Technical Report No. 3 May 6, 2013

Marginal structural modeling in health services research

Kyung Min Lee
kml@bu.edu

Boston University School of Public Health
Department of Health Policy & Management

This paper was published in fulfillment of the requirements for PM931 Directed Study in Health Policy and Management under Professor Cindy Christiansen's (cindylc@bu.edu) direction. Michal Horný, Jake Morgan, Marina Soley-Bori, and Meng-Yun Lin provided helpful reviews and comments.

Contents

Executive Summary

Statistical inferences from observational studies are often subject to confounding caused by both observed and unobserved confounding variables. Conventional methods for controlling for confounding include applying statistical techniques such as stratification and multivariable regression analysis. In the presence of time-dependent confounders, however, such techniques may still lead to biased estimates. Marginal structural modeling (MSM) uses a multi-step estimation strategy to separate confounding control from the estimation of the parameters of interest, allowing the investigator to obtain unbiased estimates. Given that there are no unmeasured confounders and the probability of treatment is positive, the estimates of a marginal structural model can be interpreted as causal. This report serves as a starting point for researchers who wish to use MSM in their studies, providing an overview of the theory behind MSM and a guidance for its implementation.

Keywords: marginal structural models, causal inference, time-dependent confounders, weighted regression, counterfactuals

1 Introduction

A causal relationship between a treatment and its associated outcome variable becomes ambiguous in the presence of a confounder; the treatment effect is confounded when one or more risk factors for the outcome are also correlated with the treatment. Randomization in clinical trials eliminates or reduces most of such confounding by randomly assigning each subject into either the control or treatment group, thereby allowing the investigator to measure the true causal effect of the treatment. Observational studies, in which randomization cannot be performed, typically address confounding by applying statistical techniques such as stratification and multivariable regression analysis. For a point treatment study in which the treatment is administered once, multivariable regression models may be sufficient to control for confounding variables. However, in longitudinal studies with repeated treatments over time, the estimates from regression models may still be biased if (1) there exists a time-dependent covariate that predicts subsequent treatment and is an independent predictor of the outcome, and (2) past treatment history predicts the covariate (Hernán, Brumback, & Robins, 2001). A covariate or risk factor is considered as a time-dependent confounder if it satisfies (1). Robins (1999) proposes a marginal structural model (MSM) as a method by which one can infer a causal relationship between a time-dependent treatment and outcome in the presence of a time-dependent confounder. MSM uses a two-step modeling strategy that separates confounder control from the structural model, avoiding over-adjustment of confounders (Joffe, Have, Feldman, & Kimmel, 2004). Since its development several studies have applied MSM to investigate the effect of medication use: aspirin use on cardiovascular deaths, methtrexate use on mortality in patients with rheumatoid arthritis, asthma rescue medication on peak expiratory flow rate, and heparin use on arteriovenous fistula surgery outcome in patients with ESRD (Choi, Hernan, Seeger, Robins, & Wolfe, 2002; Cook, Cole, & Hennekens, 2002; Joffe et al., 2004; Mortimer, Neugebauer, van der Laan, & Tager, 2005).

This paper is intended to provide a practical guide to researchers who wish to use MSM in a relatively quick manner. As such, the following sections on the theory of MSM and

its estimation strategy are not exhaustive. Refer to Robins, Greenland, and Hu (1999) and Robins (1999) for a thorough exposition of the theory and its mathematical justifications.

2 Theory

2.1 Counterfactuals

Ideally, individual causal effects can be estimated simply by comparing the outcome under the treatment ($Y_a = 1$) and the outcome under no treatment ($Y_a = 0$) for each person. However, observational data usually consist of only one of the two outcomes because each person receives *either* treatment or no treatment. This may be considered as a missing-data problem, in which counterfactual outcomes—the outcome that would have been observed had the individual received the treatment other than the one s/he actually received—are not observed (Mortimer et al., 2005). Marginal structural models estimate the average causal effect of a treatment on potential outcomes (or counterfactuals) by comparing the distributions of $Y_a = 1$ and $Y_a = 0$ on the aggregate (Joffe et al., 2004).

2.2 Causal Pathway

Figure 1 is a set of directed acyclic graphs (DAG) that depict the causal pathways between variables of a longitudinal study in which a time-dependent treatment is present. In a DAG, the nodes represent variables and the directed arrows represent direct causal effects (Pearl, 1995). The treatment A_t is time-dependent because its effect on the outcome varies depending on when it is administered. Figure 1(a) shows that both the measured and unmeasured covariates at baseline, L_0 and U_0 respectively, predict subsequent treatment A_1 and also independently predict the outcome, Y, confounding the treatment effect. Furthermore, the past treatment history A_0 predicts the subsequent covariate levels L_1 and U_1. Since U_t cannot be measured, it is impossible to control for the confounding caused by U_t. MSM uses weighted estimation to adjust for the confounding caused by L_t, assuming that there is no

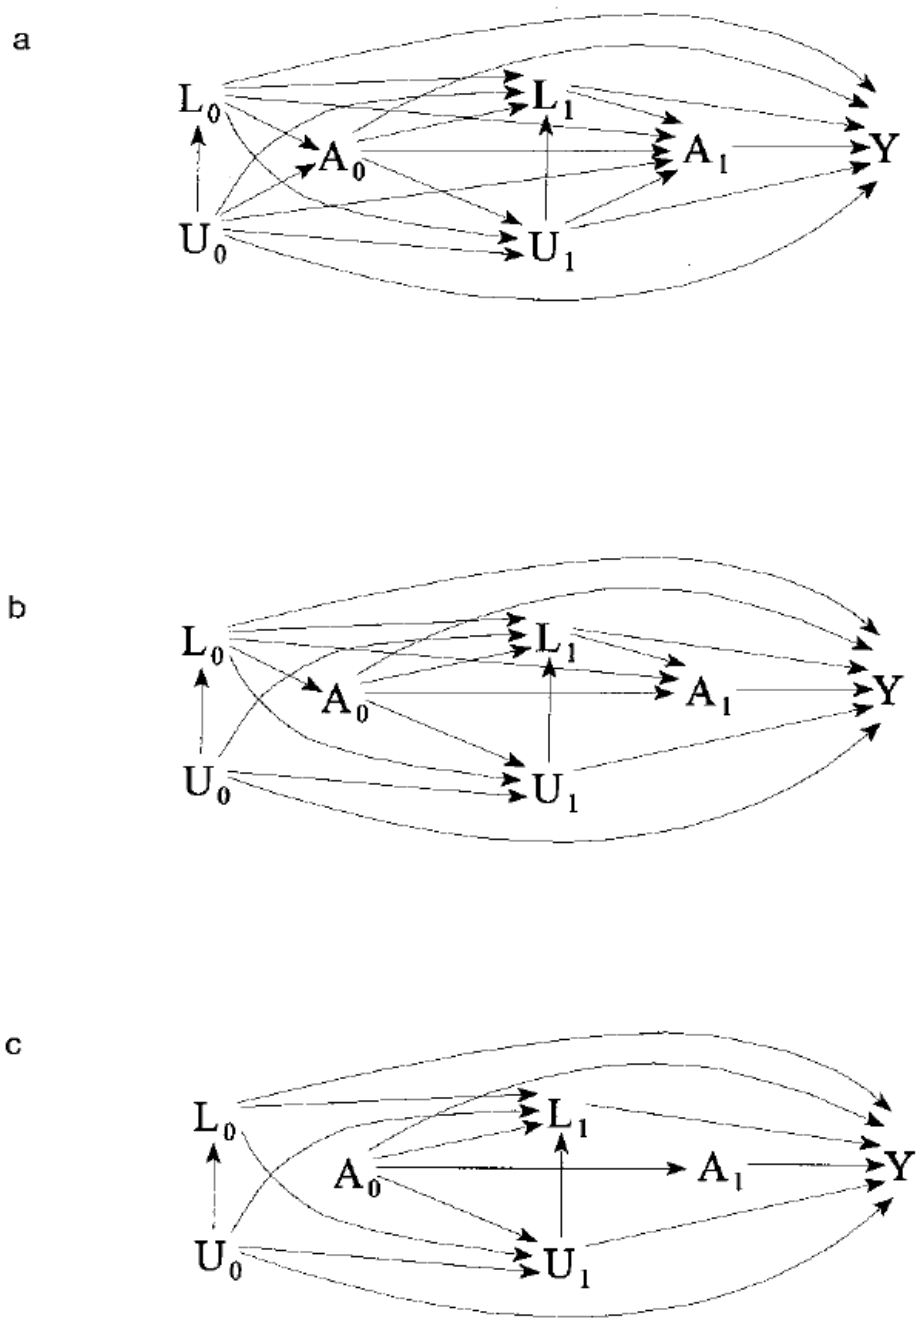

Figure 1: Directed acyclic graphs (DAG) illustrating the causal pathways between measured (L_t) and unmeasured (U_t) covariates, a time-dependent treatment (A_t), and the outcome (Y): (a) A_t is confounded by both unmeasured and measured covariates (b) A_t is confounded by only L_t (c) No confounding exists. Source: Robins et al. (2000).

unmeasured confounding as in Figure 1(b). Figure 1(c) represents the relationship in which no confounding exists.

2.3 Assumptions

MSM assumes that there exists no unmeasured confounders. In order for IPTW estimation to consistently estimate the causal effect of a time-dependent treatment, all relevant confounders should be measured (Robins, 1999). While such an assumption may seem too strong, point-treatment studies require the same assumption to make causal inferences from regression parameter estimates (Hernán, Brumback, and Robins, 2000; 2001). The assumption of no unmeasured confounders cannot be tested directly using observational data (Hernán et al., 2001). However, one may conduct a sensitivity analysis to examine the effect of the confounding caused by unmeasured confounders on the causal parameter estimates (Robins et al., 1999; Robins, 1999). Alternatively, Henneman, van der Laan, and Hubbard (2002) suggest using instrumental variables to control for unmeasured confounding.

Another critical assumption of MSM is that the probability of treatment must be nonzero. Often called the *positivity condition* or *experimental treatment assumption*, it requires that the probability of being assigned to each of the treatment options is greater than zero. In fact, w_i (defined in 2.4) is undefined when the conditional probability of treatment is 0. An example in which the positivity condition is violated is an exposure study in an occupational setting (Robins et al., 2000). If $A_t = 1$ denotes positive exposure to some industrial chemical at time t and $L_t = 1$ indicates when the worker was off duty (e.g. weekends, sick days), then $A_t = 0$ for *all* workers with $L_t = 1$. Robins et al. (2000) suggest using structural nested models for such studies. In practice, even extremely low probabilities of treatment may substantially bias the IPTW estimator (Mortimer et al., 2005). For example, subjects with certain characteristics may be practically unlikely to receive treatment. To assess the extent to which the IPTW estimates are biased due to a violation of the experimental treatment assumption, Wang, Petersen, Bangsberg, and van der Laan (2006) have developed

a diagnostic procedure that quantifies the magnitude of the bias.

2.4 Estimation

Several methods have been used to estimate the parameters in marginal structural models including inverse-probability-of-treatment weight (IPTW), double robust, and targeted maximum likelihood estimators (Odden et al., 2011).

IPTW estimator is the most commonly used estimator for MSM owing to its ease of implementation using standard statistical software packages (Mortimer et al., 2005). IPTW estimation is a two-stage process. In the first stage, weights are derived for each subject i. As its name suggests, the weights of the IPTW estimator are simply the inverse of the conditional probability of receiving treatment A given the past treatment history and covariate history:

$$w_i = \prod_{k=0}^{t} \frac{1}{P(A_{ik} = 1 \mid \bar{A}_{ik-1}, \bar{L}_{ik})} \tag{1}$$

$\bar{A}_{k-1}$ denotes treatment history through time $t-1$ and $\bar{L}_k$ denotes the covariate history through time t. $\bar{A}_{-1}$ is defined as A_0. w_i is then used to perform a weighted regression analysis in the second stage. Weighting in effect creates a pseudo-population in which no confounding exists by replicating w_i copies of each subject (Li, Evans, & Hser, 2010). Therefore, the parameter estimate of the treatment in this population can be interpreted as the true causal effect of the treatment on the outcome.

The IPTW estimator performs inefficiently if w_i has extremely large or small values. To stabilize the distribution of w_i, 1 in the numerator is replaced by the conditional probability of the treatment given the past treatment history and the baseline covariates (Robins et al., 2000):

$$sw_i = \prod_{k=0}^{t} \frac{P(A_{ik} = 1 \mid \bar{A}_{ik-1}, L_{i0})}{P(A_{ik} = 1 \mid \bar{A}_{ik-1}, \bar{L}_{ik})} \tag{2}$$

Note that if the confounding effect of $\bar{L}_{ik} = 0$ (i.e. no confounding exists), $sw_i = 1$. Both the denominator and the numerator of sw_i can be estimated using the standard statistical software packages (e.g. Stata and SAS): A logistic regression can be performed for a binary treatment variable and an OLS regression can be performed for a continuous treatment variable. Compared to sw_i, w_i tends to be more variable and skewed (Li et al., 2010). However, sw_i is still subject to skewness caused by extreme values. However, Bodnar, Davidian, Siega-Riz, and Tsiatis (2004) and Li et al. (2010) report that re-estimating the treatment effect after either removing the individuals with extreme weights or top-coding the extreme weights did not qualitatively change their findings.

The treatment model in the denominator must be correctly specified to obtain an unbiased estimate of the parameter coefficient in the subsequent regression (Mortimer et al., 2005). However, one may be tempted to include all potential confounders in the treatment model in effort to avoid a violation of the no unmeasured confounder assumption (discussed in Section 2.3). In fact, Bodnar et al. (2004), citing Robins, suggest that including more variables may be preferable to the risk of excluding relevant confounders. While it may decrease the risk of violating the no unmeasured confounders assumption, including more (and possibly irrelevant) variables increases the risk of model misspecification. Mortimer et al. (2005) propose a model building procedure that selects the best treatment model from several candidate models with different combinations of variables. In essence, the procedure chooses the treatment model that gives the best IPTW estimate of the MSM parameter determined from the validation using 10 percent of the observed data. It uses the Monte Carlo cross-validation with a modified residual sums of squares criterion to evaluate goodness of fit, which optimizes the tradeoff between bias and variance (Mortimer et al., 2005).

In a longitudinal study, some study subjects are lost to follow-up due to, among many other reasons, adverse events, personal reasons, and death. Loss to follow-up, or attrition, is a major source of selection bias because it may cause significant differences in the composition of the remaining study subjects and those who drop out. Consequently, not adjusting for loss

to follow-up may result in biased MSM parameter estimates. Censoring weights—attrition is considered as right-censoring— are typically used to account for any loss to follow-up in the observed data (Hernán et al., 2001). Conceptually, censoring is considered as another time-dependent treatment (Robins et al., 2000). Therefore, censoring weights are derived similarly to sw_i, where

$$cw_i = \prod_{k=0}^{t} \frac{P(C_{ik} = 0 \mid \bar{C}_{ik-1} = 0, A_{ik-1}, L_{i0})}{P(C_{ik} = 0 \mid \bar{C}_{ik-1} = 0, A_{ik-1}, \bar{L}_{ik})} \tag{3}$$

The only difference here is that the denominator models the probability of *not* receiving treatment (being censored) at time t instead of receiving treatment ($A_{it} = 1$) as in sw_i. The final weight is simply the product of the two weights previously derived:

$$fw_i = sw_i \cdot cw_i \tag{4}$$

fw_i is then used in a subsequent regression model to obtain the MSM parameter estimate. Fewell et al. (2004) outline how to derive fw_i and conduct a weighted regression analysis in Stata.

2.5 Limitations

MSM assumes that the treatment regime is fixed over time. For instances in which the treatment varies depending on an intermediate measure (e.g. white blood cell count), the parameter estimates of MSM may be biased (Bodnar et al., 2004). History-adjusted MSM—generalized MSM—has been proposed as an alternative approach for modeling dynamic treatment regimes (Brunelli et al., 2008; Neugebauer, Fireman, Roy, O'Connor, & Selby, 2012; Petersen, Deeks, Martin, & van der Laan, 2005).

Consistency of the IPTW estimator relies heavily on the assumption of no unmeasured confounders (Bodnar et al., 2004). Misspecification of the treatment model due to omitted confounders in deriving the IPTW can cause substantial bias in the subsequent regression

model using those weights (Mortimer et al., 2005). To address this issue, Imai and Ratkovic (2013) have proposed a new method—the covariate balancing propensity score method—that optimizes the inverse probability weights by making the treatment model robust to misspecification.

3 Applications

Cook et al. (2002) examine the effect of aspirin use on cardiovascular deaths. The data come from the Physicians' Health Study, a clinical trial that randomized about 22,000 physicians to either aspirin use every other day or placebo. In this study, non-fatal cardiovascular events such as MI, stroke, and coronary artery bypass graft are identified as a time-dependent confounder because they are predicted by prior aspirin use and predict subsequent aspirin use. They are also an independent risk factor for the outcome—cardiovascular (CV) death. Compared to the intent-to-treat analysis estimate of RR=0.99 (95% C.I.: 0.70 - 1.40), the MSM parameter estimate from a weighted pooled logistic regression suggests that aspirin use is protective against CV death (RR=0.74, 0.48 - 1.15). Although none of the estimates are statistically significant, the lower RR given by MSM indicates that the IPTW weight appears to have removed the confounding by indication—those who experience a non-fatal CV event are more likely to initiate aspirin use compared to those who do not.

While most applications of MSM are found in epidemiology and medicine, some studies have begun to apply MSM in other contexts. Do, Wang, and Elliott (2012), for example, apply MSM to investigate the effect of neighborhood poverty on mortality risk. Because neighborhood exposure is dynamic—factors that affect neighborhood such as demographics, socioeconomics, and the environment change over time—they suggest that single-point estimates are likely to be underestimates. Moreover, they point out that time-varying covariates such as income, marital status, educational attainment, and employment status predict future likelihood of living in a certain neighborhood and also independently predict health

outcomes. Therefore, adjusting for the time-varying covariates in a single regression model to estimate the neighborhood effect on mortality risk may result in biased estimates. Calling such a model the naïve model, they compare the estimates obtained from the naïve models to that of the MSM models. Whereas the naïve models do not show any significant neighborhood effect, the MSM model for the neighborhood poverty from 20% to 100% indicates a 63% increase in the odds of mortality for every 10% increase in neighborhood poverty.

4 Example

Cook et al. (2002) provide a hypothetical example in which 200,000 subjects are randomized to either aspirin or no aspirin use and followed for 2 years (Figure 2). The outcome variable is death cause by myocardial infarction (MI). The probabilities of occurrence for each state are listed in Cook et al. (2002).

Analysis	Rate ratio
Intent-to-treat	0.52
As-treated (AT)	0.66
Counterfactual AT	0.43
MSM	0.44

Table 1: Rate ratios of the rate of deaths among aspirin users to that of non-users based on the frequencies in Figure 2. The rate ratio for MSM was taken from Cook et al. (2002).

The estimated rate ratio (RR) in the **intent-to-treat** (ITT) analysis is 0.52 while the RR in the **as-treated** (AT) analysis is 0.63 (Table 1). The AT RR is greater than the ITT RR because those who experience nonfatal MI in year 1 have increased risk of MI death in year 2 compared to those who do not experience nonfatal MI (Cook et al., 2002). Both estimates are still biased because the effect of aspirin use on MI death is confounded by the nonfatal MI. The true causal RR can be estimated by assuming full compliance to the

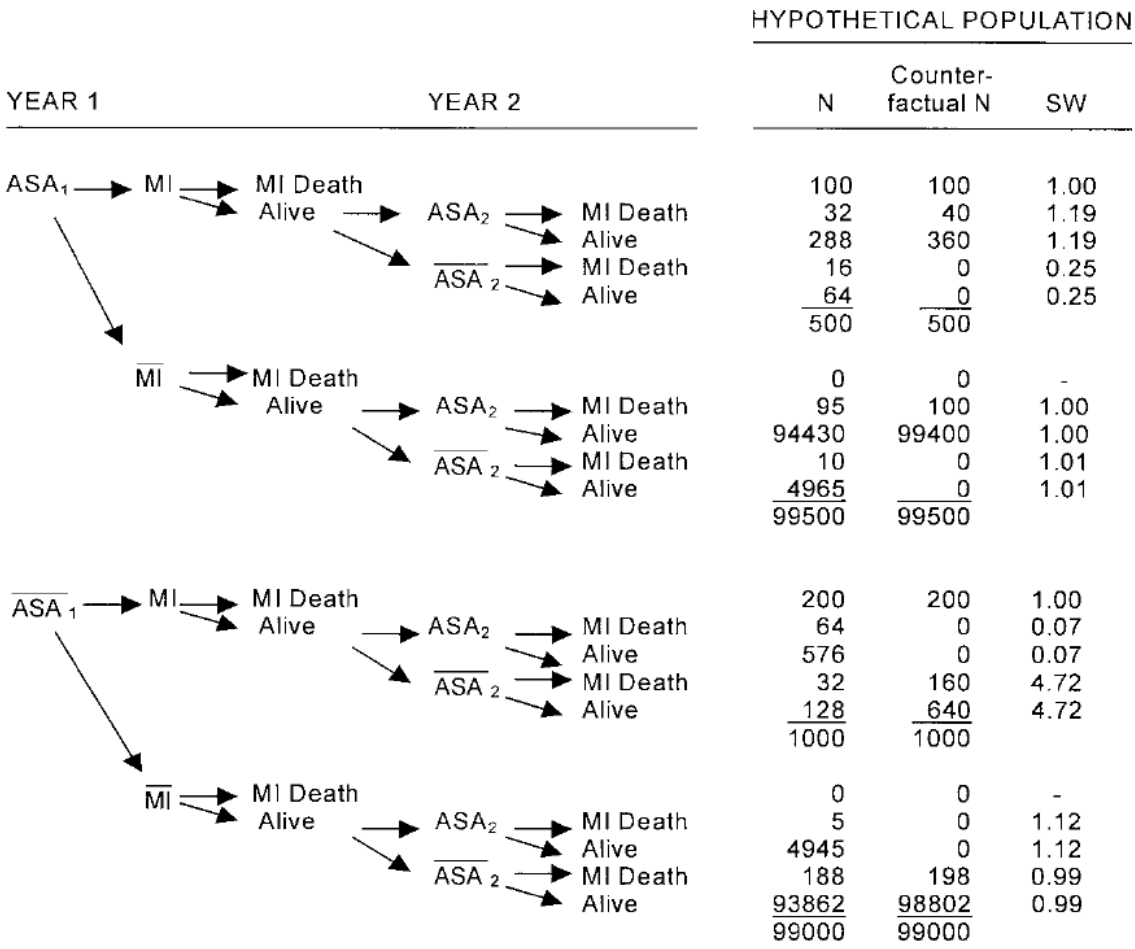

Figure 2: A hypothetical population of 200,000 randomized to aspirin (ASA_t) or no aspirin use ($\overline{ASA}_t$), where t indicates the follow-up year. MI denotes myocardial infarction, N, the frequency of subjects who followed each path, and SW, the stabilized weights for subjects in each subgroup. Source: Cook et al. (2002).

treatment assignment. The counterfactual N reflects the frequencies of each state in year 2 under full compliance (Figure 2). The counterfactual RR of 0.43—lower than the arbitrary RR of 0.5—suggests that aspirin use reduces not only the risk of MI death, but also the risk of nonfatal MI. The MSM estimate using the stabilized weights is 0.44—slightly higher but the closest to the true causal RR compared to the ITT and AT estimates (Table 1).

A reproduced version of the example, *not* written by the original authors, is available as an Excel spreadsheet (https://www.dropbox.com/s/n84tpnpax7ru4a8/COOK%20APPENDIX_mod.xlsx).

5 Suggested Reading

Cook et al. (2002) presents a relatively straightforward application of MSM in examining the effect of aspirin use on cardiovascular deaths. The hypothetical example in its appendix is described in Section 4. Most of the applications of MSM presented in this paper involve a medication treatment effect; as summarized in Section 3, Do et al. (2012) applies MSM to investigate the effect of a non-medication treatment—neighborhood poverty.

Fewell et al. (2004) provides useful information for implementing MSM using Stata. Complete with step-by-step Stata codes and abbreviated outputs in the body of the text (*not* in appendix), it guides the reader in programming Stata to run a MSM.

Robins et al. (1999) and Robins (1999) develop the theory on marginal structural modeling. It provides the intuition of the method as well as its mathematical justifications. While difficult to read due to their theoretical nature, they do represent the early works that formally developed MSM.

MSM is only one class of models that estimate causal parameters. In their book, *Causal Inference* (work in progress), Miguel Hernán and Jamie Robins synthesize the theories and methods used in modeling causality in a cohesive manner. A detailed description of the book as well as the drafts and other supplemental materials including SAS, Stata, and R programs can be found on Hernán's webpage: http://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/.

In addition to the materials from *Causal Inference*, more programs are available on the webpage from the Program on Causal Inference at Harvard School of Public Health (http://www.hsph.harvard.edu/causal/software/). The programs pertaining to this report include "Marginal structural Cox model in SAS", "Structural models for survival analysis in SAS: the MSM macro", and "Marginal structural models in Stata". Note that "Marginal structural models in Stata" does not provide the actual program file as in the first two, but merely refers to Fewell et al. (2004).

References

Bodnar, L. M., Davidian, M., Siega-Riz, A. M., & Tsiatis, A. A. (2004). Marginal structural models for analyzing causal effects of time-dependent treatments: an application in perinatal epidemiology. *Am J Epidemiol*, *159*(10), 926–34.

Brunelli, S. M., Joffe, M. M., Israni, R. K., Yang, W., Fishbane, S., Berns, J. S., & Feldman, H. I. (2008, May). History-adjusted marginal structural analysis of the association between hemoglobin variability and mortality among chronic hemodialysis patients [Article; Proceedings Paper]. *Clinical Journal of the American Society of Nephrology*, *3*(3), 777–782. (39th Annual Meeting of the American-Society-of-Nephrology, San Diego, CA, NOV 14-19, 2006)

Choi, H., Hernan, M., Seeger, J., Robins, J., & Wolfe, F. (2002, APR 6). Methotrexate and mortality in patients with rheumatoid arthritis: a prospective study. *Lancet*, *359*(9313), 1173-1177.

Cook, N., Cole, S., & Hennekens, C. (2002, JUN 1). Use of a marginal structural model to determine the effect of aspirin on cardiovascular mortality in the physicians' health study. *Am J Epidemiol*, *155*(11), 1045-1053.

Do, D. P., Wang, L., & Elliott, M. (2012). *Investigating the relationship between neighborhood poverty and mortality risk: A marginal structural modeling approach* (Tech. Rep. No. 12-763). University of Michigan Institute for Social Research.

Fewell, Z., Hernán, M. Á., Wolfe, F., Tilling, K., Choi, H., & Sterne, J. (2004). Controlling for time-dependent confounding using marginal structural models. *Stata Journal*, *4*(4), 402–420(19).

Henneman, T. A., van der Laan, M., & Hubbard, A. E. (2002). *Estimating causal parameters in marginal structural models with unmeasured confounders using instrumental variables* (No. 104). University of California, Berkeley.

Hernán, M. Á., Brumback, B., & Robins, J. M. (2000, SEP). Marginal structural models to estimate the causal effect of zidovudine on the survival of hiv-positive men. *Epidemi-*

ology, *11*(5), 561-570.

Hernán, M. Á., Brumback, B., & Robins, J. M. (2001). Marginal structural models to estimate the joint causal effect of nonrandomized treatments. *Journal of the American Statistical Association*, *96*(454), 440–448.

Imai, K., & Ratkovic, M. (2013). *Robust estimation of inverse probability weights for marginal structural models.* Princeton University.

Joffe, M. M., Have, T. R. T., Feldman, H. I., & Kimmel, S. E. (2004). Model selection, confounder control, and marginal structural models: Review and new applications. *The American Statistician*, *58*(4), 272–279.

Li, L., Evans, E., & Hser, Y. I. (2010). A marginal structural modeling approach to assess the cumulative effect of drug treatment on the later drug use abstinence. *J Drug Issues*, *40*(1), 221–240.

Mortimer, K., Neugebauer, R., van der Laan, M., & Tager, I. (2005, AUG 15). An application of model-fitting procedures for marginal structural models. *Am J Epidemiol*, *162*(4), 382-388.

Neugebauer, R., Fireman, B., Roy, J. A., O'Connor, P. J., & Selby, J. V. (2012). Dynamic marginal structural modeling to evaluate the comparative effectiveness of more or less aggressive treatment intensification strategies in adults with type 2 diabetes. *Pharmacoepidemiol Drug Saf*, *21 Suppl 2*, 99–113.

Odden, M. C., Tager, I. B., van der Laan, M. J., Delaney, J. A. C., Peralta, C. A., Katz, R., ... Shlipak, M. G. (2011). Antihypertensive medication use and change in kidney function in elderly adults: A marginal structural model analysis. *International Journal of Biostatistics*, *7*(1).

Pearl, J. (1995, DEC). Causal diagrams for empirical research. *Biometrika*, *82*(4), 669-688.

Petersen, M. L., Deeks, S. G., Martin, J. N., & van der Laan, M. (2005). *History-adjusted marginal structural models to estimate time-varying effect modification* (No. 199).

Robins, J., Greenland, S., & Hu, F. (1999, SEP). Estimation of the causal effect of a time-

varying exposure on the marginal mean of a repeated binary outcome. *Journal of the American Statistical Association*, *94*(447), 687-700.

Robins, J. M. (1999). Association, causation, and marginal structural models. *Synthese*, *121*(1-2), 151–179.

Robins, J. M., Hernán, M. Á., & Brumback, B. (2000, September). Marginal structural models and causal inference in epidemiology. *Epidemiology*, *11*(5), 550–560.

Wang, Y., Petersen, M. L., Bangsberg, D. R., & van der Laan, M. J. (2006). *Diagnosing bias in the inverse probability of treatment weighted estimator resulting from violation of experimental treatment assignment.* U.C. Berkeley Division of Biostatistics Working Paper Series.

Glossary

as-treated analysis based on subjects' observed treatment and outcome in a randomized clinical trial. While it accommodates loss to follow-up, randomization is lost because the remaining subjects in each treatment arm may no longer be random.

intent-to-treat analysis based on subjects' initial treatment assignment in a randomized clinical trial. Consequently, it does not take into account attrition (or loss to follow-up).