

History-Adjusted Marginal Structural Models: Time-Varying Effect Modification

Maya L. Petersen*

Mark J. van der Laan[†]

*Division of Biostatistics, School of Public Health, University of California, Berkeley, mayaliv@hotmail.com

[†]Division of Biostatistics, School of Public Health, University of California, Berkeley, laan@berkeley.edu

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/paper173>

Copyright ©2005 by the authors.

History-Adjusted Marginal Structural Models: Time-Varying Effect Modification

Maya L. Petersen and Mark J. van der Laan

Abstract

Marginal structural models (MSM) provide a powerful tool for estimating the causal effect of a treatment, particularly in the context of longitudinal data structures. These models, introduced by Robins, model the marginal distributions of treatment-specific counterfactual outcomes, possibly conditional on a subset of the baseline covariates. However, standard MSM cannot incorporate modification of treatment effects by time-varying covariates. In the context of clinical decision-making such time-varying effect modifiers are often of considerable interest, as they are used in practice to guide treatment decisions for an individual. In this article we introduce a generalization of marginal structural models, which we call history-adjusted marginal structural models (HA-MSM). These models allow estimation of adjusted causal effects of treatment, given the observed past, and are therefore more suitable for making treatment decisions at the individual level and for identification of time-dependent effect modifiers. We provide a practical introduction to HA-MSM relying on an example drawn from the treatment of HIV, and discuss parameters estimated, assumptions, and implementation using standard software.

1 Introduction.

Marginal structural models (MSM), introduced by Robins, (e.g., Robins (2000), Robins (1999), and van der Laan and Robins (2002)), represent a major advance in the statistical methodology for causal inference. In epidemiological and clinical research, subjects are often followed over time and longitudinal data on treatment status and covariates collected. When treatment status can change over time, conventional analytic approaches (such as standard multivariable regression methods) often fail to allow valid causal inferences about the effect of treatment. Marginal structural models address this well-recognized problem. In addition, in longitudinal studies many covariates are often collected over multiple time points, resulting in a high-dimensional modelling problem in which likelihood-based estimation approaches can be highly susceptible to model misspecification. Marginal structural models offer alternative estimators of causal effects that focus on the causal parameter of interest, increasing robustness to misspecification of other parts of the likelihood. As a result, in many research contexts marginal structural models are the best available analytic method.

Marginal structural models to date have been restricted to the estimation of the causal effects of treatment or exposure conditional on baseline covariates. Thus it has been possible to use this methodology to address questions such as "What is the effect of a treatment and how does it differ between study members with different covariate values at entry to the study?", but not "How does the effect of a treatment differ as a result of changing values of a covariate over the course of the study?". In other words, it has not been possible to use MSM to estimate modification of causal effects by time-dependent covariates. In many research settings, estimation of time-dependent effect modification is of major interest. In clinical settings in particular, treatment decisions are often modified over time as a result of changing values of a patient's covariates. Estimation of how past values of these covariates modify the future causal effect of a treatment has important implications for understanding the mechanistic action of the treatment, as well as guiding clinical decision-making.

In this paper we introduce a generalization of MSM to allow estimation of effect modification by time-varying covariates, which we call history-adjusted marginal structural models (HA-MSM). The aim of the paper is to provide a practical and intuitive understanding of this new methodology- the formal statistical theory is presented in a companion paper. Here, we provide a non-

technical discussion of the assumptions underlying this method and discuss implementation using standard software and interpretation of results. We rely throughout the paper on an example drawn from the treatment of Human Immunodeficiency Virus (HIV) to illustrate our notation and results.

2 Example: Antiretroviral therapy for the treatment of HIV infection.

Antiretroviral therapy (ART) is able to suppress HIV replication to an extent that virus in the patient's blood, as measured by plasma HIV RNA level (viral load), becomes undetectable. Unfortunately, the virus infecting many treated individuals eventually develops resistance to the drug regimen being used, resulting in a rebound of viral load to detectable levels (Ledgergerber et al. (1999), Lucas et al. (1999)). Viral rebound is frequently followed by a decline in immunologic function, as reflected by CD4 T-cell count, and by disease progression if the patient is not switched to a new therapy regimen that is able to re-suppress the virus (Panel on Clinical Practices for Treatment of HIV infection (2004)).

In contrast, some patients maintain stable CD4 T-cell counts over time, despite loss of virologic suppression (Deeks et al. (2002)). One hypothesis advanced to explain this phenomenon is that mutations necessary to confer viral resistance to some types of antiretroviral drugs result in a loss of viral fitness, and as a result, a reduced ability of the virus to deplete CD4 T-cells (Deeks (2001)). Over time, the virus may accumulate additional compensatory mutations that help restore viral fitness. In this article we address whether the presence of certain resistance mutations, which emerge over time, can modify the effect of future non-suppressive antiretroviral treatment on CD4 T-cell count depletion. Demonstration of such time-dependent effect modification would have important clinical consequences, suggesting that data on viral resistance mutations could help inform decisions about when to switch antiretroviral therapy regimen in patients with incomplete virologic suppression.

In order to address this question, we use data from a cohort of HIV-infected patients currently treated with ART. Subjects become eligible for our analysis when their viral load rebounds to a detectable level. Data on covariates and treatment are collected on each subject throughout the study until follow-up ends at time $K + 1$. Covariates at time t are denoted $L(t)$, $t =$

$0, \dots, K + 1$, and include CD4 T-cell count, the outcome of interest, denoted $Y(t)$, and viral mutation profile, denoted $S(t)$. For the sake of exposition, we treat $S(t)$ here as a binary variable, denoting the presence or absence of a set of key viral mutations thought to affect fitness. Of course, one could also make $S(t)$ a vector, where each component is an indicator that a given mutation is present. This type of treatment status is a binary variable indicating whether a subject has switched off of his or her original non-suppressive ART regimen or not, and is denoted $A(t), t = 0, \dots, K$. Thus, the data observed on each randomly sampled subject can be written: $O = (L(0), A(0), \dots, L(K), A(K), L(K + 1))$, where we assume $L(t)$ is measured before $A(t)$. For pedagogical purposes, we assume that variables are discrete valued, and that a subject can only switch to a new therapy once; we define $A(t) = 1$ as long as a subject remains on his/her original treatment regimen, and let $A(t)$ jump to 0 and remain there as soon as a subject switches to a new treatment. The method can be easily extended, however, to encompass more complex treatment patterns.

3 Marginal Structural Models and the counterfactual framework for causal inference

Marginal structural models are based on a counterfactual definition of causal effects. A counterfactual outcome is defined as the outcome that would have been observed for a given individual under a specific treatment, whether or not the individual did in fact receive that treatment. In the counterfactual framework, the causal effect of a treatment regimen on an individual is defined as the difference in that individual's counterfactual outcome if he/she received the treatment regimen vs. did not receive the treatment regimen. Marginal structural models are models of how the population distribution of these counterfactual outcomes changes as result of changes in treatment.

Introduction of some standard notation used in marginal structural models helps to define the assumptions on which they are based. We denote a time-dependent process $X(t)$, measured sequentially over time from the beginning of the study to time t , as $\bar{X}(t)$. We thus denote treatment over the course of the study ($t = 0, \dots, K$) as $\bar{A}(K) = (A(0), \dots, A(K))$, and covariates measured over the course of the study $\bar{L}(K + 1)$. We use \mathcal{A} to denote all possible longitudinal treatment regimens, and $\mathcal{A}(t)$ to denote all possi-

ble longitudinal treatment regimens through time t . Finally, we use $\bar{L}_{\bar{a}}(t)$ to denote a counterfactual covariate process, consisting of the values of all measured covariates up till time t that would have been observed if a subject had followed a treatment regimen $\bar{A} = \bar{a}$.

Marginal structural models assume that the outcome or covariate process we observe for each individual is equivalent to the counterfactual outcome or covariate process for that individual under the treatment regimen he/she actually received (Consistency Assumption: $O = (\bar{A}(K), \bar{L}_{\bar{A}}(K+1))$). Under this assumption, causal inference can be treated as a missing data problem. If no data were missing, we would observe for each individual the counterfactual covariate process he/she would have followed under each possible treatment regimen. If we had access to these full data we could use them to fit the marginal structural model defining our causal effect of interest. Instead, the observed data consist only of each individual's covariates under his/her observed treatment regimen.

Marginal structural models are thus defined using the full counterfactual data, but must be estimated using the observed data. In order to make them identifiable, we assume sequential randomization of the treatment with respect to subject's past measured covariates (Sequential Randomization Assumption: $A(t) \perp Y_{\bar{a}} | \bar{A}(t-1), \bar{L}(t), t = 0, \dots, K$). Under the sequential randomization assumption (SRA), treatment assignment at each time point is allowed to depend on any measured characteristics of the subject's past at that time point, but not on any unmeasured characteristics that also affect outcome. In other words, there must be no unmeasured confounders of the effect of treatment on outcome.

Finally, we assume that the counterfactual values of covariates cannot be affected by treatment that occurs after they are measured (Temporal Ordering Assumption: $L_{\bar{a}}(t) = L_{\bar{a}(t-1)}(t)$).

Applied to our HIV example, $\bar{A}(K)$ is a vector of binary variables, consisting of one for each time point until a subject switches therapy, and zero thereafter. For each possible time until switching, $(\bar{L}_{\bar{a}}(t)(K+1))$ denotes the counterfactual CD4 T-cell counts, viral load, mutations, etc. over time that would have been observed if the subject had switched therapy at the time implied by $\bar{A} = \bar{a}$. The full data for an individual thus consist of $X^{Full} = (\bar{L}_{\bar{a}}(K+1), \bar{a} \in \mathcal{A})$. Our outcome of interest for a given time point t is the counterfactual CD4 T-cell count measured m months in the future under the switching time indicated by \bar{a} , denoted $Y_{\bar{a}}(t+m)$. Under the temporal ordering and consistency assumptions, the observed data for a subject

can be represented as his/her actual switch time and the counterfactual CD4 T-cell count, viral mutation process, etc. corresponding to that switch time. In addition, we assume that there are no unmeasured variables that predict, at any time point, both probability of switching treatment and also final CD4 T-cell count (the SRA).

4 Definition of history-adjusted marginal structural models: What do HA-MSM estimate?

HA-MSM model some parameter of the counterfactual outcome if the study population were to follow their observed treatment history up till time j , followed by a specified counterfactual future treatment history until outcome is measured, conditional on a subset of (possibly time-varying) covariates and/or treatment history measured before time j . In this article, we will focus on HA-MSM concerned with the mean of these counterfactual outcomes; however, the same framework can be readily adopted to model the median or any other parameter. Before we discuss the interpretation of this parameter, we introduce some final notation to facilitate its formal definition.

We denote a future longitudinal treatment regimen, beginning at time j and continuing until the outcome is measured m time points later, as $\underline{a}(j, j+m-1) \equiv (a(j), a(j+1), \dots, a(j+m-1))$, for $j = 0, \dots, K+1-m$. Let $V(j) \subset (\bar{L}(j), \bar{A}(j-1))$ be a subset of a subject's treatment and covariate history up till time j . For each time point in the study for which the outcome m time points later is defined, $j = 0, \dots, K+1-m$, HA-MSM model some parameter of the counterfactual outcome $Y_{\bar{A}(j-1), \underline{a}(j, j+m-1)}(j+m)$, conditional on $V(j)$, under each possible future treatment regimen. We note that, in the case that the set of possible future treatment regimens beginning at time j depends on a subject's treatment history up till time j , we must include treatment history in the subset of a subject's past on which we condition ($\bar{A}(j-1) \subset V(j)$), so that the set of possible future treatments is the same for all subjects with a given $V(j)$ value. For example, if only a single change in treatment is allowed over the course of the study, whether or not such a treatment has already occurred will restrict whether or not it can occur in the future, and so when considering counterfactuals indexed by possible future switch times, we must condition on a subject not having already switched treatment.

HA-MSM in this paper are concerned with inference for the following parameter:

$$E_0(Y_{\bar{A}(j-1), \underline{a}(j, j+m-1)}(j+m) \mid V(j)), \quad j = 0, \dots, K+1-m \quad (1)$$

where subscript 0 denotes the true parameter value.

To clarify both the notation and the interpretation of this parameter, we rely again on our HIV example. Future antiretroviral treatment after time j until outcome is measured, denoted $\underline{a}(j, j+m-1)$, consists of a vector of treatment decisions $(a(j), \dots, a(j+m-1))$, where $a(t) = 0$ if a subject has switched treatment at or before time t , and otherwise $a(t) = 1$. This vector of future treatment decisions exists for each subject beginning at each time point $j = 0, \dots, K+1-m$.

We summarize $\underline{a}(j, j+m-1)$ as $c(j) \equiv \sum_{l=j}^{j+m-1} a(l)$, which represents the future time (after time point j) that the subject will spend on his/her original failing therapy before either switching or the outcome is measured. The viral mutation profile at time j (presence or absence of mutations affecting viral fitness) is denoted $S(j)$, a subset of the full covariate history measured over time, $\bar{L}(j)$. For each time point j , we are interested in the mean counterfactual CD4 T-cell counts m months later among individuals who have not yet switched therapy, if they were to follow their observed treatment history up till time j (not switching), and then switch therapy at a specified time after j . In addition, we are interested in how these counterfactual CD4 T-cell counts may differ depending on the presence of key viral mutations at time j .

For example, for $j = 6$

$$E_0 \left(Y_{\bar{A}(5)=1, c(6)=0}(6+m) \mid \bar{A}(5) = 1, S(6) = 1 \right), \quad (2)$$

denotes the mean counterfactual CD4 T-cell count m months later among the subgroup of individuals who have not switched therapy and who have key viral mutations present at time point 6, if that subgroup had followed its observed treatment history through time 5 (not switching), and then switched therapy at time 6. If we knew the counterfactual outcomes under an immediate switch of treatment for every member of this subgroup, we could simply take the mean of these outcomes. However, in the observed data, we only observe an immediate switch time for some (*non-random*) members of this subgroup.

Although the assumptions outlined above ensure that our counterfactuals of interest are identifiable from the observed data, in most settings the high dimensions of the data require that we assume some model for (1) to make estimation feasible using finite sample sizes. In other words, to reduce the dimensions of the problem we assume some model for the dependence of the counterfactual outcome on future and past treatment regimens, conditional on a subset of the observed past. Formally, we assume a model $m_{\beta_0}(j, \underline{a}(j, j+m-1) | V(j))$ for our parameter of interest $E_0(Y_{\bar{A}(j-1), \underline{a}(j, j+m-1)}(j+m) | V(j))$, $j = 0, \dots, K+1-m$. The model adopted is flexible and should be driven by our knowledge of the subject matter. To illustrate one potential model, we return to our HIV example.

Recall that we are interested in answering the question “Given a subject’s current viral mutation profile and the length of time he/she has already spent on non-suppressive ART, how does additional time on non-suppressive therapy affect CD4 T-cell count m months later?”. We hypothesize that the effect of future time until switching will differ depending on a subject’s current viral mutation profile; in the presence of mutations that restore viral fitness, non-suppressive therapy is hypothesized to result in a faster depletion of CD4 T-cells. We further hypothesize that the effect of future time until switching will differ depending how long a subject has already spent on non-suppressive therapy.

To address these questions, we might assume the following model :

$$E_0(Y_{\bar{A}(j-1), \underline{a}(j, j+m-1)}(j+m) | \bar{A}(j-1), S(j)) = I(A(j-1) = 1)(\beta_0 + \beta_1 * c(j) + \beta_2 * S(j) + \beta_3 * j + \beta_4 * c(j) * S(j) + \beta_5 * c(j) * j), \quad (3)$$

$j = 0, \dots, K+1-m$. In other words, we might assume that, among individuals who have not yet switched treatment ($A(j-1) = 1$), counterfactual CD4 T-cell count m months later depends linearly on additional time until switching ($c(j)$), but the magnitude of this effect may differ depending on the duration a patient has already spent on non-suppressive therapy (j) and the current mutation profile of the HIV virus ($S(j)$). This model allows us to estimate the effect of each additional month until switching to a new therapy on CD4 T-cell count m months later, among patients who have been on their current non-suppressive therapy for different durations, and among patients infected with virus that does and does not harbor key mutations. For example, by testing whether $\beta_4 = 0$ we are testing our hypothesis that

the causal effect of increased future exposure time differs depending on a subject's current viral mutation profile.

5 Implementation of HA-MSM: the Inverse Probability of Treatment Weighted estimator.

Several HA-MSM estimators are available; here, we focus on the Inverse Probability of Treatment Weighted (IPTW) estimator, which can be implemented using standard software. The IPTW estimator can be understood as simply a weighted least squares estimator. For each time point j in the study, each subject receives a weight. The denominator of this weight is informally the subject's probability of receiving the treatment that he or she actually received, from time point j until the outcome is measured. If a subject has a longitudinal treatment regimen beginning at time point j that occurs frequently in the data among subjects with his/her covariate and treatment history, he/she receives a small j -specific weight. In contrast, if the subject has an unusual longitudinal treatment regimen given his/her covariates, the subject will receive a large weight. Applied to our HIV example, it is likely that the majority of patients will switch therapy quickly following a decline in their CD4 T-cell counts. A subject that did not switch therapy despite a recent decline in CD4 T-cell count would thus have a small predicted probability of receiving his/her observed treatment, and receive a large weight. The goal of the weighting process is to create a new dataset, in which the effect of exposure time on outcome is no longer confounded. The least squares estimator is then calculated using this weighted dataset.

In addition to the assumptions discussed above, the IPTW estimator relies on the additional assumption that for each time point, for every subpopulation defined by treatment and covariate history up to that time point, every possible treatment that is compatible with the treatment history up till that time point must have positive probability of occurring (Experimental Treatment Assignment assumption: For all possible $\bar{A}(t-1), \bar{L}(t)$, $P(A(t) = a(t) | \bar{A}(t-1), \bar{L}(t)) > 0$ for all $a(t) \in a^*(t) : (\bar{A}(t-1), a^*(t)) \in \mathcal{A}(t)$) In our example, we assume that all individuals who have not yet switched treatment at each time point have some positive probability of both remaining on their current treatment and of switching to a new treatment at that time point,

regardless of their covariate values.

The first step in implementing the IPTW estimator is to assume a model for the treatment mechanism, a predictive model of treatment at each time point t , given the observed past up till that time point.

$$g(A(t)|\bar{A}(t-1), \bar{L}(t)), t = 0, \dots, K \quad (4)$$

The parameters of the model can be estimated using maximum likelihood. In specifying the treatment mechanism, one can choose to specify a separate model for treatment at each time point, or have a single model for all time points, with or without including time as a covariate in the model. For example, one might choose a single logistic regression model for the treatment decision (switch therapy or not) made at every time point t ,

$$P(A(t) = 1 | A(t-1) = 1, CD4(t)) = \frac{1}{1 + \exp^{-(\theta_0 + \theta_1(CD4(t)) + \theta_2(t))}}, t = 0, \dots, K \quad (5)$$

where $CD4(t)$ CD4 T-cell count at time t . If we believe the model (5), we assume that the decision about whether or not to switch treatment at each time point is made based on a patient's most recent CD4 T-cell count and the time that has elapsed. Recall that a subject switches therapy only once; once a subject switches, he/she is no longer at risk of switching in the future. Thus, when fitting our model of the probability of staying on therapy at a given time point ($A(t) = 1$), we condition on not having already switched before that time point ($A(t-1) = 1$).

For the IPTW estimator to be consistent, the estimate of the treatment mechanism must be consistent. In addition, enough covariates must be included in the treatment model so that outcome is independent of treatment assignment conditional on the variables in the model (or in other words, there must not be additional confounders that do not appear in the model of the treatment mechanism).

For each time point $j = 0, \dots, K + 1 - m$, the model of the treatment mechanism (5) is used to estimate the denominator of the j -specific weight: $\prod_{l=j}^{j+m-1} g(A(l) | \bar{A}(l-1), \bar{L}(l))$. In our example, for subjects who do not switch therapy before the outcome is measured m months later, the denominator of the j -specific weight is $\prod_{l=j}^{j+m-1} P(A(l) = 1 | A(l-1) = 1, CD4(l))$. For subjects who have not switched therapy by time j , but who switch at some point $T = j + C(j)$ before the outcome is measured ($C(j) < m$), the denominator of the j -specific weight is

$(1 - P(A(T) = 1 \mid A(T - 1) = 1, CD4(T))) \prod_{l=j}^{T-1} P(A(l) = 1 \mid A(l - 1) = 1, CD4(l))$. For subjects that have already switched therapy by time j , the denominator of the j -specific weight is always 1.

For example, suppose we have a subject who switched therapy at time point 6. To model the weight for this subject corresponding to $j = 5$, $C(j) = 1$, one begins by using the model of the treatment mechanism and the subject's covariate and treatment history to estimate the probability that the subject followed his/her observed treatment history beginning at time point 5. Using the model described above (5), one first estimates the probability that the subject did not switch therapy at time 5, given the subject's CD4 T-cell count at time 5 and the elapsed time ($P(A(5) = 1 \mid A(4) = 1, CD4(5))$). One then uses the observed CD4 T-cell count at time 6, $CD4(6)$, and the elapsed time, $t = 6$, to estimate the probability of switching therapy at time 6 ($1 - P(A(6) = 1 \mid A(5) = 1, CD4(6))$). By repeating this process for each time point until a subject switches, one generates an estimate of the probability of the subject receiving his/her observed treatment at each time point (after a subject switches, the probability of receiving his/her observed treatment is 1); the product of these probabilities provides an estimate of the probability of the subject receiving his/her observed treatment over time, from time point j until the outcome is measured.

Choice of a numerator for the weights will not affect the consistency of the IPTW estimator, as long as the numerator is only a function of treatment history and baseline covariates at time j . A standard choice of numerator is $g^*(\underline{A}(j, j + m - 1) \mid V(j)) \equiv \prod_{l=j}^{j+m-1} g^*(A(l) \mid \bar{A}(l - 1), V(j))$. To estimate this numerator for each subject's j -specific weights one simply fits a model of the treatment decision at each time point $t = j, \dots, K$, analogous to the model of the treatment mechanism fit for the denominator, but now including only those covariates contained in the subset of baseline measurements at time j . Applied to the treatment mechanism in our HIV example, (5), the model used to generate the numerators would be a function of the elapsed time, t (the treatment mechanism specified in (5) does not include any baseline covariates, $V(j)$).

Finally, once each subject has been assigned a set of $K + 1 - m$ weights, the simple least squares estimator of the parameter of interest can be estimated on the pooled data, using standard software, with each subject contributing $K + 1 - m$ weighted lines of data. For example, in estimating the parameters of the simple linear model described above (3), each j -specific line of data would include

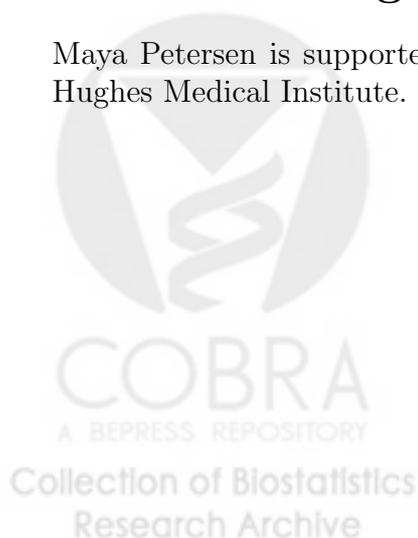
- An indicator that the subject remained on non-suppressive therapy at the previous time point ($I(A(j-1) = 1)$)
- Elapsed time since loss of suppression occurred (j)
- Remaining time until the subject either switched therapy or the outcome was measured ($C(j)$)
- The subject's current viral mutations ($S(j)$)
- The subjects CD4 T-cell count 4 months later ($Y(j+m)$)
- The subject's j -specific weight One then simply runs a weighted least squares regression on the pooled data, providing the software with the weights.

6 Discussion.

The history-adjusted marginal structural models presented in this paper represent an important generalization of marginal structural model methodology. Marginal structural models are well-established as powerful tools for causal inference, particularly in the the setting of longitudinal data. In this article we have introduced an extension of marginal structural models to identify and estimate time-dependent causal effect modification. The methods outlined in this paper demonstrate how history-adjusted marginal structural models can be implemented using standard software. We believe this method will prove to be an extremely useful tool for using observational data to improve clinical decision-making based on time-varying covariates.

7 Acknowledgements

Maya Petersen is supported by a Predoctoral Fellowship from the Howard Hughes Medical Institute.



References

- S. G. Deeks. Virologic and immunologic consequences of discontinuing combination antiretroviral drug therapy in HIV-infected patients with detectable viremia. *New England Journal of Medicine*, 344(7):472–480, 2001.
- S. G. Deeks, J. D. Barbour, R. M. Grant, and J. M. Martin. Duration and predictors of CD4 T-cell gains in patients who continue combination therapy despite detectable plasma viremia. *AIDS*, 16(2):201–207, 2002.
- B. Ledergerber, M. Egger, M. Oravil, and et. al. Clinical progression and virologic failure on highly active antiretroviral therapy in HIV-1 patients: A prospective cohort study. *Lancet*, 353(9156):863–868, 1999.
- G. M. Lucas, R. E. Chaisson, and R. D. Moore. Highly active antiretroviral therapy in a large urban clinic: Risk factors for virologic failure and adverse drug reactions. *Annals of Internal Medicine*, 131(2):81–87, 1999.
- Panel on Clinical Practices for Treatment of HIV infection. Guidelines for the use of antiretroviral agents in HIV-1-infected adults and adolescents. Technical report, Department of Health and Human Services, 2004.
- J.M. Robins. Robust estimation in sequentially ignorable missing data and causal inference models. In *Proceedings of the American Statistical Association: Section on Bayesian Statistical Science*, pages 6–10, 1999.
- J.M. Robins. Marginal structural models versus structural nested models as tools for causal inference. In *Statistical models in epidemiology, the environment, and clinical trials (Minneapolis, MN, 1997)*, pages 95–133. Springer, New York, 2000.
- M.J. van der Laan and J.M. Robins. *Unified methods for censored longitudinal data and causality*. Springer, New York, 2002.

