

Assessing the Effectiveness of Antiretroviral Adherence Interventions

Using Marginal Structural Models to Replicate the Findings of Randomized Controlled Trials

Maya L. Petersen, MS,* Yue Wang, MS,* Mark J. van der Laan, PhD,* and David R. Bangsberg, MD, MPH†‡

Summary: Randomized controlled trials of interventions to improve adherence to antiretroviral medications are not always feasible. Marginal structural models (MSM) are a statistical methodology that aims to replicate the findings of randomized controlled trials using observational data. Under the assumption of no unmeasured confounders, 3 MSM estimators are available to estimate the causal effect of an intervention. Two of these estimators, G-computation and inverse probability of treatment weighted (IPTW), can be implemented using standard software. G-computation relies on fitting a multivariable regression of adherence on the intervention and confounders. Thus, it is related to the standard multivariable regression approach to estimating causal effects. In contrast, IPTW relies on fitting a multivariable logistic regression of the intervention on confounders. This article reviews the implementation of these methods, the assumptions underlying them, and interpretation of results. Findings are illustrated with a theoretic data example in which MSM are used to estimate the effect of a behavioral intervention on adherence to antiretroviral therapy.

Key Words: causal inference, counterfactual, HIV, G-computation, inverse probability of treatment, randomized controlled trial

(*J Acquir Immune Defic Syndr* 2006;43:S96–S103)

Randomized controlled trials remain the “gold standard” for assessing the effectiveness of interventions aimed at improving adherence to antiretroviral medications. Such trials may be impractical to implement for many potential

interventions, however. Generally, a separate trial is required to investigate each intervention, making the evaluation of multiple interventions costly and time-consuming. In addition, randomized trials measure the efficacy of an intervention applied in a controlled setting, which is not always representative of the effectiveness of the same intervention applied in clinical practice.

In many cases, an intervention can become the standard of care despite the absence of compelling evidence demonstrating its efficacy. For such practices, it may be considered unethical to assign patients to the control arm of a trial. For example, randomized controlled trials have provided equivocal evidence that behavioral interventions (eg, case management, pharmacist-based education, and psychoeducational interventions) improve adherence to antiretroviral therapy.^{1–4} Such interventions are increasingly considered the standard of care, making additional randomized trials less likely.

Observational data offer a rich alternative for estimating the effectiveness of adherence interventions. The same clinical cohort can be used to investigate the potential of several proposed interventions. As in randomized trials, accurate monitoring of the adherence outcome is essential. In addition, as we discuss in this article, estimation of causal effects from observational data further requires accurate measurement of all confounders.

Estimation of causal effects, whether using randomized or observational data, requires sufficient experimentation in the assignment of the intervention of interest. If all patients in the study population receive an intervention of interest, then, clearly, no information is available in the data to estimate the effect of the intervention. In other words, estimation of causal effects is not possible without a control series of some sort. Thus, observational data provide an alternative approach to estimating the causal effects of interventions that are considered “standard care,” and thus not amenable to study using randomized trials, but only if such interventions are not, in fact, practiced uniformly by all practitioners. For example, despite growing consensus in favor of behavioral adherence interventions, the use of such interventions remains far from ubiquitous because of a range of reasons, including availability of the interventions and physician training and beliefs. Thus, observational data could be used to study the effects of these interventions.

From the *Division of Biostatistics, University of California at Berkeley, Berkeley, CA; †Epidemiology and Prevention Interventions Center, Division of Infectious Diseases, San Francisco General Hospital, San Francisco, CA; and ‡Positive Health Program, San Francisco General Hospital, University of California at San Francisco, San Francisco, CA. Funded by National Institute of Mental Health (NIMH) grants 54907 and 63011. D. R. Bangsberg received additional funding from The Doris Duke Charitable Foundation. M. L. Petersen is supported by a Pre-Doctoral Fellowship from the Howard Hughes Medical Institute. M. J. van der Laan is supported by NIMH grant R01 GM071397.

Reprints: Maya L. Petersen, MS, Division of Biostatistics, School of Public Health, University of California at Berkeley, Earl Warren Hall, Room 7360, Berkeley, CA 94720–7360 (e-mail: mayaliv@berkeley.edu).

Copyright © 2006 by Lippincott Williams & Wilkins

Marginal structural models (MSM) are a statistical methodology that aims to replicate the findings of a randomized controlled trial using observational data.^{5,6} These models, developed by Robins, estimate the difference in mean adherence that would have been observed if the entire cohort had received an intervention versus if the entire cohort had not received the intervention. Such theoretic outcomes are referred to as “counterfactuals.” This article provides a practical introduction to MSM, discussing the assumptions behind these models, their implementation, and the interpretation of results. The focus is on 2 MSM estimators, G-computation and inverse probability of treatment weighted (IPTW), which can be implemented using standard software. Concepts are illustrated with a theoretic example of a behavioral adherence intervention, attendance at a pharmacist-based adherence clinic. In such an example, MSM can be used to estimate the effect of clinic attendance on adherence to antiretrovirals in an observational clinical cohort.

The article begins by describing the hypothetical data used to estimate the effect of pharmacist-based adherence clinic attendance on adherence. Next, it provides a nontechnical introduction to the counterfactual framework for causal inference and uses this framework to review the assumptions required to estimate causal effects based on observational data. The G-computation estimator is introduced, and its relation to standard multivariable regression is discussed. Next, the IPTW estimator is presented. Implementation of both estimators is illustrated using data from a hypothetical clinical cohort that can be worked through using a calculator or spreadsheet (R code for implementing these estimators is provided in the Appendix). Finally, additional assumptions required by the alternative estimators are compared, and the advantages of using MSM versus more common analytic approaches are highlighted.

EFFECT OF PHARMACIST-BASED ADHERENCE EDUCATION ON ADHERENCE

The impact of a behavioral intervention, such as a pharmacist-based educational adherence clinic, on antiretroviral adherence could be assessed by determining participation in one or several clinics over time and measuring adherence longitudinally using monthly pill counts. The adherence measure would be defined as the difference between the current and previous pill counts divided by the prescribed number of doses for the same period.⁷ Participation in the clinic could be determined by patient and pharmacist interview. For ease of exposition, we assume that the clinics of interest are held at the beginning of each month. Data on multiple potential confounders, including disease stage, recreational drug use (crack, intravenous drug, and alcohol use), depression, housing status, age, gender, ethnicity, and housing status would also be collected.

To address the question “How does clinic participation at the beginning of a month affect adherence during the next 30 days?” a data set would be created that consisted of a data point for each person-month during follow-up for which clinic attendance and subsequent adherence were measured. Thus, each person could contribute several data points. The observed

data for a given person-month would consist of a binary intervention (A = clinic attendance at the beginning of the month), a continuous outcome (Y = adherence for the month), and a set of covariates (W).

COUNTERFACTUALS AND CAUSAL INFERENCE

The causal effect of an intervention can be defined using the concept of counterfactuals. A counterfactual outcome, Y_a , is defined as the outcome an individual would have had under a specific intervention, a . Thus, 2 counterfactual outcomes exist for each person-month in the study: Y_1 is the adherence that would have been observed for that month if the subject had participated in the clinic at the beginning of the month ($a = 1$), and Y_0 is the adherence that would have been observed over the month if the subject had not participated in the clinic ($a = 0$). These outcomes are termed *counterfactual* because only 1 outcome is observed for a given person-month: if the subject participated in the clinic, Y_1 is observed, and if the subject did not participate, Y_0 is observed.

The causal effect of an intervention for a given person-month is defined as $Y_1 - Y_0$, or the difference in the counterfactual outcome if the individual had received the intervention that month versus the outcome if the same individual had not received the intervention. The causal effect in the cohort is simply the mean of these data point-specific effects: $E(Y_1 - Y_0)$. If we observed each patient's adherence with and without clinic participation for each month, we could simply compare these outcomes to estimate the causal effect of the intervention. We only ever observe 1 of these counterfactual outcomes for a given person-month, however. Thus, the counterfactual framework turns the problem of estimating causal effects into a problem of missing data.

Counterfactuals illustrate why randomized controlled trials can be used to estimate causal effects. In a randomized trial, random assignment to clinic participation or not ensures that members of both intervention groups are representative samples of the study population. As a result, the observed adherence among people who participate in the clinic is representative of the counterfactual adherence if everyone in the study had participated. Similarly, the adherence among people who do not participate in the clinic is representative of the counterfactual adherence if everyone in the study had not participated. Thus, the difference in mean adherence observed between 2 randomized intervention groups is equivalent to the difference in counterfactual outcomes, or the causal effect of mean intervention: $E(Y|A = 1) - E(Y|A = 0) = E(Y_1 - Y_0)$, where $E(Y|A = 1)$ is the mean adherence among the random set of subjects who attended the adherence clinic and $E(Y|A = 0)$ is the mean adherence among the random set of subjects who did not attend.

CHALLENGE OF OBSERVATIONAL DATA

Counterfactuals further illustrate the challenges of estimating causal effects using observational data. When intervention status is not assigned randomly, members of a treatment group are unlikely to be representative of the study population. Thus, observed adherence among people who

attended a pharmacist-based adherence clinic is generally not representative of the counterfactual adherence that would have been observed if everyone in the study had attended the clinic. Consider a data example based on participation in a pharmacist-based adherence clinic but with prior adherence as the single confounder. For simplicity of calculations, we consider prior adherence as a binary variable: $W = 1$ if prior adherence was $<95\%$ (“low prior adherence”) and $W = 0$ if prior adherence was $>95\%$ (“high prior adherence”). Table 1 gives hypothetical data for this example. In the hypothetical data, individuals who had low prior adherence were more likely to attend the clinic (these individuals were more likely to be referred to the clinic by clinicians concerned about their adherence levels) (odds ratio = 4.0) and also had lower average adherence after clinic attendance (on average 45% lower). In this example, we expect that the mean observed adherence among the nonrandom subgroup of people who attended the clinic, $E(Y|A = 1)$, underestimates the mean counterfactual adherence if everyone in the study had attended the clinic, $E(Y_1)$; thus, $E(Y|A = 1) - E(Y|A = 0) \neq E(Y_1 - Y_0)$.

Prior adherence in this example illustrates the problem of confounding in observational data. In general, confounding occurs when a common cause affects receipt of an intervention and the adherence outcome (Fig. 1). Unlike the simple example presented previously, however, the presence of multiple confounders often makes the direction of the resulting bias in effect estimates unpredictable.

To ensure the identifiability of causal effects from observational data, a key assumption is needed. MSM assume that, within strata defined by all measured covariates, intervention assignment is randomized. In other words, among individuals who are identical with respect to all measured covariates, the observed adherence of individuals who received the intervention is representative of counterfactual adherence under the receipt of the intervention for that stratum (and, the observed adherence of individuals who did not receive the intervention is representative of counterfactual adherence under no intervention for that stratum). This assumption, known as the randomization assumption, or assumption of no unmeasured confounders, can be stated as $A \perp Y_a | W$ (ie, the intervention is independent of counterfactual outcomes given measured covariates). It is not an assumption that is testable using the data but, rather, relies on the background knowledge of the investigator.


TABLE 1. Hypothetic Observed Data for Simplified Example

ID	Intervention (A)	Prior Adherence (W)	Adherence (Y)
1	1	1	0.7
2	0	0	0.8
3	1	1	0.4
4	1	0	1
5	0	1	0.4
6	0	0	0.7

Six subjects contribute 1 record each. We are interested in estimating the effect of an adherence intervention (A) on adherence (Y). Assume a single confounder, $W =$ prior adherence ($W = 1$ for prior adherence $<95\%$, $W = 0$ for prior adherence $\geq 95\%$). The behavioral adherence intervention was associated with 6.7% higher adherence (before adjusting for confounding).

Effect of interest: -----

A. General causal structure for confounding



B. Example of confounding of the effect of a behavioral intervention on future adherence by prior adherence




FIGURE 1. Confounding of the effect of exposure on outcome.

When considering whether a proposed set of covariates, W , is sufficient to ensure that the randomization assumption holds, it is crucial to give close attention to the issue of temporal ordering. Simply put, confounding arises because of covariates that affect intervention assignment; thus, confounders must occur before rather than after intervention assignment. Inclusion in W of covariates that occur after intervention assignment, and that are affected by the intervention, can bias estimates of effect.⁵ Thus, in the adherence clinic analyses, potential confounders, W , for a given person-month could include non-time-varying covariates (eg, gender, ethnicity) and time-varying covariates measured before the decision to refer a patient to the adherence clinic for that month (eg, prior adherence, CD4 T-cell count, drug use).

Under the randomization assumption, MSM aim to use statistical methods to replicate the results that would have been observed in a randomized controlled trial. Several MSM approaches are available to estimate this causal effect. Specifically, there are 3 MSM estimators: G-computation, IPTW, and double-robust (DR).⁸⁻¹⁰ These estimators rely on distinct models and assumptions to estimate the same causal effect.

MARGINAL STRUCTURAL MODELS: G-COMPUTATION ESTIMATOR

The G-computation estimator (when implemented in a point treatment setting, as described here) relies on the same modeling approach as standard multivariable regression, a commonly used method for estimating causal effects using observational data. When the outcome of interest is continuous and the intervention does not interact with covariates to affect the outcome, both approaches estimate the results that would have been seen in a randomized controlled trial. In nonlinear models (eg, commonly used logistic regression) and in the presence of intervention-covariate interactions, however, standard multivariable regression estimates an adjusted, or conditional, causal effect, whereas G-computation estimates the marginal

causal effect, or the results that would have been seen in a randomized trial. To see how these effects differ, consider the following example.

The first step in implementing the G-computation estimator is to fit a multivariable regression model of the outcome, given the intervention and all covariates, $\hat{E}(Y|A, W)$. In our hypothetical example, adherence is regressed on clinic participation and possible confounders. This step corresponds to the standard multivariable regression approach.

In G-computation, however, this regression model is then used to predict 2 counterfactual outcomes for each person-month, given covariate values for that month (W): the counterfactual adherence if the individual had participated in the adherence clinic at the beginning of the month, $\hat{Y}_1 = \hat{E}(Y|A = 1, W)$, and the counterfactual adherence if the individual had not participated in the clinic at the beginning of the month, $\hat{Y}_0 = \hat{E}(Y|A = 0, W)$. A new data set is constructed that contains the predicted counterfactual outcome for each person-month in the presence and absence of the intervention; thus, the new data set contains twice the number of rows as the initial data.

This process is analogous to running an ideal experiment, in which the investigator first assigns each individual in the cohort to attend the clinic and observes the resulting adherence and then assigns the identical cohort to not attend the clinic and observes the resulting adherence. Instead, in implementing the G-computation estimator, the investigator sets clinic attendance equal to 1 in the regression model and records the predicted adherence for all person-months and then sets clinic attendance equal to 0 and records the predicted adherence for all person-months. In the newly constructed data set, these predicted counterfactual outcomes are then regressed on the intervention. Figure 2 illustrates implementation of the G-computation process using the hypothetical data presented in Table 1; corresponding R code is provided in the Appendix.

In this example, G-computation estimates that clinic participation by a random sample of the population would be expected to increase adherence in the subsequent month by 20% compared with that of a random control group that did not attend the clinic. In contrast, multivariable regression estimates a conditional causal effect of 25%–10% prior low adherence. In other words, standard multivariable regression does not provide an effect estimate for the entire study population but, rather, provides an effect estimate that differs depending on an individual's prior adherence. Because a linear model was used in this example, in this case, the G-computation estimate is equivalent to taking a weighted average of the conditional causal effect estimate from the multivariable regression model, weighted with respect to the distribution of prior adherence in the study population.

MARGINAL STRUCTURAL MODELS: INVERSE PROBABILITY OF TREATMENT-WEIGHTED ESTIMATOR

The IPTW estimator controls for confounding using an approach that does not depend on fitting a multivariable regression of adherence on intervention and confounders. Instead, the IPTW estimator recognizes that confounding can

be viewed as a problem of biased sampling. If an intervention were assigned randomly, covariates would have the same distribution in the intervention and control groups. To return to our earlier example, if clinic participation had been randomized, people with low prior adherence would occur with the same expected frequency among those who did and did not attend the adherence clinic. In this example, however, individuals with low prior adherence are overrepresented among those who attended the clinic and underrepresented among those who did not attend.

The IPTW estimator aims to create a reweighted data set in which the intervention is randomized. To accomplish this, individuals are assigned larger weights if their observed intervention status is rare given their covariates, and are assigned smaller weights if their observed intervention status is common given their covariates. In our simple example, individuals with high prior adherence who attended the clinic get larger weights and individuals with high prior adherence who did not attend the clinic get smaller weights. In the reweighted sample, prior adherence is distributed evenly among subjects who did and did not attend the clinic.

Implementation of the IPTW estimator begins with fitting a multivariable regression model of the probability of receiving an intervention, given covariates. This model, called the treatment mechanism, can be written as $g(A|W) = P(A = a|W)$. In our example, logistic regression of clinic attendance on covariates can be used to model the treatment mechanism. The model of the treatment mechanism is then used to predict each individual's probability of receiving his or her observed intervention. Subjects are assigned weights equal to the inverse of this predicted probability. As a result, individuals with underrepresented intervention status, given their covariates, get larger weights. For example, person-months for which subjects attended the adherence clinic are assigned weights inversely proportional to the probability of attending the clinic, given covariate values for those person-months ($W = w$): $\hat{P}(A = 1|W)$. Similarly, person-months for which individuals did not attend the clinic are assigned weights inversely proportional to the probability of not attending, given covariate values for those person-months (W): $1 - \hat{P}(A = 1|W)$. Once the weights for each person-month have been assigned, the IPTW estimator is calculated using standard weighted least squares regression of adherence on intervention. Figure 3 illustrates IPTW implementation using the hypothetical data presented in Table 1; corresponding R code is provided in the Appendix.

Implementation of IPTW in the hypothetical data example estimates that clinic attendance by a random sample of the population would be expected to increase adherence by 20% compared with that of a random control group that did not participate in the clinic.

COMPARISON OF MARGINAL STRUCTURAL MODEL ESTIMATORS

In addition to the randomization assumption, the consistency of the MSM estimators relies on correct specification of the models used to control confounding. The G-computation estimator relies on correct specification of the

1. Fit a linear regression model of observed adherence given the adherence intervention and prior adherence.

$$E(Y | A, W) = 0.75 + 0.25 \times \text{Intervention} - 0.35 \times \text{Prior Adherence} - 0.1 \times \text{Intervention} \times \text{Prior Adherence}$$

$$E(Y | A = 1, W = 1) = 0.55$$

$$E(Y | A = 0, W = 1) = 0.4$$

$$E(Y | A = 1, W = 0) = 1.0$$

$$E(Y | A = 0, W = 0) = 0.75$$

2. Use this linear regression model to predict counterfactual adherence for each subject in the presence and absence of the intervention.

ID	Intervention (a)	Counterfactual Adherence (Ya)
1	0	0.4
2	0	0.75
3	0	0.4
4	0	0.75
5	0	0.4
6	0	0.75
1	1	0.55
2	1	1.0
3	1	0.55
4	1	1.0
5	1	0.55
6	1	1.0

3. Regress the counterfactual adherence on the intervention in new dataset.

The behavioral intervention improves adherence 20%

$$E(Y_1 - Y_0) = 0.20$$

FIGURE 2. Implementation of the G-computation estimator using hypothetical data from Table 1.

model regressing adherence on intervention and covariates, $E [Y | A, W]$, whereas the IPTW estimator relies on correct specification of the model for the treatment mechanism, $g (A | W)$. In the hypothetical data example, the multivariable regression model of adherence on intervention and confounders and the model of the treatment mechanism are perfectly specified (in fact, in this simple example, no parametric models are needed), yielding identical results between the estimators. In practice, however, model selection can be a challenging issue for either approach, particularly when the set of potential confounders is large. One approach is to use an aggressive search algorithm and cross-validation when fitting these models.

Beyond the general requirement for sufficient experimentation required by any attempt to estimate causal effects, the IPTW estimator, in particular, relies on the assumption that confounders do not perfectly predict intervention assignment. This assumption, called experimental treatment assignment (ETA), requires experimentation in the use of the intervention within every stratum of confounders. The assumption can be stated as follows: $P(A = a | W) > 0$ for all covariate values W , for each possible level of the intervention.

For example, to estimate of the effect of adherence clinic attendance, IPTW would require that each subject in the study have some positive probability of attending and not attending

the clinic, regardless of the values of his or her confounding variables. This requirement stems from the fact that if, for example, no individuals with high prior adherence actually participated in the adherence clinic, the data contain no information about adherence after clinic participation among these individuals. If the ETA assumption is violated, the G-computation approach can still be used; however, it relies on extrapolation.

Note that the assumption of ETA is not the same as requiring that some subjects in each stratum receive no adherence intervention at all. The control or reference level of the intervention is determined by the investigator. For example, a realistic question that might be addressed using MSM is whether a targeted behavioral intervention (eg, attendance at a pharmacist-based adherence clinic) improved adherence as compared to standard provider-based counseling.

The IPTW and G-computation estimators rely on consistent estimation of the model used to control confounding. An alternative DR MSM estimator is also available that uses the model of $E (Y | A, W)$ and the model of the treatment mechanism but remains consistent if either model is correctly specified.^{8,11} Thus, the DR estimator is maximally robust to model misspecification. Unlike G-computation and IPTW, however, the DR estimator requires nonstandard software to implement.

1. Fit a logistic regression model of the treatment mechanism (probability of receiving the intervention given confounders): $P(A = 1 | W) = \frac{1}{1 + \exp(-(-0.6931 - 1.3863 \times W))}$

2. Use this model to predict the probability of receiving the intervention for subjects that received it, and the probability of not receiving the intervention for subjects that did not receive it.

$$P(A = 1 | W = 1) = 0.67$$

$$P(A = 0 | W = 1) = 1 - 0.67 = 0.33$$

$$P(A = 1 | W = 0) = 0.33$$

$$P(A = 0 | W = 0) = 1 - 0.33 = 0.67$$

ID	Intervention (A)	Adherence (Y)	Weights = 1/Predicted
			Probability of Observed Intervention
1	1	0.7	1.5
2	0	0.8	1.5
3	1	0.4	1.5
4	1	1.0	3.0
5	0	0.4	3.0
6	0	0.7	1.5

3. Regress the observed adherence on the observed intervention status, using weights. The behavioral intervention improves adherence 20%

$$E(Y_1 - Y_0) = 0.20$$

FIGURE 3. Implementation of the IPTW estimator using hypothetical data from Table 1.

DISCUSSION

Although the G-computation and IPTW estimators presented can be implemented using standard software, the standard error estimates and resulting probability values provided by the software are not accurate. A relatively simple approach to constructing accurate confidence intervals is to use a nonparametric bootstrap. In applying this method, 500 (or more) bootstrap samples are constructed by resampling (with replacement) from the study population. If subjects can contribute more than 1 data point, sampling must be based on subject rather than on data point. In each bootstrap sample, the models for $E(Y | A, W)$ and the treatment mechanism are refit, and the G-computation and IPTW estimators are recalculated. The variability of the estimators across bootstrap samples provides an estimate of their standard errors.

MSM offer several advantages over standard analytic approaches to estimating causal effects. Use of different estimators, which rely on distinct models to control for confounding, can improve robustness to model misspecification. In addition, MSM allow the investigator to focus on the question of interest. If the aim is to replicate the findings of a randomized trial, the investigator may not be interested in the effect of an intervention conditional on all confounders, as is estimated by standard multivariable regression. For example, the investigator might fit a standard multivariable regression model, in which the effect of a behavioral intervention varied according to a subject's age, gender, and current CD4 T-cell

count. Such effect modification may, of course, be of interest in itself and would be revealed when fitting the model of $E(Y | A, W)$ that underlies the G-computation estimator. MSM allow the investigator to go further, however, and estimate the difference in average adherence that would have been observed if the entire study population had been randomized to receive the behavioral intervention of interest versus the control level of intervention.

In addition, in settings with longitudinal treatments, MSM are often the only valid analytic approach to the estimation of causal effects. This is true when time-dependent confounding is present or, in other words, when the effect of future treatment is confounded by covariates that are themselves affected by past treatment. In this setting, the coefficients from a standard multivariable regression generally do not have a clear causal interpretation.⁵

For example, the effect of attendance at a pharmacist-based adherence clinic on adherence may be cumulative over time. Attendance at a single clinic may result in little or no immediate effect, while repeated clinic attendance significantly improves adherence. Researchers might wish to estimate the effect of duration or frequency of attendance at a pharmacist-based adherence clinic on average adherence over a 6-month period after cohort enrollment, or perhaps the effect of cumulative exposure to adherence clinic education on time to virologic failure. In such a setting, it is quite possible for providers to assess their patients' adherence over time and be

more likely to enroll in the clinic those patients whose adherence declines over the course of follow-up. Thus, the adherence of patients over the course of follow-up can result in confounding by indication. Traditional multivariable regression does not allow the researcher to adjust for such confounding; including adherence measured during follow-up in the multivariable regression model would amount to adjusting away part of the effect of interest. In contrast, the MSM estimators presented here could be used to provide an estimate of the effect of cumulative participation in the adherence clinic while adjusting for this time-dependent confounding (the implementation of the estimators in this time-dependent setting would differ slightly from the point treatment implementation illustrated here).

This article has outlined how MSM can be implemented to estimate the causal effect of a proposed adherence intervention. The example estimates the effect of a single binary point treatment on a continuous outcome. The same MSM methods can be applied to a wide range of interventions, however, including continuous and multicomponent interventions, and to a wide range of outcomes, including time-to-event or binary outcomes. As a result, MSM provide a powerful and broadly applicable tool for causal inference.

REFERENCES

- Rathbun RC, Farmer KC, Stephens JR, et al. Impact of an adherence clinic on behavioral outcomes and virologic response in treatment of HIV infection: a prospective, randomized, controlled pilot study. *Clin Ther*. 2005;27:199–209.
- Rawlings MK, Thompson MA, Farthing CF, et al. Impact of an educational program on efficacy and adherence with a twice-daily lamivudine/zidovudine/abacavir regimen in underrepresented HIV-infected patients. *J Acquir Immune Defic Syndr*. 2003;34:174–183.
- Tuldra A, Fumaz CR, Ferrer MJ, et al. Prospective randomized 2-arm controlled study to determine the efficacy of a specific intervention to improve long-term adherence to highly active antiretroviral therapy. *J Acquir Immune Defic Syndr*. 2000;25:221–228.
- Milam J, Richardson JL, McCutchan A, et al. Effect of a brief antiretroviral adherence intervention delivered by HIV care providers. *J Acquir Immune Defic Syndr*. 2005;40:356–363.
- Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology*. 2000;11:550–560.
- Hernan MA, Brumback BA, Robins JM. Estimating the causal effect of zidovudine on CD4 count with a marginal structural model for repeated measures. *Stat Med*. 2002;21:1689–1709.
- Bangsberg DR, Hecht FM, Charlebois ED, et al. Adherence to protease inhibitors, HIV-1 viral load, and development of drug resistance in an indigent population. *AIDS*. 2000;14:357–366.
- Robins JM. Robust estimation in sequentially ignorable missing data and causal inference models. In: *1999 Proceedings of the American Statistical Association*. American Statistical Association, Alexandria, VA. 2000:6–10.
- Robins JM. Marginal structural models versus structural nested models as tools for causal inference. In: Halloran ME, Berry D, eds. *Statistical Models in Epidemiology, the Environment, and Clinical Trials*. New York: Springer-Verlag; 1999:95–134.
- Robins JM. A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *J Chron Dis*. 1987;40:139S–161S.
- Neugebauer R, van der Laan M. Why prefer double robust estimates? Illustration with causal point treatment studies. In: *University of California, Berkeley, Division of Biostatistics Working Paper Series*. Berkeley: University of California; 2002. Available at: <http://www.bepress.com/ucbbiostat/paper115/>.

APPENDIX

R CODE FOR GENERATING HYPOTHETICAL DATA AND IMPLEMENTING IPTW AND G-COMPUTATION ESTIMATORS

```
#Generate data (Table 1)
```

```
#Confounding of the effect of the behavioral intervention on adherence by prior adherence
```

```
AWY<-
```

```
rbind (c(1,1,1,0.7), c(2,0,0,0.8), c(3,1,1,0.4), c(4,1,0,1.0), c(5,0,1,0.4), c(6,0,0,0.7))
```

```
AWY<-as.data.frame (AWY)
```

```
names (AWY) <-c("ID", "Intervention", "Prior.Adherence", "Adherence")
```

```
#Look at crude (unadjusted) associated between the intervention and adherence
```

```
glm (Adherence~Intervention, data = AWY)
```

```
#Implement IPTW Estimator (Figure 2)
```

```
#1. Model the treatment mechanism
```

```
rx.mech<-glm (Intervention~Prior.Adherence, family = binomial, data = AWY)
```

```
rx.mech
```

```
#2. Fit each individual's predicted probability of receiving the intervention he in fact received
```

```
AWY$PAW<-if else (AWY$Intervention == 1, predict.glm (rx.mech, type = "response"), 1-predict.glm (rx.mech, type = "response"))
```

```
#weight by the inverse of these predicted probabilities AWY$w<-1/AWY$PAW
```

```
#3. Fit standard weighted least squares regression glm (Adherence~Intervention, weight = w, data = AWY)
```

#Implement G-computation Estimator (Figure 3)

#1. Model $E(Y|A,W)$

```
predY<-glm (Adherence~Intervention*Prior.Adherence, data = AWY)
```

```
predY
```

#2. For each record predict counterfactual outcome in presence and absence of the intervention

```
n<-dim(AWY) [1]
```

```
y0<-predict.glm(predY,newdata =
```

```
data.frame (Prior.Adherence = AWY$Prior.Adherence, Intervention = rep(0,n)))
```

```
y1<-predict.glm(predY,
```

```
newdata = data.frame (Prior.Adherence = AWY$Prior.Adherence,
```

```
Intervention = rep(1,n)))
```

#3. Make new dataset with counterfactual treatments and outcomes

```
y.pool <- c(y0,y1)
```

```
a.pool <- c(rep(0,n),rep(1,n))
```

```
gcomp<-data.frame("ID" = c(1:6,1:6), "Intervention" = a.pool,
```

```
"Counterfactual.adherence" = y.pool)
```

#4. Fit standard regression in new dataset of counterfactuals

```
glm(Counterfactual.adherence~Intervention, data = gcomp)
```