

Comparison of Dynamic Treatment Regimes via Inverse Probability Weighting

Miguel A. Hernán¹, Emilie Lanoy², Dominique Costagliola² and James M. Robins^{1,3}

¹Department of Epidemiology, Harvard School of Public Health, Boston, MA, U.S.A.,

²U720 INSERM et Université Pierre et Marie Curie, Paris, France, and

³Department of Biostatistics, Harvard School of Public Health, Boston, MA, U.S.A.

(Received September 4, 2005; Accepted November 7, 2005)

Abstract: Appropriate analysis of observational data is our best chance to obtain answers to many questions that involve dynamic treatment regimes. This paper describes a simple method to compare dynamic treatment regimes by artificially censoring subjects and then using inverse probability weighting (IPW) to adjust for any selection bias introduced by the artificial censoring. The basic strategy can be summarized in four steps: 1) define two regimes of interest, 2) artificially censor individuals when they stop following one of the regimes of interest, 3) estimate inverse probability weights to adjust for the potential selection bias introduced by censoring in the previous step, 4) compare the survival of the uncensored individuals under each regime of interest by fitting an inverse probability weighted Cox proportional hazards model with the dichotomous regime indicator and the baseline confounders as covariates. In the absence of model misspecification, the method is valid provided data are available on all time-varying and baseline joint predictors of survival and regime discontinuation. We present an application of the method to compare the AIDS-free survival under two dynamic treatment regimes in a large prospective study of HIV-infected patients. The paper concludes by discussing the relative advantages and disadvantages of censoring/IPW versus g-estimation of nested structural models to compare dynamic regimes.

Several randomized (Hammer *et al.* 1997) and observational (Cole *et al.* 2003; Sterne *et al.* 2005) studies have shown that continuous use of highly active antiretroviral therapy (HAART) improves the acquired immunodeficiency syndrome (AIDS)-free survival of human immunodeficiency virus (HIV)-infected patients compared with no use of HAART.

The next logical step is to determine the optimal time at which to initiate HAART to maximize expected survival or perhaps AIDS-free survival. Current guidelines generally recommend that HAART be initiated the first time CD4 cell count drops below 350 cells/ μ l. Formally, this recommendation corresponds to the treatment regime “start HAART when the CD4 cell count first drops below 350 cells/ μ l, then always treat.” This regime is an example of a dynamic treatment regime. A dynamic treatment regime is one in which the treatment assigned depends upon a subject’s evolving time-dependent covariate history. In contrast, most previous studies have compared non-dynamic regimes, like “always treat” and “never treat,” which assign treatment (or no treatment) regardless of the subject’s history. It is widely argued that a dynamic regime in which initiation

of HAART depends optimally on the subject’s CD4 cell count should result in longer expected survival than any non-dynamic regime: if one begins therapy before it is really needed (say, when the CD4 cell count exceeds 500), there is additional time for viral resistance or drug-related toxicities to develop. On the other hand, if therapy is withheld for too long (say, until the CD4 cell count drops below 200 cells/ μ l), permanent irreversible damage to the immune system may have already occurred.

Thus, there is great current interest in answering questions like the following: Is the (dynamic) regime “start HAART when CD4 cell count first drops below 500 cells/ μ l, then always treat” better than the (dynamic) regime “start HAART when CD4 cell count first drops below 200 cells/ μ l, then always treat”? It is, however, unlikely that long-term trials will be ever conducted to compare each of the possible cutoff points defined by the time-varying covariate CD4 cell count (500 and 200 are only two examples). Hence appropriate analysis of observational data is our best chance to obtain answers to many questions that involve dynamic treatment regimes.

This article describes a simple method introduced in Robins (1993) to compare dynamic treatment regimes by artificially censoring subjects and then adjusting for selection bias due to censoring by inverse probability weighting (IPW). We present an application of the method to compare the AIDS-free survival under two dynamic HAART re-

Author for correspondence: Miguel Hernán, Department of Epidemiology, Harvard School of Public Health, 677 Huntington Avenue, Boston, MA 02115, U.S.A. (fax +1 617 432 1884, e-mail miguel_hernan@post.harvard.edu).

gimes in a large prospective study of HIV-infected patients. The article concludes by discussing the relative advantages and disadvantages of censoring/IPW versus g-estimation of nested structural models to compare dynamic regimes. Because, as we will see, censoring/IPW can be similarly applied to compare either non-dynamic or dynamic treatments, we start by reviewing the simpler setting in which the question of interest involves non-dynamic regimes only.

Comparing non-dynamic regimes

To fix ideas, let us consider the two non-dynamic regimes “always treat” and “never treat” with HAART. Our goal is to decide which of these two non-dynamic regimes results in a longer AIDS-free survival time in HIV-infected patients. For convenience, we will often use the terms treatment and survival (or mortality) rather than HAART and AIDS-free survival (or AIDS/mortality).

A randomized experiment to compare these two regimes could be designed as follows. First, recruit a cohort of patients recently infected with HIV and that have never received antiretroviral therapy. Second, assign them randomly to either of the two regimes “always treat” ($A=1$) or “never treat” ($A=0$). Third, administer them their assigned treatment regime, HAART or placebo, until the end of follow-up (say, five years after randomization), AIDS, or death, whichever occurs earlier. Finally, choose a survival analysis method to compare the mortality between the two groups. For example, use a proportional hazards model to contrast the mortality rates or an accelerated failure time model to contrast the median survival times in each group.

When the data from this randomized experiment are analyzed under the intention to treat (ITT) principle, the treatment variable does not change with time. That is, for the purposes of the analysis, each subject is considered to remain in the treatment group ($A=1$ or $A=0$) to which she was assigned at baseline, regardless of whether she actually adhered to the assigned regime. A consequence of working under the ITT approach is that one does not really compare the regimes “always treat” and “never treat”, but rather the regimes “intent to always treat” and “intent to never treat”. In the presence of a treatment effect, the two contrasts will only be identical when all participants fully comply with the assigned treatment.

A non-randomized (observational) study to compare the two regimes “always treat” and “never treat” could be designed and conducted as follows. First, identify a cohort of patients recently infected with HIV and that have never received antiretroviral therapy. Second, develop a system to identify their treatment status over time, e.g., $A(t)=1$ if they received treatment or $A(t)=0$ if they did not receive treatment at time t after study entry. Third, record data on treatment status and confounders until the end of follow-up (say, five years), AIDS, or death, whichever occurs earlier. Throughout this article we will assume that the collected information included data on all the time-dependent confounders for the effect of antiretroviral therapy on mor-

tality, that is, on all time-varying risk factors for mortality that are used to decide whether to treat a patient at time t . Finally, choose a survival analysis method to compare the mortality between the groups. The problem here is what groups? Unlike in the randomized experiment with full compliance, participants in an observational study do not necessarily follow the pre-specified regimes “always treat” or “never treat.” Some individuals may happen to follow those regimes but many others will follow different regimes like “do not treat during the first year, then always treat.” There could be as many different regimes as participants in the observational study.

The two main solutions to the problem of multiplicity of regimes in observational studies with time-varying treatments are:

1. Modeling the causal effect of the time-varying treatment on mortality. For example, one can use a proportional hazards model to model the mortality rate, or an accelerated failure time model to model the survival time, under all possible non-dynamic treatment regimes. If this “dose-response” model is correctly specified, then one simply needs to use the model to estimate the hazard rate, or the median survival time, under each of the two regimes of interest and compare them. We have previously applied this approach to compare the regimes “always treat” and “never treat” by using a marginal structural Cox model (Cole *et al.* 2003; Sterne *et al.* 2005) or a nested structural accelerated failure time model (Hernán *et al.* 2005). These methods allowed us to appropriately adjust for measured time-dependent confounding, even in the presence of time-varying confounders that are affected by – or share common causes with – prior treatment history (Hernán *et al.* 2004), because the parameters of marginal structural models are estimated by inverse probability weighting (IPW) and the parameters of nested structural models by g-estimation. (Appropriate adjustment for time-dependent selection bias due to loss to follow-up required additional IPW whether we used marginal or nested structural models.) Both the hazard ratio estimated from a marginal structural Cox model and the median survival time ratio estimated from a nested structural accelerated failure time model indicated that the regime “always treat” was superior to the regime “never treat.” This finding was consistent with results from randomized trials (Hammer *et al.* 1997).
2. Artificially censoring individuals when they deviate from one of the two regimes of interest. For example, a patient who was untreated at baseline, $A(0)=0$, will be considered to be following the regime “never treat” ($A=0$) and we will censor her when/if she starts therapy during the follow-up. Because the treatment variable is effectively forced to be non time-varying – as soon as it varies, the person is censored – there is no time-varying confounding and therefore no need for IPW or g-estimation to appropriately adjust for such confounding. However, the censoring itself may introduce time-dependent selection bias and IPW is therefore needed to adjust for such bias.

Comparing dynamic treatment regimes

The strategies available for comparing dynamic treatment regimes are conceptually identical to those described above for non-dynamic treatment regimes. To fix ideas, let us consider the two dynamic treatment regimes “start HAART when CD4 cell count first drops under 500 cells/ μl , then always treat” and “start HAART when CD4 cell count first drops under 200 cells/ μl , then always treat.” Our goal is to decide which of these two dynamic regimes results in longer AIDS-free survival of HIV-infected patients.

A randomized experiment to compare these two regimes could be designed as follows. First, recruit a cohort of patients infected with HIV that have never received antiretroviral therapy and whose CD4 cell count has just fallen below 500 cells/ μl for the first time. Second, assign them at random to either of the two regimes “always treat” ($A=1$) or “wait until CD4 cell count is below 200 cells/ μl , then always treat” ($A=0$). Third, administer patients in each treatment group their assigned treatment regime until the end of follow-up (say, five years), AIDS, or death, whichever occurs earlier. Finally, choose a survival analysis method to compare the mortality between the two groups. For example, use a proportional hazards model to contrast the mortality rates or an accelerated failure time model to contrast the median survival times in each group. Note that the design and analysis of the randomized experiment is similar for dynamic and non-dynamic treatment regimes. In both cases, individuals are randomized to one of two (or more) treatment strategies and then the data are analyzed by comparing the survival between the groups defined by the assigned treatment regime. The variable treatment regime is non time-varying in both cases.

An observational analysis to compare these two dynamic regimes could be designed and conducted like the one described in the previous section (non-dynamic regimes) except that a patient’s study entry is now defined as the time when her CD4 cell count first falls below 500 cells/ μl . Again, patients will not generally follow the regimes of interest and the two main approaches to deal with the multiplicity of regimes in the observational study are:

1. Modeling the causal effect of the time-varying treatment on mortality. Unlike for non-dynamic regimes, for dynamic regimes the model must include not only the time-varying treatments but also the time-varying covariates used to assign treatment in the regimes. In our example, we would need to model the time-varying treatment and its interaction with the time-varying covariate CD4 cell count. Nested structural models can naturally include these interaction terms and thus estimate the effect of dynamic treatment regimes, but marginal structural models (Joffe *et al.* 2001; van der Laan & Petersen 2004), are less useful for estimating the effect of dynamic treatment regimes.
2. Artificially censoring those who deviate from one of the two regimes of interest and use IPW estimators. We now describe how to apply the censoring/IPW approach to

compare two dynamic regimes in a cohort of HIV-infected individuals.

Application of censoring/IPW to an observational cohort

We defined a cohort comprised by the 2344 HIV-infected patients included in the French Hospital Database on HIV (FHDH) (Grabar *et al.* 2000) who had their first CD4 cell count measurement below 500 cells/ μl between 1 January 1996 and 30 June 2004, and who had never received antiretroviral therapy before that measurement. We followed these patients from their first CD4 cell count measurement below 500 cells/ μl (baseline) until a diagnosis of AIDS, death, or June 2004, whichever occurred earlier. Data on HAART use, as well as on time-dependent covariates (e.g., CD4 cell count) were recorded throughout the follow-up. Treatment A was 1 if the patient started HAART within one month of baseline, and 0 otherwise. There were 131 patients in the group $A=1$, and 2217 in the group $A=0$.

We censored patients when they initiated a non-valid regime, that is, when they did not start HAART use the first time their CD4 cell count dropped below either 500 or 200 cells/ μl . In the $A=1$ group no patients were censored because all of them initiated a valid regime. In the $A=0$ group, 655 patients were censored when they started HAART before their CD4 cell count dropped below 200 cells/ μl or when they failed to start HAART within one month of their first CD4 cell count measurement below 200 cells/ μl . A Cox model (or any other model for survival analysis) that compares the mortality between these two censored groups may not lead to causal estimates for two reasons.

First, the risk factor profile may vary between those who started and did not start HAART at baseline. For example, doctors and patients may be more likely to prescribe HAART to patients known to be infected with HIV for a long duration. If time since infection predicts a shorter survival time, independently of CD4 count, it is possible that subjects assigned to the $A=1$ group at baseline may have a worse prognosis than those assigned to $A=0$. This baseline *confounding* may be fully adjusted for (by stratification/regression or IPW) only if all baseline risk factors for mortality that predict HAART initiation were measured and recorded in the database. We adjusted for the measured confounding by including, in addition to the binary treatment regime indicator A , the following baseline (non time-varying) covariates in the Cox model: time since diagnosis of HIV infection, sex, age, risk group, and history of use of prophylaxis for *Pneumocystis carinii pneumonia*. See Appendix for details.

Second, among patients in the group $A=0$ with CD4 cell count still greater than 200 cells/ μl at time t after entry, those who initiate treatment and thus become censored at t likely will likely have, on average, lower CD4 cell counts than those who do not initiate and remain uncensored at time t . If so, individuals censored at t will have a worse prognosis than those who remain uncensored. Similarly, among those with a CD4 cell count less than 200 cells/ μl

for the first time, those who did not start HAART and thus become censored likely have, on average, a higher CD4 cell count than those who did start HAART and remained uncensored. This time-dependent *selection bias* in the artificially censored data may be fully adjusted for by IPW only if all time-varying risk factors for mortality that predict HAART initiation in the $A=0$ group were recorded. We adjusted for the measured selection bias by fitting a weighted non time-varying Cox model for mortality (including the baseline covariates listed above and the treatment indicator A) that used inverse probability weights.

The technical details of the estimation of the weights have been described elsewhere (Hernán *et al.* 2000). Briefly, the patient-time-specific weights are estimates of the inverse of the patient's probability of remaining uncensored, that is, of the probability of following her baseline regime. These probabilities were estimated by fitting a pooled logistic model in which HAART initiation was the outcome variable and that was restricted to individuals who had not yet initiated HAART use. The covariates included in the model for HAART initiation were the baseline variables listed above, treatment indicator, and time-varying CD4 cell count (for patients with treatment indicator $A=0$). We also estimated inverse probability of censoring weights in order to adjust for the potential time-dependent selection bias due to non-artificial censoring by loss to follow-up, and stabilized the two sets of weights to improve the statistical efficiency of our hazard ratio estimate. We use the robust variance estimate to compute a conservative 95% confidence interval for the hazard ratio. See the Appendix for a description of the models used in the analysis. Conceptually, IPW creates a pseudo-population in which everybody follows one of the two regimes of interest (Hernán *et al.* 2004) and the Cox model for death is fit to the pseudo-population.

Eighty-five patients developed AIDS or died during the follow-up (before censoring). The estimated hazard ratio for “treat when CD4 cell count first drops under 500 cells/ μ l” versus “treat when CD4 cell count first drops under 200 cells/ μ l” from our weighted Cox model was 0.5 (95% confidence interval: 0.2,1.1). In contrast, the standard unweighted Cox model that inappropriately adjusted for time-dependent selection bias by adding the time-varying covariate CD4 cell count yielded a hazard ratio of 0.9 (95% confidence interval: 0.4, 1.8).

Discussion

This article describes how to compare the effect on event-free survival of two (dynamic or non-dynamic) treatment regimes by using artificial censoring followed by IPW. The basic strategy can be summarized in four steps: 1) define two regimes of interest, 2) artificially censor individuals when they stop following one of the regimes of interest, 3) estimate inverse probability weights to adjust for the possibly informative censoring in the previous step, 4) compare the survival of the uncensored individuals under each re-

gime of interest in a weighted analysis adjusted for baseline covariates. The idea is to create a pseudo-population in which nobody initiates a treatment regime other than one of the two regimes of interest. We used dynamic HAART regimes as an example, but the IPW method can be applied to compare any treatment regimes using observational data.

Modeling the effect of the time-varying treatment (e.g., by g-estimation of nested structural models) is an alternative approach to censoring/IPW. Both approaches, structural modeling and censoring/IPW, require the following two assumptions: (i) no unmeasured baseline confounding and no unmeasured time-dependent selection bias (that is, data on all joint risk factors for mortality and treatment initiation were collected), (ii) no model misspecification. These assumptions are not required in randomized experiments, and can be weakened in observational studies by the use of doubly-robust estimators (Bang & Robins 2005).

The main conceptual difference between censoring/IPW and structural modeling is that censoring/IPW does not impose a structural model for the effect of treatment across regimes, and therefore the validity of the inference is more robust to the assumption of no model misspecification for the effect of treatment on survival. Censoring/IPW simply discards a patient's data after she deviates from one of the two regimes of interest, whereas modeling uses those data to estimate the mortality under the regimes of interest by borrowing information from subjects who do not always follow those regimes. For example, if one proposes a structural model in which the survival time increases (or decreases) monotonically with the duration of treatment, then the median survival time under the regime “always treated” must be greater (less) than the median survival time under any other non-dynamic regime, and the median survival time under the regime “never treat” must be less (greater) than the median survival time under any other non-dynamic regime. Therefore the data generated after some patients stop following either of the regimes of interest, “always treat” or “never treat,” help estimate the survival times of those who always follow the regimes of interest.

Because censoring/IPW does not utilize the data on the survival times of the artificially censored subjects, it is a less (statistically) efficient approach than structural modeling. This is a paradigmatic bias-variance trade-off: the censoring/IPW approach is robust to bias from misspecification of a structural model, but may be an impractical method in many applications because of poor power that results in uninformative, excessively wide, confidence intervals. On the other hand, structural models provide more efficient estimates, but these are only correct under the additional assumption that the model for the effect of treatment on the outcome is correct. When the null hypothesis of no treatment effect is true, any nested structural model is correctly specified and this additional assumption is fulfilled. Thus nested structural models can be used to provide tests of the null hypothesis of no treatment effect without concern for model misspecification. Indeed, because g-estimation does not require the use of estimated weights in the

denominator, tests of the null based on nested structural models may be more robust than those based on IPW.

To further clarify the bias-variance trade-off, consider an observational study with 100 time points (e.g., months) and a dichotomous treatment (yes, no) at each point. For simplicity, let us restrict our discussion to non-dynamic regimes. Each study participant follows one of 2^{100} possible treatment regimes. Now choose two regimes at random. The censoring/IPW approach may be of little help to contrast those two regimes because no realistic data set can be expected to include a sufficient number of patients who followed the regimes. We would need to supplement the information in the data with the a priori assumptions encoded in a structural model. However, if there exist a few specific regimes that are frequently followed by a large number of patients, censoring/IPW estimation could be used to compare these regimes. In our example, the 95% confidence interval for the weighted hazard ratio was wide, even though the FHDH is the largest database of HIV patients. Our findings suggest a beneficial effect of initiating treatment when the CD4 cell count drops below 500 versus delaying treatment until the CD4 cell count drops below 200 cells/ μl , but this issue needs to be further explored in collaborative research with other HIV cohorts.

Another relative advantage of censoring/IPW over structural modeling is that the method that can be easily implemented using standard software and minor programming although, when more than two regimes are of interest, the analysis needs to be rerun several times. For example, suppose we were interested in identifying the best strategy among the regimes “start treatment the first time that CD4 count drops below x cells/ μl ,” with x taking values between 500 and 200 in steps of 10 cells/ μl . One possible strategy would be to choose a regime of reference (say, 500) and conduct 29 analyses (i.e., estimate 29 hazard ratios) to compare 500 versus 200, 490 versus 200, 480 versus 200, and so on. The hazard ratio estimates from these 29 analyses can be inspected to determine the best time to start treatment by identifying the CD4 cell count level that led to the lowest hazard ratio. Realistically, however, sampling variability may be too large for this approach to provide an informative answer. Fitting a curve (say, a global or local polynomial smoother) to the 29 points may be a better strategy to identify the best CD4 cell count level, but then the analyst needs to propose a model across treatment regimes. More elaborate censoring/IPW methodology can be proposed to compare $k > 2$ regimes in a single analysis that incorporates the model across treatment regimes. However, this methodology is more difficult to implement with standard software.

Our observational IPW analysis had a built-in intention-to-treat flavour; although we censored patients when they initiated treatment at times not prescribed by the regimes of interest, we did not censor them when they stopped treatment. Therefore, because some do actually discontinue treatment, as in a randomized experiment with non-compliance analyzed under the ITT principle, our effect estimate is

likely to be closer to the null than the effect under continued compliance. However, IPW can be used in both randomized experiments and observational studies to estimate the effect of continued compliance with the “assigned” treatment regime by censoring any patient when she first deviates from her regime. In our example, that implies additionally censoring those in group $A=1$ when they stop treatment at any time. IPW estimates of the effect of continued compliance are only guaranteed to be unbiased under additional assumptions, regardless of whether the data came from a randomized experiment or an observational study. The additional assumptions are no unmeasured joint risk factors for treatment discontinuation and mortality, and no misspecification of the model for treatment discontinuation. In the absence of full compliance, randomized experiments with adjustment for non-compliance need to be analyzed using the same methods as observational studies. However, we note that if, as is common, discontinuation of therapy is predominantly due to severe toxicity rather than to non-medical reasons, then the effect of continued compliance will not be the effect of public health interest and the additional censoring of patients who stop therapy is undesirable.

In summary, strong assumptions are necessary in order to ensure that censoring/IPW estimates and g-estimates can be interpreted causally. At worst, these observational methods will be useful as a first-stage approach to identify promising regimes, which can then be compared in randomized experiments. 2^{100} possible regimes is too high dimensional a problem to be addressed by conducting randomized experiments, but observational analyses based on IPW and g-estimation may help reduce the dimension of the problem. The recent development of so-called “optimal regime structural nested models” (Murphy 2003; Robins 2004) based on dynamic programming may improve upon early methods, including censoring/IPW. Although dynamic programming methods are not guaranteed to find the optimal regime when the set of regimes is restricted to those of the form “start treatment the first time that CD4 count drops below x cells/ μl ,” these methods can be used to estimate the optimal regime among a much larger set of candidate regimes. This larger set includes any regimes in which the decision to start HAART at time t is a function of past CD4 history. The regimes so identified can then be selected for subsequent confirmatory face to face comparison in a randomized trial.

Appendix

Let T denote the time in months from baseline to death (or AIDS, whichever occurs earlier), A the regime assigned at baseline (1: start; 0: delay start until CD4 cell count < 200 cells/ μl), V the vector of baseline covariates listed in the text, and G an indicator of whether the baseline eligibility criteria are met (1: first time CD4 cell count < 500 μl , no prior use of antiretroviral therapy; 0: otherwise).

Consider the Cox proportional hazards model:

$$\lambda_{\tau}(t|A, V, G=1) = \lambda_0(t) \exp[\beta_1 A + \beta_2 V]$$

where $\lambda_{\tau}(t)$ is the mortality hazard (rate) at time t , $\lambda_0(t)$ is the “baseline” hazard at time t , and β_1 is the log mortality rate ratio for group $A=1$ versus $A=0$. This model assumes that β_1 is constant over time and across values of the baseline covariates. We fit this model to the person-time that results from censoring patients who initiate a treatment regime other than the two defined by A . Note that effect modification by baseline factors V can be investigated by adding interaction terms $\beta_3 VA$ to the model. The patients' contributions were weighted by the time-varying inverse probability weights that are described in the next paragraph. To overcome software limitations regarding time-varying weights in Cox models, we actually estimated the parameters β by fitting a pooled logistic regression model (Thompson 1977). We computed robust variances and thus obtained conservative 95% confidence intervals.

Let $N(t)$ denote an indicator of artificial censoring by initiating a non-valid treatment regime at t (1: censored, 0: uncensored), $C(t)$ an indicator of censoring by loss to follow-up at t (1: censored, 0: uncensored), and $L(t)$ the most recently available CD4 cell count measurement at t in cells/ μl (which we use to model CD4 cell count history from baseline to time t , $\bar{L}(t)$, in the models described below). A subject in the risk set defined at time t , was weighted by the inverse probability weight $SW^N \times SW^C$. The weight SW^N is defined as

$$SW^N(t) = \prod_{k=0}^t \frac{\Pr[N(k)=0|C(k)=0, N(k-1)=0, A, V, G=1]}{\Pr[N(k)=0|\bar{L}(k), C(k)=0, N(k-1)=0, A, V, G=1]}$$

where the denominator is, informally, the subject's probability of remaining uncensored by deviation from protocol through time t given baseline and time-varying confounders. We estimated this probability by fitting a pooled logistic model, as previously described (Hernán *et al.* 2000 & 2001). The weight SW^C is defined as

$$SW^C(t) = \prod_{k=0}^t \frac{\Pr[C(k+1)=0|N(k)=0, C(k)=0, A, V, G=1]}{\Pr[C(k+1)=0|N(k)=0|\bar{L}(k), C(k)=0, A, V, G=1]}$$

where the denominator is, informally, the subject's probability of remaining uncensored by loss to follow-up through time t given baseline and time-varying confounders. We also estimated this probability by fitting a pooled logistic model (Hernán *et al.* 2001). The numerators of both weights serve a stabilizing purpose to reduce the variance of the estimate of the parameters β , and are also estimated by fitting pooled logistic models.

Finally, for comparison purposes, we also estimated the parameters of the unweighted Cox proportional hazards model

$$\lambda_{\tau}(t|A, V, G=1) = \lambda_0(t) \exp[\beta_1 A + \beta_2 V + \beta_3 L(t)]$$

that includes the time-varying confounder CD4 cell count as a covariate.

Acknowledgements

Miguel Hernán was partly supported by NIH grant R01 HL080644-01. James Robins was partly supported by NIH grant R01 AI32475.

References

- Bang, H. & J. M. Robins: Doubly-robust estimation in missing data and causal inference models. *Biometrics* 2005, **61**, 962–973.
- Cole, S. R., M. A. Hernán, J. M. Robins, K. Anastos, J. Chmiel, R. Detels, C. Ervin, J. Feldman, R. Greenblatt, L. Kingsley, S. Lai, M. Young, M. Cohen & A. Munoz: Effect of highly active antiretroviral therapy on time to acquired immunodeficiency syndrome or death using marginal structural models. *Amer. J. Epidemiol.* 2003, **158**, 687–694.
- Grabar, S., V. Le Moing, C. Goujard, C. Leport, M. D. Kazatchkine, D. Costagliola & L. Weiss: Clinical outcome of patients with HIV-1 infection according to immunologic and virologic response after 6 months of highly active antiretroviral therapy. *Ann. Intern. Med.* 2000, **133**, 471–473.
- Hammer, S. M., K. E. Squires, M. D. Hughes, J. M. Grimes, L. M. Demeter, J. S. Currier, J. J. Eron Jr., J. E. Feinberg, H. H. Balfour Jr., L. R. Deyton, J. A. Chodakewitz & M. A. Fischl: A controlled trial of two nucleoside analogues plus indinavir in persons with human immunodeficiency virus infection and CD4 cell counts of 200 per cubic millimeter or less. AIDS Clinical Trials Group 320 Study Team. *New Engl. J. Med.* 1997, **337**, 725–733.
- Hernán, M. A., B. Brumback & J. M. Robins: Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* 2000, **11**, 561–570.
- Hernán, M. A., B. Brumback & J. M. Robins: Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *J. Amer. Statist. Ass.* 2001, **96**, 440–448.
- Hernán, M. A., S. R. Cole, J. Margolick, M. Cohen & J. M. Robins: Structural accelerated failure time models for survival analysis in studies with time-varying treatments. *Pharmacoepidemiol. Drug Saf.* 2005, **14**, 477–491.
- Hernán, M. A., S. Hernández-Díaz & J. M. Robins: A structural approach to selection bias. *Epidemiology* 2004, **15**, 615–625.
- Joffe, M. M., J. Santanna & H. Feldman: Partially marginal structural models for causal inference [abstract]. *Amer. J. Epidemiol.* 2001, **153**, S261.
- Murphy, S. A.: Optimal dynamic treatment regimes (with discussion). *J. Royal Statist. Soc. Series B* 2003, **65**, 331–366.
- Robins, J. M.: Analytic methods for estimating HIV treatment and cofactor effects. In: *Methodological issues of AIDS Mental Health Research*. Eds.: D. G. Ostrow & R. Kessler. Plenum Publishing, New York, 1993, pp. 213–290.
- Robins, J. M.: Optimal structural nested models for optimal sequential decisions. In: *Proceedings of the Second Seattle Symposium in Biostatistics: Analysis of Correlated Data*. Eds: D. Y. Lin & P. Heagarty. Springer, Seattle, 2004.
- Sterne, J. A., M. A. Hernán, B. Ledergerber, K. Tilling, R. Weber, P. Sendi, M. Rickenbach, J. M. Robins, M. Egger & Swiss HIV Cohort Study: Long-term effectiveness of potent antiretroviral therapy in preventing AIDS and death: A prospective cohort study. *Lancet* 2005, **366**, 378–384.
- Thompson, W. A., Jr.: On the treatment of grouped observations in life studies. *Biometrics* 1977, **33**, 463–470.
- van der Laan, M. J. & M. L. Petersen: History-adjusted marginal structural models and statically-optimal dynamic treatment regimes. U.C. Berkeley Division of Biostatistics Working Paper Series. Working Paper 158, 2004. <http://www.bepress.com/ucbbiostat/paper158>