000-00011>.
- Robins, J. M., and Rotnitzky, A. (1992). Recovery of information and adjustment for dependent censoring using surrogate markers. In: *AIDS Epidemiology*, N. P. Jewell, K. Dietz, and V. T. Farewell (Eds.), 297–331. Boston: Birkhäuser, chapter 3.
- Robins, J. M., Rotnitzky, A., and van der Laan, M. J. (2000b). Discussion of on profile likelihood by Murphy and Van der vaart. *Journal of the American Statistical Association*, 95:477–482.
- Rotnitzky, A., and Robins, J. (2005). Inverse probability weighted estimation in survival analysis. Harvard University, <http://www.biostat.harvard.edu/robins/publications/IPW-survival-encyclopedia-submitted-corrected.pdf>. Accessed 26, Oct. 2016.
- Samet, J. H., Freedberg, K. A., Savetsky, J. B., Sullivan, L. M., Padmanabhan, L., and Stein, M. D. (2003). Discontinuation from HIV medical care: Squandering treatment opportunities. *Journal of Health Care for the Poor and Underserved*, 14:244–255. http://muse.jhu.edu/content/crossref/journals/journal_of_health_care_for_the_poor_and_underserved/v014/14.2.samet.html.
- Satten, G. A., and Datta, S. (2001). The Kaplan-Meier estimator as a weighted average. *The American Statistician*, 55:207–210. <http://www.jstor.org/stable/2685801>.
- Satten, G. A., and Datta, S. (2004). Marginal analysis of multistage data. In: *Handbook of Statistics: Advances in Survival Analysis*, 23 Edition, N. Balakrishnan and C. Rao (Eds.), 559–574. North Holland: Elsevier, chapter 32.
- Schnitzer, M. E., Moodie, E. E. M., van der Laan, M. J., Platt, R. W., and Klein, M. B. (2014). Modeling the impact of hepatitis C viral clearance on end-stage liver disease in an HIV co-infected cohort with targeted maximum likelihood estimation. *Biometrics*, 70:144–152. http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3954273&tool=pmcentrez&render_type=abstract.
- Stringer, J., Zulu, I., Levy, J., and Stringer, E. (2006). Rapid scale-up of antiretroviral therapy at primary care sites in Zambia: Feasibility and early outcomes. *JAMA*, 296:782–793. <http://jama.jamanetwork.com/article.aspx?articleid=203173>.
- Taubman, S. L., Robins, J. M., Mittleman, M. A., and Hernán, M. A. (2009). Intervening on risk factors for coronary heart disease: An application of the parametric g-formula. *International Journal of Epidemiology*, 38:1599–1611. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2786249&tool=pmcentrez&rendertype=abstract>.
- Tchetgen, E. J. T., and VanderWeele, T. J. (2012). On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21:55–75. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=4216807&tool=pmcentrez&rendertype=abstract>.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58:267–288.
- Tsiatis, A. (2006). *Semiparametric Theory and Missing Data*. New York: Springer.
- UNAIDS (2013). *Global Report: UNAIDS report on the global AIDS epidemic 2013*, Technical report.
- Van Damme, W., Kober, K., and Laga, M. (2006). The real challenges for scaling up ART in sub-Saharan Africa. *AIDS*, 20:653–656.
- van der Laan, M. J., and Dudoit, S. (2003). Unified cross-validation methodology for Unified cross-validation methodology for selection among estimators and a general cross-validated adaptive epsilon-net estimator: Finite sample oracle inequalities and examples, Technical Report 130.
- van der Laan, M. J., and Gruber, S. (2011). Targeted minimum loss based estimation of an intervention specific mean outcome. The Berkeley Electronic Press. Technical Report 290.
- van der Laan, M. J., Polley, E. C., and Hubbard, A. E. (2007). Super Learner, U.C. Berkeley Division of Biostatistics Working Paper Series, 1–20.
- van der Laan, M. J., and Robins, J. M. (2003). *Unified Methods for Censored Longitudinal Data and Causality*. New York: Springer.
- van der Laan, M. J., and Rose, S. (2011). *Targeted Learning*. New York: Springer.

- van der Laan, M. J., Rose, S., and Gruber, S. (2009). Readings on targeted maximum likelihood estimation, Technical report, Bepress, <http://www.bepress.com/ucbbiostat/paper254>.
- van der Laan, M. J., and Rubin, D. (2006). Targeted maximum likelihood learning, U.C. Berkeley Division of Biostatistics Working Paper Series, 1–87.
- van der Vaart, A. W., Dudoit, S., and van der Laan, M. J. (2006). Oracle inequalities for multi-fold cross validation. *Statistics & Decisions*, 24:351–371. <http://www.degruyter.com/view/j/stnd.2006.24.issue-3/stnd.2006.24.3.351/stnd.2006.24.3.351.xml>.
- Westreich, D., and Cole, S. R. (2010). Invited commentary: Positivity in practice. *American Journal of Epidemiology*, 171:674–677. Discussion 678–81, <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2877454&tool=pmcentrez&rendertype=abstract>.
- World Health Organization (2013a). Consolidated Guidelines on the Use of Antiretroviral Drugs for Treating and preventing HIV Infection, June, London, http://apps.who.int/iris/bitstream/10665/85321/1/9789241505727_eng.pdf.
- World Health Organization (2013b). Global update on HIV treatment 2013: Results, impact, and opportunities, Technical Report June, Geneva.
- Zheng, W., and van der Laan, M. (2010). Asymptotic theory for cross-validated targeted maximum likelihood estimation, U.C. Berkeley Division of Biostatistics Working Paper Series, <http://biostats.bepress.com/ucbbiostat/paper273/>.
- Zwahlen, M., Harris, R., May, M., Hogg, R., Costagliola, D., de Wolf, F., Gill, J., Fätkenheuer, G., Lewden, C., Saag, M., Staszewski, S., d’Arminio Monforte, A., Casabona, J., Lampe, F., Justice, A., von Wyl, V., and Egger, M. (2009). Mortality of HIV-infected patients starting potent antiretroviral therapy: Comparison with the general population in nine industrialized countries. *International Journal of Epidemiology*, 38:1624–1633. <http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3119390&tool=pmcentrez&rendertype=abstract>.