Practice of Epidemiology

Time-modified Confounding

Robert W. Platt, Enrique F. Schisterman, and Stephen R. Cole

According to the authors, time-modified confounding occurs when the causal relation between a time-fixed or time-varying confounder and the treatment or outcome changes over time. A key difference between previously described time-varying confounding and the proposed time-modified confounding is that, in the former, the *values* of the confounding variable change over time while, in the latter, the *effects* of the confounder change over time. Using marginal structural models, the authors propose an approach to account for time-modified confounding when the relation between the confounder and treatment is modified over time. An illustrative example and simulation show that, when time-modified confounding is present, a marginal structural model with inverse probability-of-treatment weights specified to account for time-modified confounding remains approximately unbiased with appropriate confidence limit coverage, while models that do not account for time-modified confounding are biased. Correct specification of the treatment model, including accounting for potential variation over time in confounding, is an important assumption of marginal structural models. When the effect of confounders on either the treatment or outcome changes over time, time-modified confounding should be considered.

bias (epidemiology); confounding factors (epidemiology); structural model

Confounding occurs when treatment and disease share a common cause (1). *Time-varying* confounding occurs when there is a *time-varying* cause of disease that brings about changes in a time-varying treatment (2, 3). *Time-varying confounding affected by prior treatment* occurs when subsequent values of the time-varying confounder are caused by prior treatment (4). Of course, the measured (time-fixed or -varying) confounder may be a proxy for the underlying causal confounder. Here, we consider "time-modified confounding," which occurs when there is a *time-fixed* or *time-varying* cause of disease that also affects subsequent treatment, but where the *effect* of this confounder on either the treatment or outcome changes over time.

A key difference between time-varying confounding and time-modified confounding is that, in the former, the *values* of the confounding variable change over time while, in the latter, the *effects* of the confounder change over time. Further, time-modified confounding may occur with a time-fixed or time-varying covariate, while time-varying confounding occurs only with time-varying covariates. Time-varying confounding and time-modified confounding may therefore occur simultaneously in the case of a time-varying covariate. Time-modified confounding is not to be confused with effect measure modification, where the effect of treatment on outcome differs by levels of an effect modifier (even if the effect modifier is also a time-varying confounder (5)). The term "modified" was chosen because the effect of the *confounder* on either the treatment or outcome is modified over time.

Time-varying effects of treatment have been considered in the epidemiologic and statistical literature (6–8); indeed, an assessment (9) of the proportional hazards assumption (10) is an assessment of a departure from a constant effect of exposure on outcome. Our purpose here is to define and illustrate time-modified confounding. Studying the effect of breastfeeding on infants' weight gain, we provide an example and evaluate the impact of time-modified confounding on the bias and variability of estimates of causal effect, using simulation.

CAUSAL EFFECTS

We begin by defining causal effects using potential outcomes (11, 12). For an outcome *Y*, we denote the potential

Correspondence to Dr. Robert W. Platt, Montreal Children's Hospital Research Institute, 4060 Ste. Catherine Street West, No. 205, Westmount, QC, Canada H3Z 2Z3 (e-mail: robert.platt@mcgill.ca).

outcome Y_x as the outcome for a given treatment level x; there are as many potential outcomes as there are levels of treatment (note that we use "treatment" throughout to refer to an exposure that may or may not be an assigned treatment). Specifically, the average causal effect for treatment level x, compared with level $x' \neq x$, is defined as $E(Y_x) - E(Y_{x'})$, where the expectations are taken over the same subjects under different levels of treatment. Instead of a difference, one could make the contrast a ratio or an odds ratio; we do the latter in the example below. We could also define an average covariate-conditional causal effect as $E(Y_x | Z = z) - E(Y_{x'} | Z = z)$, the causal effect at a specific level z of covariate Z. The average covariate-conditional causal effect does not necessarily equal the average causal effect for certain contrasts (in particular, the odds ratio), even in the absence of effect modification by Z (13, 14).

CAUSAL DIAGRAMS

Pearl (15) formalized causal diagrams as directed acyclic graphs, giving investigators powerful tools for bias assessment, providing the rules of causal diagrams are followed. A set of these rules is succinctly given in the Appendix of the work by Hernán et al. (16). Causal diagrams link variables by single-headed (i.e., directed) arrows that represent direct causal effects. The absence of an arrow between 2 variables, on the other hand, is a strong claim of no direct causal effect of the former variable on the latter. To represent chains of causation in time, Pearl's original formalization of causal diagrams does not allow a directed path (i.e., a trail of arrows) to point back to a prior variable (i.e., the diagrams are acyclic). For a diagram to represent causal effects as defined in the prior section, all common causes of any pair of variables included on the graph must also be included on the graph.

CONFOUNDING

Confounding refers to settings in which treatment and outcome share a common cause, which may be represented by a single variable or a combination of variables. Figure 1A describes a simple case of confounding. Let $X(0)$ represent treatment at time 0 (using values in parentheses to represent time ordering), $Y(2)$ outcome at time 2, and $Z(0)$ a confounding variable occurring temporally prior to $X(0)$ that has a direct causal effect on both $X(0)$ and $Y(2)$; note that, in this figure, $X(0)$ has no causal effect on $Y(2)$. Robins (17) and Hernán et al. (16) demonstrated the need to consider substantive knowledge in decisions on adjustment for confounders. In Figure 1A, control for $Z(0)$ through regression, stratification, or restriction provides a consistent estimate of the $Z(0)$-conditional causal effect of treatment $X(0)$ on outcome $Y(2)$ (an estimator is consistent if it converges in probability to the true value as the sample size tends to infinity). In the following section, we consider a series of extensions of the confounding definition to settings involving time-varying variables and effects.
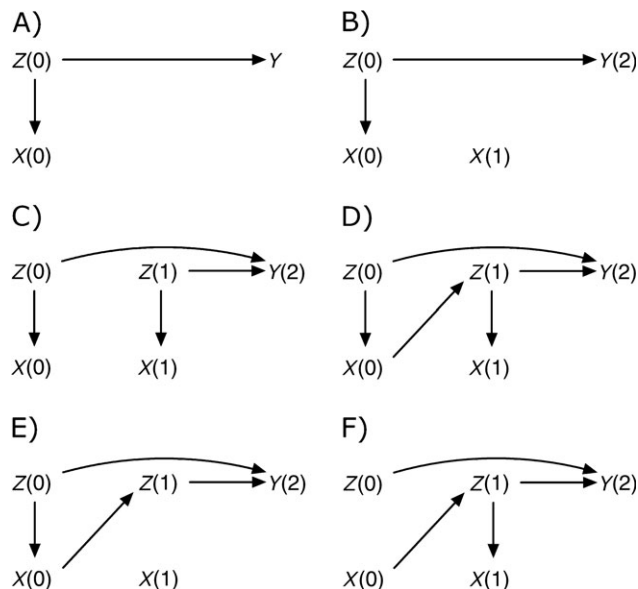

Figure 1. Causal diagrams representing confounding (A), time-modified confounding by a time-fixed covariate (B), time-varying confounding (C), time-varying confounding affected by prior treatment (D), and time-modified confounding (E and F).

Time-modified confounding

Figure 1B extends the setting in Figure 1A to the case where treatment X now varies over time, but $Z(0)$ confounds only the $X(0)$–$Y(2)$ relation. This is a simple case of time-modified confounding. Control for $Z(0)$ will give a consistent estimate of the $Z(0)$-conditional causal effect of $X(0)$ on $Y(2)$, while the effect of $X(1)$ on $Y(2)$ can be estimated consistently from the crude (unadjusted) model.

Time-varying confounding

Figure 1C describes a setting with time-varying (or time-dependent) confounding. $X(0)$ and $X(1)$ represent treatment at times 0 and 1, and $Z(0)$ and $Z(1)$ represent (possibly a set of) time-varying confounding variables measured temporally prior to times $t = 0$ and $t = 1$, respectively. In the following, we consider only 2 time points; the argument generalizes to more time points. At each time point, the confounders have a direct causal effect on treatment. As in Figure 1, A and B, Figure 1C is constructed such that there is no direct causal effect of $X(0)$ or $X(1)$ on $Y(2)$, the outcome measured at time 2. To simplify exposition, we measure outcome only at time 2; however, our claims apply equally if the outcome were measured at each time point. By using Figure 1C to estimate the total (i.e., direct and indirect) causal effect of $X(0)$ on $Y(2)$, adjustment for $Z(0)$ is necessary. To estimate the total causal effect of $X(1)$ on $Y(2)$, we must adjust for $Z(1)$. In the setting described by Figure 1C, standard statistical methods (e.g., Cox regression with time-varying treatment and covariates) can consistently estimate, as previously defined, the $Z(t)$-conditional causal effect of treatment $X(t)$, $t = 0$ or 1, on outcome $Y(2)$ (4).

Figure 1D describes an extension of Figure 1C, where now the time-varying confounder is *affected* by prior treatment. In this case, adjustment for $Z(1)$ is necessary to estimate the effect of $X(1)$ but blocks a causal effect of $X(0)$. In such settings, Robins et al. (18) showed that estimation of the total effect of $X(t)$ was biased when standard statistical methods were used. Robins et al. (4) have proposed a series of methods to address this problem, including marginal structural models, which will be described below after we further define time-modified confounding.

Time-modified confounding by time-varying factors

Figure 1E describes a further extension of Figure 1C, which illustrates time-modified confounding. In this case, the effect of the confounding variables $Z(t)$ on treatment $X(t)$ differs over time t; in the case of Figure 1E, there is a direct causal effect of $Z(0)$ on $X(0)$, but there is no direct causal effect of $Z(1)$ on $X(1)$. Figure 1F describes a companion scenario of time-modified confounding, where $Z(0)$ has no direct causal effect on $X(0)$, although $Z(1)$ does have a direct causal effect on $X(1)$. Two other companion scenarios of time-modified confounding could be drawn, alternating absence and presence of the effects of $Z(0)$ and $Z(1)$ on $Y(2)$, respectively.

Figure 1E could represent an observational study of a time-varying treatment $X(t)$, such as use of a pharmacologic therapy, where $Z(0)$ and $Z(1)$ represent an indication for early treatment, such as access to care. In this setting, having the indication may be positively associated with use of a novel pharmacologic agent (at time 0), but as inequities in use become apparent and corrected, access to care may no longer be associated with use of the agent at time 1.

Figure 1F could represent a randomized trial of a treatment taken over time, in which $X(t)$ represents treatment actually received at time t (i.e., irrespective of randomization). In this setting, we assume that there is no effect of $Z(0)$ on $X(0)$ by design, because of randomization and full initial compliance. However, it is plausible that compliance behavior changes over time, and that the dependence of compliance on covariates changes over time as well, so that $Z(1)$ has a direct causal effect on $X(1)$, or that $Z(1)$ and $X(1)$ share a common cause.

As an aside, in both of these simple cases (i.e., Figure 1, E and F), the causal effect of $X(t)$ on $Y(2)$ could be consistently estimated by standard methods (such as linear or logistic regression). For example, in Figure 1E, adjusting for $Z(0)$ is sufficient to control confounding and to provide an unbiased estimate of the causal effect. In Figure 1F, the simple cross-tabulation of $X(0)$ and $Y(2)$ would provide an unbiased estimate of the causal effect of $X(t)$ on $Y(2)$. However, these approaches are based on the knowledge that the hypothesized diagram is correct. In practice, one may be unlikely to estimate the effect of $X(t)$ without using all measured exposure and confounding information. Such adjustment may introduce bias. Moreover, as the dimension of the problem grows with additional time points, summaries of exposures and confounders (such as cumulative averages (19)) are needed and often preclude such simple solutions. In practice, these simple solutions would fail if there were a causal

effect of $Z(1)$ on $X(1)$ (4). Marginal structural models can provide consistent estimates in either case.

Time-modified confounding is not restricted to settings where an effect is present at one time and absent at another time. Indeed, such examples will be rare relative to examples where the *size* or the *direction* of the effect changes over time. Causal diagrams are better suited to illustrate all-or-none effects, because arrows encode the presence or absence of a direct causal effect rather than the size or direction of the effect. One could use a signed causal diagram (20, 21), under certain assumptions, to describe settings where the effect direction changes over time.

In Figure 1, D–F, $Z(1)$ is a confounding variable affected by previous treatment $X(0)$. As previously noted, in such settings marginal structural models can be used to consistently estimate the total causal effect of $X(t)$ on $Y(2)$. In the next section, we describe how to use marginal structural models to account for time-modified confounding in the presence of time-varying confounding.

MARGINAL STRUCTURAL MODELS

Marginal structural models (4, 22) are models for the marginal expectation of a potential outcome as a function of a specified treatment regimen. For example, if Y is an outcome and $X(t)$ is a time-varying treatment, then a marginal structural model is specified as

$$E[Y_{\bar{x}}] = f(\bar{x}),$$

where $\bar{x}$ refers to the history of treatment $X(t)$, and $f(.)$ is a defined function, typically a (perhaps transformed) linear combination of components of x. To estimate the parameters of a marginal structural model, we compute stabilized weights (22). We first fit a model for the probability of receiving treatment x and then weight individuals by the inverse probability of receiving the observed treatment given the measured treatment and confounder histories. These weights are stabilized to improve efficiency by a function *not* including the variables for which one wishes the weights to remove confounding. Specifically, for a categorical exposure, the weights are defined as follows:

$$W_i(t) = \prod_{s=0}^{t} P[X_i(s) = x_i(s)] / P[X_i(s) = x_i(s) \mid \bar{X}_i(s-1), \bar{Z}_i(s)],$$

where $\bar{X}_i(s)$ represents treatment history from baseline to time s. These weights can be extended to continuous treatments by replacing the probability mass functions with the corresponding densities (4). These inverse probability-of-treatment weights can be multiplied by inverse probability-of-censoring weights when censoring is informative by measured variables (22); we do not consider censoring here. An unadjusted model for the outcome as a function of treatment is then fit to the *weighted* sample. If the functional form of the treatment model is correctly specified and the other assumptions of the marginal structural model (i.e., consistency (23, 24), positivity (25), exchangeability (26))

Table 1.   Odds Ratios for Continued Breastfeeding in the First Year of Life[a]

Variable	Odds Ratio for Continued Breastfeeding						
	1 Month	2 Months	3 Months	6 Months	9 Months	12 Months	Pooled
Intervention	2.09	1.49	1.47	1.42	1.34	1.25	1.47
Past smoking	0.51	0.58	0.84	0.72	0.91	1.03	0.65
Maternal age, years	1.12	1.07	1.08	1.04	1.09	1.06	1.07
Maternal age^2, years^2	1.00	1.00	1.00	1.00	1.00	1.00	1.00
Child's weight, kg	3.61	2.04	4.04	3.01	1.06	1.23	0.52
Child's weight^2, kg^2	0.84	0.97	0.91	0.93	1.00	0.99	1.03

[a] Models are adjusted for hospital region, maternal education, current drinking and smoking, history of atopy and breastfeeding, child's sex, and time-varying health status.

hold, then the estimate of the marginal effect of treatment has a causal interpretation.

A critical step in fitting a marginal structural model is the development of a model for the probability of treatment. When treatment is binary and time varying, logistic regression models for each time point are typically pooled together in a single model fit (27) for the probability of treatment given the history of the treatment and confounders. The coefficients in this pooled logistic model represent the effect of confounders on treatment conditional on past treatment and averaged over time points. In Figure 1, C and D, for example, assuming that the association between confounder $Z(t)$ and treatment $X(t)$ is constant over time, then the single coefficient from the pooled logistic model will accurately portray this time-stable association. However, in Figure 1, E and F, where this association varies over time, the single coefficient for $Z(t)$ in the pooled logistic model will represent a weighted average of a null association (of $Z(0)$ on $X(0)$) and a nonnull association (of $Z(1)$ on $X(1)$) and, hence, fail to appropriately account for time-varying confounding.

A simple method for correcting this potential residual confounding bias is to fit the logistic model for treatment separately at each time point, rather than pooling across time. A cost of this correction is a loss in efficiency that will appear in the form of increased variability in the estimated weights. In real data settings (i.e., finite samples), a compromise must be sought between control of confounding bias and loss of precision along the lines previously discussed (25, 28). We note that marginal structural models (29) require that the model for the weights be correctly specified; however, to our knowledge, little attention has been paid to the issue of time-modified confounding in specifying models for the weights.

EXAMPLE

Moodie et al. (30) considered the causal effect of breastfeeding on infant weight gain in a cluster-randomized trial of a breastfeeding promotion intervention (31). Maternity hospitals and their affiliated polyclinics were randomized to a breastfeeding promotion intervention or to standard care. The data include 17,046 mother-infant pairs from 31 sites, all of whom started breastfeeding. Breastfeeding status was

recorded at 1, 2, 3, 6, 9, and 12 months of life, with weight at 12 months as the outcome.

Here, we consider a marginal structural linear model for infant weight at 12 months, as a function of breastfeeding regimen, with regimens of the form "breastfeed until month j and then stop," with $j = 1, 2, 3, 6, 9, 12$. To fit the marginal structural model, we computed inverse probability-of-continued-breastfeeding weights at each time point. We generated stabilized inverse probability weights using 1) pooled weighting, such that all time points were treated in the same model, and 2) separate models for the inverse probability weights at each time point. Table 1 presents selected odds ratios from the probability-of-continued-breastfeeding models under the 2 specifications. In the pooled model, the odds ratio for past maternal smoking is 0.65, while in the separate models the odds ratio is 0.51 at 1 month and 1.03 at 12 months, indicating substantial time-modified confounding. The corresponding estimated causal effects of breastfeeding to 12 months (relative to early weaning) are 0.10 kg (95% confidence limits: −0.01, 0.22) in the standard marginal structural model and −0.09 kg (95% confidence limits: −0.18, 0.01) in the marginal structural model allowing for time-modified confounding. The reversal of effect is somewhat surprising; however, the large differences in the probability-of-treatment models for the 2 marginal structural models (and the probable misspecification of the pooled model) likely account for this difference.

SIMULATION STUDY

We now consider a simulation study with time-modified confounding. For illustrative purposes, we consider simple settings, reflecting some diagrams in Figure 1. Consider a study of the proportion of patients with hypercholesterolemia whose cholesterol levels have not changed after 8 months of treatment with a novel pharmacologic agent versus a standard agent. We assume that the study is conducted within levels of all major time-fixed predictors of treatment failure, so that there is no unmeasured confounding. We considered 3 scenarios, corresponding to Figure 1, C–E. In all 3 cases, at study entry ($t = 0$), patients with an indicator of poor liver function (e.g., alanine transaminase, >32 IU/L for women or >50 IU/L for men) were placed on the novel therapy at 4 times the odds of being placed on

Table 2. Odds Ratios, 95% Confidence Limit Coverage, and Simulation Standard Errors for Cumulative Average Treatment With a Novel (Versus Standard) Cholesterol Therapy and 8-Month Risk of Treatment Failure for Unadjusted, Adjusted, Time-pooled-weighted, and Time-stratified-weighted Marginal Structural Logistic Regression Models

Scenario and Method	Odds Ratio[a]	95% Confidence Limit Coverage, %	Standard Error[b]
Scenario 1 (Figure 1C), causal odds ratio = 1.0[c]			
Crude	2.07	N/A[d]	0.381
Adjusted	1.00	95.6	0.416
Standard marginal structural model	1.01	94.7	0.427
Proposed marginal structural model	1.01	94.7	0.427
Scenario 2 (Figure 1D), causal odds ratio = 1.35			
Crude	1.95	N/A	0.365
Adjusted	1.01	N/A	0.453
Standard marginal structural model	1.39	94.7	0.400
Proposed marginal structural model	1.39	94.7	0.399
Scenario 3 (Figure 1E), causal odds ratio = 1.35			
Crude	2.23	N/A	0.384
Adjusted	1.00	N/A	0.437
Standard marginal structural model	1.61	N/A	0.392
Proposed marginal structural model	1.37	94.8	0.407

Abbreviation: N/A, not applicable.

[a] Geometric mean of 5,000 estimated odds ratios.

[b] Simulation standard error for the log odds ratio (i.e., standard deviation of 5,000 estimated log odds ratios).

[c] In scenario 1, where the true effect is null, the statistical power column gives the type 1 error.

[d] Confidence limit coverage and statistical power were provided for only approximately unbiased estimators.

a standard therapy. Having poor liver function was a strong predictor of treatment failure (i.e., the odds ratios for $Z(0) = 6$ and for $Z(1)$ given $Z(0) = 2$), irrespective of which treatment a patient received. In scenario 1, treatment has no effect on either liver function at 4 months or outcome. Scenario 2 modifies the first scenario by the addition of a causal effect of therapy on liver function at 4 months. Finally, scenario 3 is a modification of scenario 2, in which we remove the direct causal effect between liver function at 4 months and treatment just after 4 months. For each scenario, 5,000 simulated data sets of 600 subjects were generated. Details of the data structure are given in the Appendix.

Table 2 summarizes results from a series of analyses for each scenario. We present the geometric mean odds ratio, the estimated 95% confidence limit coverage probability, and the standard error of the log odds ratio for each analysis for 4 models. First, we present results from an unadjusted logistic regression model where treatment failure is the outcome and the sole regressor is the cumulative average treatment, defined simply as $[X(0) + X(1)]/2$. Second, we present results from a $Z(t)$-adjusted logistic regression model, where treatment failure is again the outcome and the regressors are the cumulative average treatment, along with $Z(0)$ and $Z(1)$. Third, we present results from a *standard* marginal structural logistic regression model, where treatment failure is

again the outcome and the sole regressor is the cumulative average treatment, but the patient contributions are weighted by the inverse probability of treatment, *pooled* over times 0 and 1. Fourth and last, we present results from a marginal structural logistic regression model, where treatment failure is again the outcome and the sole regressor is again the cumulative average treatment; however, now the patients' contributions are weighted by the inverse probability of treatment, which is *stratified* by time.

In scenario 1 (Figure 1C), there is no direct or indirect causal effect of treatment $X(t)$ on outcome $Y(2)$. On the basis of the theory of causal diagrams (15), we expect the crude analyses to be biased and the adjusted and both marginal structural model analyses to be unbiased, with the adjusted analysis being most efficient. As seen in Table 2, the crude analysis is biased away from the null with a geometric mean odds ratio of 2.07. The adjusted, standard, and proposed marginal structural models were each approximately unbiased, with geometric mean odds ratios of 1.00, 1.01, and 1.01 and 95% confidence limit coverages of 96%, 95%, and 95%, respectively. Relative to the adjusted analysis, the efficiencies of the standard and proposed marginal structural models were 1.05 and 1.06, respectively.

In scenario 2 (Figure 1D), there is an indirect causal effect of treatment $X(0)$ on outcome $Y(2)$ mediated by the time-varying confounder $Z(1)$, which is expected to yield an odds

ratio of 1.35. Again, on the basis of causal diagram theory, we expect both the crude and adjusted analyses to be biased and both marginal structural model analyses to be unbiased, with the standard marginal structural model being the more efficient of the 2. As seen in Table 2, the crude and adjusted analyses are biased away and toward the null, with geometric mean odds ratios of 1.95 and 1.01, respectively. The standard and proposed marginal structural models were each biased slightly away from the null, with geometric mean odds ratios of 1.39 and 1.39 and 95% confidence limit coverages of 95% and 95%, respectively. In both cases, the lower bound of a 95% simulation interval was 1.37. The efficiency of the proposed marginal structural models relative to the standard marginal structural model was 1.00.

In scenario 3 (Figure 1E), there is again an indirect causal effect of $X(0)$ on $Y(2)$ mediated by $Z(1)$, which is expected to yield an odds ratio of 1.35. On the basis of theory, we expect the crude, adjusted, and standard marginal structural model analyses to be biased and the proposed marginal structural model analysis to be unbiased. As seen in Table 2, the crude, adjusted, and standard marginal structural model analyses are biased with geometric mean odds ratios of 2.23, 1.00, and 1.61, respectively. The proposed marginal structural model was again biased slightly away from the null, with a geometric mean odds ratio of 1.37 and 95% confidence limit coverage of 95%. The lower bound of a 95% simulation interval was 1.36.

We expected to observe a loss of efficiency when estimating the proposed, as compared with the standard, marginal structural model in scenario 2, when the standard marginal structural model was correctly specified. In Figure 2, we plot the inverse probability weights from the standard marginal structural model by the corresponding weights from the proposed marginal structural model for 2,000 samples of size 250 from scenario 2. We chose a sample size of 250 to exaggerate the efficiency loss in a small-sample setting. Three groupings of points are apparent in the plot. Individuals whose treatment was always as predicted by the treatment model fall in the lower-left quadrant along the 45-degree reference line. Individuals whose treatment was always the *opposite* of that predicted by the treatment model fall into the upper-right quadrant along the 45 degree line. Individuals whose treatment was sometimes predicted correctly fall into a cluster in the center of the plot; the higher variability of the proposed weights (the wider spread along the X-axis than the Y-axis) is evidence of the loss of efficiency that occurs when the stratified weights are used when not needed.

DISCUSSION

Time-modified confounding occurs when there is a time-fixed or time-varying cause of disease that also influences subsequent treatment, and when the *effect* of this confounder on either the treatment or outcome changes over time. Time-modified confounding differs from time-varying confounding because, in the former, the magnitude or direction of the effect, rather than the value of the variables, changes over time. Time-modified confounding differs from
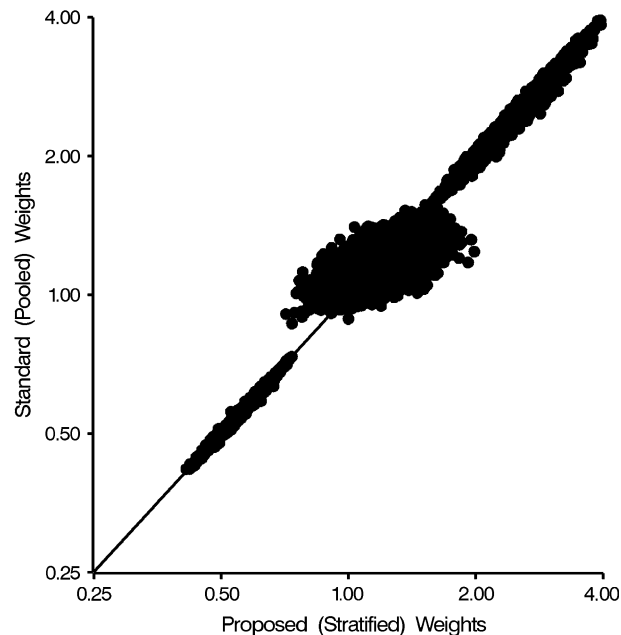

Figure 2. Scatterplot of standard (pooled) inverse probability weights against proposed (time-stratified) inverse probability weights, for example, scenario 2 (Figure 1C).

effect measure modification, because we refer to the magnitude of the effect of the confounder on the treatment (or outcome) changing over time, whereas the latter refers to the effect of the treatment on outcome differing by levels of the modifier. It is possible to conceive of a time-modified confounder that is also an effect modifier, but that is beyond the scope of this paper.

Time-modified confounding may often be present in epidemiologic studies of time-varying treatments. For example, Figure 1E with one modification (the arrow from $X(0)$ to $Z(1)$ would be absent) could represent an observational study of a time-varying treatment $X(t)$, such as use of a pharmacologic therapy with a time-fixed confounder, for example, sex (such that $Z(0)$ and $Z(1)$ will be the same value for each participant, $Z(t) = Z$). In this setting, being male may be positively associated with use of a novel pharmacologic agent (at time 0), but as inequities in use become apparent and corrected, sex may no longer be associated with use of the agent at time 1.

Even in the absence of time-modified confounding and assuming that the number of time points is not excessive, there may be little loss in using the stratified weights (i.e., in assuming that time-modified confounding is present). In Figure 2, with the exception of individuals whose treatment is always as predicted by the model, the pooled weights are slightly smaller than the stratified weights, as evidenced by a slightly smaller (but close to 1.0) mean weight (not shown). The slightly smaller, but less variable, standard weights appear to trade a small amount of bias for improved efficiency. More work is needed to explore the finite-sample properties of marginal structural models.

We have provided an example indicating the potential for time-modified confounding to affect estimates of causal effects and have emphasized the need to correctly specify the treatment model when fitting marginal structural models. Our example and simulations illustrate the magnitude of bias possible in typical epidemiologic settings, but they are not exhaustive. More simulations and worked examples are needed. In conclusion, when the *effect* of a time-fixed or time-varying confounder on either treatment or outcome changes over time, attention must be given to the possibility of time-modified confounding.

REFERENCES

1. Greenland S, Morgenstern H. Confounding in health research. *Annu Rev Public Health*. 2001;22:189–212.
2. Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods—application to control of the healthy worker survivor effect. *Math Model*. 1986; 7:1393–1512.
3. Robins J. A graphical approach to the identification and estimation of causal parameters in mortality studies with sustained exposure periods. *J Chronic Dis*. 1987;40(suppl 2):139S–161S.
4. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology*. 2000;11(5):550–560.
5. Robins JM, Hernán MA, Rotnitzky A. Effect modification by time-varying covariates. *Am J Epidemiol*. 2007;166(9):994–1002.
6. Platt RW, Joseph KS, Ananth CV, et al. A proportional hazards model with time-dependent covariates and time-varying effects for analysis of fetal and infant death. *Am J Epidemiol*. 2004;160(3):199–206.
7. Abrahamowicz M, MacKenzie T, Esdaile JM. Time-dependent hazard ratio: modeling and hypothesis testing with application in lupus nephritis. *J Am Stat Assoc*. 1996;91(436):1432–1439.
8. Altman DG, De Stavola BL. Practical problems in fitting a proportional hazards model to data with updated measurements of the covariates. *Stat Med*. 1994;13(4):301–341.
9. Shepherd BE. The cost of checking proportional hazards. *Stat Med*. 2008;27(8):1248–1260.
10. Cox DR. Regression models and life tables (with discussion). *J R Stat Soc (B)*. 1972;34:187–220.
11. Holland PW. Statistics and causal inference. *J Am Stat Assoc*. 1986;81(396):945–960.
12. Hernán MA. A definition of causal effect for epidemiological research. *J Epidemiol Community Health*. 2004;58(4):265–271.
13. Greenland S. Absence of confounding does not correspond to collapsibility of the rate ratio or rate difference. *Epidemiology*. 1996;7(5):498–501.
14. Miettinen OS, Cook EF. Confounding: essence and detection. *Am J Epidemiol*. 1981;114(4):593–603.
15. Pearl J. Causal diagrams for empirical research. *Biometrika*. 1995;82(4):669–688.
16. Hernán MA, Hernández-Díaz S, Werler MM, et al. Causal knowledge as a prerequisite for confounding evaluation: an application to birth defects epidemiology. *Am J Epidemiol*. 2002;155(2):176–184.
17. Robins J. Data, design and background knowledge in etiologic inference. *Epidemiology*. 2001;12(3):313–320.
18. Robins JM, Armitage P, Colton T. Structural nested failure time models. In: *The Encyclopedia of Biostatistics*. Hoboken, NJ: Wiley; 1997:4372–4389.
19. Hu FB, Stampfer MJ, Rimm E, et al. Dietary fat and coronary heart disease: a comparison of approaches for adjusting for total energy intake and modeling repeated dietary measurements. *Am J Epidemiol*. 1999;149(6):531–540.
20. Vanderweele TJ, Hernán MA, Robins JM. Causal directed acyclic graphs and the direction of unmeasured confounding bias. *Epidemiology*. 2008;19(5):720–728.
21. Vanderweele TJ. The sign of the bias of unmeasured confounding. *Biometrics*. 2008;64(3):702–706.
22. Hernán MA, Brumback B, Robins JM. Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology*. 2000;11(5):561–570.
23. Cole SR, Frangakis CE. The consistency statement in causal inference: a definition or an assumption? *Epidemiology*. 2009;20(1):3–5.
24. Hernán MA, Taubman SL. Does obesity shorten life? The importance of well-defined interventions to answer causal questions. *Int J Obes (Lond)*. 2008;32(suppl 3):S8–S14.
25. Cole SR, Hernán MA. Constructing inverse probability weights for marginal structural models. *Am J Epidemiol*. 2008;168(6):656–664.
26. Greenland S, Robins J. Identifiability, exchangeability, and epidemiological confounding. *Int J Epidemiol*. 1986;15(3):413–419.
27. D'Agostino RB, Lee ML, Belanger AJ, et al. Relation of pooled logistic regression to time dependent Cox regression

analysis: the Framingham Heart Study. *Stat Med*. 1990;9(12): 1501–1515.

28. Lefebvre G, Delaney JA, Platt RW. Impact of mis-specification of the treatment model on estimates from a marginal structural model. *Stat Med*. 2008;27(18):3629–3642.

29. 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, NY: Springer; 1999:95–134.

30. Moodie EEM, Platt RW, Kramer MS. Estimating response-maximized decision rules with applications to breastfeeding. *J Am Stat Assoc*. 2009;104(485):155–165.

31. Kramer MS, Chalmers B, Hodnett ED, et al. Promotion of Breastfeeding Intervention Trial (PROBIT): a randomized trial in the Republic of Belarus. *JAMA*. 2001;285(4):413–420.

32. Maldonado G, Greenland S. Simulation study of confounder-selection strategies. *Am J Epidemiol*. 1993;138(11):923–936.

33. White H. *Estimation, Inference, and Specification Analysis*. Cambridge, United Kingdom: Cambridge University Press; 1993.

APPENDIX

Details of the Simulation Study

Example data were generated following a Markovian decomposition of Figure 1, C–E, respectively. Five thousand samples each of 600 patients per sample were generated. First, $Z(0)$ was taken as a Bernoulli random variable, with probability of 0.50. Second, $X(0)$ was taken as a Bernoulli random variable with probability of $1/(1 + \exp\{-[-\log(2) + \log(4) \times z(0)]\})$, such that the marginal probability was approximately 0.50. Third, $Z(1)$ was taken as a Bernoulli random variable with probability of $1/(1 + \exp\{-[-\lambda_0 + \lambda_1 x(0)]\})$, where in scenario 1 $\lambda_0 = \lambda_1 = \log(1)$ and in scenarios 2 and 3 $\lambda_0 = \log(4)$ and $\lambda_1 = \log(8)$, such that the marginal probability of $Z(1)$ was approximately 0.50. Fourth, $X(1)$ was taken as a Bernoulli random variable with probability of $1/(1 + \exp\{-[-\eta_0 + \eta_1 x(0)]\})$, where in scenarios 1 and 2 $\eta_0 = \log(2)$ and $\eta_1 = \log(4)$ and in scenario 3 $\eta_0 = \eta_1 = \log(1)$, such that the marginal probability of $Z(1)$ was approximately 0.50. Fifth and last, $Y(2)$ was taken as a Bernoulli random variable with probability of $1/(1 + \exp\{-[-\log(40) + \log(6) \times z(0) + \log(2) \times z(1)]\})$, such that the marginal probability was approximately 0.10.

In the generated data, treatment $X(t)$ has no direct causal effects on $Y(2)$, as in Figure 1, C–E, but in scenarios 2 and 3 (Figure 1, D and E), $X(0)$ does have an indirect causal effect on $Y(2)$ mediated through covariate $Z(1)$. In scenario 1, there is no total (direct and indirect) causal effect of $X(t)$ on $Y(2)$. In scenarios 2 and 3, the total causal effect of $X(t)$ on $Y(2)$ is therefore equal to the direct causal effect of $X(0)$ on $Y(2)$. For reference, in scenarios 2 and 3, we obtain the total causal effect of $X(t)$ on $Y(2)$ as $\simeq \log(0.3)$ by use of the Kullback-Leibler Information Criterion coefficient (32). Generally, this coefficient is the maximum likelihood estimate for a specified model, such that it is the closest possible estimate to the true maximum likelihood estimate (33). Here, the Kullback-Leibler Information Criterion coefficient estimates the population-average cause effect. The Kullback-Leibler Information Criterion parameter was obtained from a logistic model for $Y(2)$ regressed on $X(0)$ with inverse probability-of-$X(0)$-treatment weights conditional only on $Z(0)$, with a sample size of 1 million patients.