Biometrics 60, 315–323
June 2004

Estimating Mean Response as a Function of Treatment Duration in an Observational Study, Where Duration May Be Informatively Censored

Brent A. Johnson

Department of Biostatistics, University of North Carolina, Chapel Hill, North Carolina 27599-7420, U.S.A.
email: bjohnson@bios.unc.edu

and

Anastasios A. Tsiatis

Department of Statistics, North Carolina State University, Raleigh, North Carolina 27695, U.S.A.
email: tsiatis@stat.ncsu.edu

Summary. After a treatment is found to be effective in a clinical study, attention often focuses on the effect of treatment duration on outcome. Such an analysis facilitates recommendations on the most beneficial treatment duration. In many studies, the treatment duration, within certain limits, is left to the discretion of the investigators. It is often the case that treatment must be terminated prematurely due to an adverse event, in which case a recommended treatment duration is part of a policy that treats patients for a specified length of time or until a treatment-censoring event occurs, whichever comes first. Evaluating mean response for a particular treatment-duration policy from observational data is difficult due to censoring and the fact that it may not be reasonable to assume patients are prognostically similar across all treatment strategies. We propose an estimator for mean response as a function of treatment-duration policy under these conditions. The method uses potential outcomes and embodies assumptions that allow consistent estimation of the mean response. The estimator is evaluated through simulation studies and demonstrated by application to the ESPRIT infusion trial coordinated at Duke University Medical Center.

Key words: Censored covariates; Censored treatment; Confounding; Infusion length; Inverse weighting; Observational study; Propensity score; Survival analysis.

1. Introduction

The ESPRIT (Enhanced Suppression of the Platelet IIb/IIIa Receptor with Integrilin Therapy) trial, which motivated this research, targeted patients with coronary artery disease scheduled to undergo percutaneous coronary intervention with stent implantation in a native coronary artery. The main objective of ESPRIT was to compare eptifibatide (Integrilin) therapy to placebo on the basis of the composite binary endpoint of death, myocardial infarction (MI), or urgent target vessel revascularization within 30 days. The study enrolled 2064 eligible patients who were randomized to either study drug (1040) or placebo (1024) regimen. The experimental treatment regimen consisted of an eptifibatide bolus and a continuous eptifibatide infusion for 18–24 hours, with a similar regimen for the placebo group. The study protocol also required that patients experiencing serious complications, such as abrupt closure, no reflow, or coronary thrombosis immediately discontinue the infusion process to receive appropriate medical attention; we define these protocol-defined adverse events as infusion-terminating events, or more generally as treatment-terminating events.

The main study analysis suggested that the eptifibatide regimen in the study is superior to placebo. The event proportion for the composite endpoint was 10.5% for placebo versus 6.8% for eptifibatide ($p = 0.0034$). A natural follow-up question posed by the investigators involved the length of infusion of eptifibatide that should be recommended for future patients. Because infusion cannot continue after a treatment-terminating event, a recommendation to infuse for t units of time necessarily implies that treatment would be discontinued after drug was administered for t units of time or when a treatment-terminating event occurs, whichever comes first. Thus, more precisely, the investigators were interested in comparing the event proportions for different infusion length "policies," where a "policy" dictates a particular infusion length t that may be possibly censored by events that require treatment termination.

Formally then, interest focuses on characterizing the mean response, here equal to the probability of whether a patient experiences an event (death, MI, or target vessel revascularization) within 30 days, as a function of treatment-duration (infusion length) policy. For a particular policy involving

duration t, then, the parameter of interest may be conceptualized as the mean response if all patients in the population were to be treated according to this policy. Clearly, if a study were conducted in which patients were assigned at random to receive different treatment-duration policies at the beginning of the study, estimation of this parameter for a particular policy would be straightforward following an intention-to-treat principle.

Data from a study that did not randomize patients to different treatment-duration policies, such as the data from patients in the ESPRIT trial who received eptifibatide, include the length of infusion for each patient and whether the infusion was terminated because of a protocol-defined adverse event or because of physician discretion. One difficulty in estimating the mean response for different treatment-duration policies using the observed data is that, unlike a well-designed randomized treatment-duration study, when a patient has their treatment terminated because of an adverse event then we do not know what treatment-duration policy was intended for that patient. Also, as in most observational studies, because infusion length was left to the discretion of the physician rather than dictated by design (randomization), patients receiving different durations may not be prognostically similar. It is intuitively apparent that simply averaging responses among individuals observed to have completed infusion of length t, either including or excluding patients experiencing a treatment-terminating event prior to t, would yield a biased estimator of the desired policy mean.

In this article, we demonstrate how methods developed for estimating causal parameters in observational studies (e.g., Robins et al., 1994, 2000; Hernan et al., 2000, 2001; Satten et al., 2001) may be used for our problem. Specifically, we show that a treatment-duration policy is a specific example of a dynamic treatment regime and delineate the conditions and assumptions that enable us to use the theory developed by Murphy, van der Laan, and Robins (2001) to derive a consistent estimator of the mean response as a function of treatment-duration policy from a sample of observed data such as those in the ESPRIT infusion trial.

This article is organized as follows. The assumptions and methods are developed in Section 2, and large sample properties of the estimator are outlined in Section 3. In Section 4, we apply the procedure to the ESPRIT infusion data. We report on simulation studies to evaluate small sample properties in Section 5.

2. Assumptions and Methods

Although the ESPRIT trial involves a dichotomous outcome, the following development allows for a general response variable that could be continuous or discrete. In order to conceptualize properly the question of interest, it is useful to introduce the idea of potential (counterfactual) random variables (Rubin, 1974; Rosenbaum and Rubin, 1983). Specifically, we define the potential random variable C to represent the time at which a randomly selected individual from our population, if continuously treated, would have a treatment-terminating event. We also define the potential random process $(Y_t^\star, t \le C)$, where $Y_t^\star$ denotes the response of the individual, if treatment were terminated at time t, for $t \le C$. In terms of these potential random variables, the policy of treating an individual for t units of time or until a treatment-terminating event results in the response $Y^\star_{t\wedge C}$, where $t \wedge C$ denotes the minimum of t and C. This may also be written as

$$Y^\star_{t\wedge C} = Y_t^\star I(C \ge t) + Y_C^\star I(C < t), \tag{1}$$

and the parameter of interest is the population mean response for this treatment-duration policy, namely, $\mathrm{E}(Y^\star_{t\wedge C})$.

The variables defined above are referred to as potential random variables, or counterfactuals, because, contrary to the fact, they may not actually be observed. In contrast, for a randomly selected individual from our population, the observable random variables are (Y, U, Δ), where Y denotes the observed response, U denotes the actual treatment duration, and Δ is an indicator variable such that $\Delta = 0$ if treatment duration was stopped due to a treatment-terminating event and $\Delta = 1$ if stopped due to physician discretion.

For simplicity, we will assume that when patients do not experience treatment-terminating events ($\Delta = 1$), then their actual treatment duration U realizes one of k finite values, $t_1, \ldots, t_k$. However, when treatment is stopped because of an intervening event ($\Delta = 0$), U can take values along a continuum of time. Under these conditions, the specific parameters on which we focus are the population mean responses for the k treatment-duration policies, $\mu_j = \mathrm{E}(Y^\star_{t_j\wedge C}), j = 1, \ldots, k$.

A key assumption is that the observed response Y may be written in terms of the potential outcomes as

$$Y = \left\{ \sum_{j=1}^{k} Y^\star_{t_j} I(U = t_j, \Delta = 1) \right\} + Y_C^\star I(\Delta = 0). \tag{2}$$

In words, assumption (2) says that if a patient experiences a treatment-terminating event, then the observed response is $Y_C^\star$; otherwise, the observed response is $Y^\star_{t_j}$ corresponding to the realized value of actual treatment duration U.

Along with the data (Y, U, Δ), we assume that additional, possibly time-dependent covariate, information is also available. Let $Z(u)$ denote the value of a vector of covariates for an individual at time u, where u is measured as time from the individual's entry into the study. We define $\bar{Z}(x)$ to be the history of covariate information up to and including time x, i.e., $\{Z(u),\, u \le x\}$, and use $\bar{Z}_j$ to denote $\bar{Z}(t_j)$. We define the jth discrete cause-specific hazard function to be the conditional probability $\lambda_j(\bar{Z}_j) = P(U = t_j, \Delta = 1 \,|\, U \ge t_j, \bar{Z}_j), j = 1, \ldots, k$. The key assumption that allows us to estimate consistently the parameter $\mu_j = \mathrm{E}(Y^\star_{t_j\wedge C})$ from a random sample of observed data is given by

$$P\big(U = t_j, \Delta = 1 \,|\, U \ge t_j, \bar{Z}_j, C, Y_x^\star, t_j \le x \le C\big) = \lambda_j(\bar{Z}_j), \quad j = 1, \ldots, k. \tag{3}$$

In words, assumption (3) implies that the decision to terminate or continue a patient's treatment at time t_j—given the patient has continuously received treatment up to and including time t_j without a treatment-censoring event, and given the patient's covariate history up to and including time t_j—does not depend on future prognosis. Such an assumption is plausible if information about an individual through time t_j, which may be prognostic and which an investigator may use to make decisions on treatment duration, are captured in

the data $\bar{Z}_j$. The assumption given by (3) is the "sequential randomization assumption" or the assumption of "no unmeasured confounders" discussed by Robins (1997).

It is shown in detail in the Appendix that the treatment-duration policy "treat for t_j units of time" is a specific example of a dynamic treatment regime as defined by Murphy et al. (2001) and μ_j is the mean response for this dynamic treatment regime. As such, under the sequential randomization assumption (3), we also demonstrate in the Appendix that the theory developed by Murphy et al. (2001) can be used to find a consistent estimator for μ_j. Specifically, we show that

$$\mathrm{E}\left[(Y-\mu_j)\times\left\{\frac{I(U=t_j,\Delta=1)}{f_j(\bar{Z}_j)}+\frac{I(U<t_j,\Delta=0)}{K_{[U]}\big(\bar{Z}_{[U]}\big)}\right\}\right]=0, \tag{4}$$

under assumption (3), where

$$f_j(\bar{Z}_j)=\lambda_j(\bar{Z}_j)\prod_{m=1}^{j-1}\{1-\lambda_m(\bar{Z}_m)\} \tag{5}$$

and

$$K_j(\bar{Z}_j)=\prod_{m=1}^{j}\{1-\lambda_m(\bar{Z}_m)\}. \tag{6}$$

We also define $[U]=\max_j\{j:t_j<U\}$; therefore, $\bar{Z}_{[U]}$ refers to all covariate information up to and including time t_j just prior to U, and $K_{[U]}(\bar{Z}_{[U]})$ will be the product of $(1-\lambda_m)$ from $m=1$ to $m=[U]$.

Thus for a sample of data $\{Y_i,U_i,\Delta_i,\bar{Z}(U_i),i=1,\ldots,n\}$, if the probabilities $f_j(\bar{Z}_j)$ and $K_{[U]}(\bar{Z}_{[U]})$ were known, then a natural estimator for μ_j would be obtained by solving the estimating equation

$$\sum_{i=1}^{n}(Y_i-\hat{\mu}_{jn})\left\{\frac{I(U_i=t_j,\Delta_i=1)}{f_j(\bar{Z}_{ij})}+\frac{I(U_i<t_j,\Delta_i=0)}{K_{[U_i]}\big(\bar{Z}_{[U_i]}\big)}\right\}=0, \tag{7}$$

where $\bar{Z}_{ij}$ refers to the covariate information up to and including time t_j for the ith individual. This yields the estimator

$$\hat{\mu}_{jn}=\frac{\sum_{i=1}^{n}Y_iw_{ij}}{\sum_{i=1}^{n}w_{ij}},$$

$$w_{ij}=\frac{I(U_i=t_j,\Delta_i=1)}{f_j(\bar{Z}_{ij})}+\frac{I(U_i<t_j,\Delta_i=0)}{K_{[U_i]}\big(\bar{Z}_{[U_i]}\big)},$$

which is a weighted average of the responses.

Remark. If we view the probability $f_j(\bar{Z}_{ij})$ as the propensity score (Rosenbaum and Rubin, 1983) for the ith individual to have treatment terminated at time t_j, then, if there were no treatment-terminating events, the estimator would equal $\sum Y_iw_{ij}/\sum w_{ij}$, where $w_{ij}=I(U_i=t_j)/f_j(\bar{Z}_{ij})$. This is a weighted average of the responses for individuals with treatment duration t_j and weights equal to the inverse of their propensity score. Such inverse propensity score estimators have been suggested by Cassel, Särndal, and Wretman (1983) and Rosenbaum (1987) to adjust for confounding of treatment with baseline time-independent covariates. With censoring, the weighted average also includes responses from individuals who have a treatment-censoring event at time $C<t_j$. However, their contribution is weighted by

$$\frac{1}{K_{[U_i]}\big(\bar{Z}_{[U_i]}\big)}=\frac{1}{f_j(\bar{Z}_{ij})}\frac{f_j(\bar{Z}_{ij})}{K_{[U_i]}\big(\bar{Z}_{[U_i]}\big)}.$$

Intuitively, this weight can be viewed as the inverse propensity score multiplied by the conditional probability that the individual would have treatment stopped at time t_j given that treatment duration was known to be greater than C. Heuristically, this shows how the response of an individual who has treatment censored at time C is "distributed to the right" (Blight, 1970; Turnbull, 1974, 1976) to estimate μ_j for all $\{j:t_j>C\}$.

Because (7) contains the unknown probabilities $f_j(\bar{Z}_j)$ and $K_{[U]}(\bar{Z}_{[U]})$, which are functions of the unknown parameters $\lambda_j(\bar{Z}_j),j=1,\ldots,k$, these probabilities must be estimated from the data. This requires that we posit a model for the discrete hazards $\lambda_j(\bar{Z}_j,\gamma),j=1,\ldots,k-1$, as a function of a finite-dimensional parameter vector γ. Assuming (3) holds, it is straightforward to derive the observed-data likelihood of $D_i=\{U_i,\Delta_i,\bar{Z}(U_i)\},i=1,\ldots,n$, as $\prod_{i=1}^{n}L(\gamma;D_i)$, where

$$L(\gamma;D_i)=\prod_{j=1}^{k-1}\left\{\frac{\lambda_{ij}(\gamma)}{1-\lambda_{ij}(\gamma)}\right\}^{I(U_i=t_j,\Delta_i=1)}\{1-\lambda_{ij}(\gamma)\}^{I(U_i\geq t_j)}, \tag{8}$$

and $\lambda_{ij}(\gamma)=\lambda_j(\bar{Z}_{ij},\gamma)$. The estimator for γ is obtained by maximizing this likelihood. Of course, the exact form of (8) depends on how we model $\lambda_j(\bar{Z}_j)$ as a function of $\bar{Z}_j$.

Generalized linear models are often used for their interpretability and general applicability. We consider one such class of generalized linear models, where

$$\lambda_j(\bar{Z}_j)=F\big(\alpha_j+\beta_j^t\bar{Z}_j\big),\quad j=1,\ldots,k-1,$$

and $F(t)$ is the logistic function, i.e., $e^t/(1+e^t)$. This is similar to the continuation ratio logit model (Agresti, 1990, p. 319), in which case (8) becomes

$$\prod_{j=1}^{k-1}\prod_{i\in R_j}\frac{\exp\big\{\big(\alpha_j+\beta_j^t\bar{Z}_{ij}\big)I(U_i=t_j,\Delta_i=1)\big\}}{1+\exp\big(\alpha_j+\beta_j^t\bar{Z}_{ij}\big)},$$

where $R_j=\{i:U_i\geq t_j\}$. This likelihood is similar to Cox's partial likelihood for continuous time proportional hazards models.

To summarize, the proposed estimator for μ_j, $j=1,\ldots,k$ is given by

$$\hat{\mu}_{jn}=\frac{\sum_{i=1}^{n}Y_i\hat{w}_{ij}}{\sum_{i=1}^{n}\hat{w}_{ij}},$$

$$\hat{w}_{ij}=\frac{I(U_i=t_j,\Delta_i=1)}{f_j(\bar{Z}_{ij},\hat{\gamma})}+\frac{I(U_i<t_j,\Delta_i=0)}{K_{[U_i]}\big(\bar{Z}_{[U_i]},\hat{\gamma}\big)},$$

where $\hat{\gamma}$ is the maximum likelihood estimator which maximizes (8).

3. Large Sample Properties

We derive the large sample properties of $\hat{\mu}_{jn}$ under the assumption that when $\Delta = 1$, actual treatment duration, U, can take only one of a finite number of values $t_1 < \cdots < t_k$. Hence, k is assumed fixed as the number of patients n goes to infinity. As mentioned previously, this is an approximation to the truth when $\Delta = 1$ and U is, in fact, continuous and the data are grouped into k categories by partitioning treatment duration into intervals, with t_j representing the midpoint of the jth interval. A rigorous approach to this problem would assume that $\mathrm{E}\{Y^{\star}_{t\wedge C}\}$ is a smooth function of t and that the number of intervals k increases as a function of n. Such a technical development is beyond the scope of the paper and, we believe, would not add additional practical insight into the problem. We make additional comments on the effect of partitioning in Section 6.

Because the proposed estimator $\hat{\mu}_{jn}$ uses the estimated discrete hazard $\lambda_j(\bar{Z}_{ij}, \hat{\gamma})$, it is convenient to define the estimator as the first element in the solution to the system of equations

$$\sum_{i=1}^{n} \begin{pmatrix} \psi_{\mu_j}(Y_i, D_i, \hat{\mu}_{jn}, \hat{\gamma}_n) \\ \psi_{\gamma}(D_i, \hat{\gamma}_n) \end{pmatrix} = 0, \tag{9}$$

where

$$\psi_{\mu_j}(Y_i, D_i, \mu_j, \gamma) = (Y_i - \mu_j)\left\{ \frac{I(U_i = t_j, \Delta = 1)}{f_j(\bar{Z}_{ij}; \gamma)} + \frac{I(U_i < l_j, \Delta = 0)}{K_{[U_i]}\big(\bar{Z}_{[U_i]}; \gamma\big)} \right\},$$

$$\psi_{\gamma}(D_i, \gamma) = \frac{\partial}{\partial \gamma} \log L(\gamma; D_i).$$

Note that we have characterized the proposed estimator as an M-estimator (Huber, 1964), whose asymptotic properties are well known. Hence, under suitable regularity conditions, $\hat{\mu}_{jn}$ can be shown to be consistent and asymptotically normal when the model for the discrete hazards (8) is correctly specified.

A consistent estimator for the asymptotic variance of the limiting normal distribution can be derived using standard arguments (Carroll, Ruppert, and Stefanski, 1995, Section A.3.6) and is given by

$$n^{-1} \sum_{i=1}^{n} \left[(Y_i - \hat{\mu}_{jn})^2 \left\{ \frac{I(U_i = t_j, \Delta = 1)}{f_j^2(\bar{Z}_{ij}; \hat{\gamma}_n)} + \frac{I(U_i < t_j, \Delta = 0)}{K^2_{[U_i]}\big(\bar{Z}_{[U_i]}; \hat{\gamma}_n\big)} \right\} - \hat{H}_i \big\{ \hat{\mathrm{E}}\big(S_{\gamma} S_{\gamma}^{t}\big) \big\}^{-1} \hat{H}_i^{t} \right],$$

where $\hat{\mathrm{E}}(S_{\gamma} S_{\gamma}^{t})$ is a consistent estimator of the Fisher information for γ, and

$$\hat{H}_i = (Y_i - \hat{\mu}_{jn}) \left\{ \frac{I(U_i = t_j, \Delta = 1) \dfrac{\partial f_j(\bar{Z}_{ij})}{\partial \gamma}}{f_j^2(\bar{Z}_{ij}; \hat{\gamma}_n)} + \frac{I(U_i < t_j, \Delta = 0) \dfrac{\partial K_{[U_i]}\big(\bar{Z}_{[U_i]}\big)}{\partial \gamma}}{K^2_{[U_i]}\big(\bar{Z}_{[U_i]}; \hat{\gamma}_n\big)} \right\}.$$

The asymptotic results above pertain to the marginal distribution of $\hat{\mu}_{jn}$ for a fixed j. It would be straightforward to consider the system of estimating equations

$$\sum_{i=1}^{n} \begin{pmatrix} \psi_{\mu_1}(Y_i, D_i, \hat{\mu}_{1n}, \hat{\gamma}_n) \\ \vdots \\ \psi_{\mu_k}(Y_i, D_i, \hat{\mu}_{kn}, \hat{\gamma}_n) \\ \psi_{\gamma}(D_i, \hat{\gamma}_n) \end{pmatrix} = 0, \tag{10}$$

simultaneously, in order to derive the joint asymptotic normal distribution of $(\hat{\mu}_{1n}, \ldots, \hat{\mu}_{kn})$. This may be useful, for example, if formal tests of contrasts of mean response for the different treatment-duration policies are desired.

4. Analysis of the ESPRIT Infusion Trial

We demonstrate the proposed methods by application to data from patients in the ESPRIT trial who received eptifibatide. The outcome of interest is a composite endpoint of death, MI, or revascularization within 30 days of the initiation of treatment. The data are discretized by taking t_j to be the midpoint of five intervals I_j, namely $I_j = \{(t_{j-1} + t_j)/2, (t_j + t_{j+1})/2\}$ for $\mathbf{t} = (t_1, t_2, t_3, t_4, t_5) = (16, 18, 20, 22, 24)$, and we redefine $U_i = t_j$ for any patient, where $U_i \in I_j$ and $\Delta_i = 1$. Four patients did not have both observed random variables (U_i, Δ_i) and were excluded from all subsequent analyses. There are seven patients who completed infusion before 15 hours and are assigned to the first group ($U_i = 16$), and seven patients who completed infusion after 25 hours and are assigned to the last group ($U_i = 24$). The inclusion or exclusion of these 14 observations does not appreciably change the results. The frequency for the number of patients who completed infusion at t_j is 61, 479, 194, 85, and 111 for $j = 1, \ldots, 5$. Among patients censored because of an infusion-terminating event, 89 were censored before 16 hours, 11 between 16 and 18 hours, and 6 patients between 18 and 20 hours.

As in most observational studies, the assumption of no unmeasured confounders is key in deriving unbiased estimators for the mean response as a function of treatment-duration policy. As such, the ESPRIT infusion trial investigators were asked to identify factors that they believed would influence treatment duration. The investigators confirmed that treatment would be discontinued with certainty if a patient experienced a protocol-defined adverse event, but otherwise, could not identify any variables that they believed would affect the decision of the participating physicians in any systematic way to terminate treatment. Thus the investigators believed that there were no obvious measured or unmeasured confounders. We further investigated the issue of measured and unmeasured confounders using a series of analyses that we will describe shortly.

Numerous measurements were collected for every patient enrolled in the ESPRIT trial, only a few of which appear prognostic for both infusion length and the 30-day endpoint. The baseline variables that we include in our analyses are diabetes (0/1), percutaneous transluminal coronary angioplasty (PTCA, 0/1), angina (0/1), heparin (0/1), and weight (in kg). A descriptive summary of these variables is given in Table 1.

One potentially important time-dependent covariate included in our analysis was enzyme level. Enzyme levels were

Table 1
Summary statistics for the ESPRIT trial

	Event within 30 days	
	No ($n = 77$)	Yes ($n = 959$)
Diabetes	0.17	0.20
PTCA	0.34	0.22
Angina	0.31	0.40
Heparin	0.18	0.13
Weight	83.99	85.17

recorded at baseline and then at each subsequent 6-hour interval up to and including 24 hours. We defined the time-dependent covariate at time t_j as the last enzyme level measured just prior to time t_j.

Naive estimators for μ_j commonly employed in practice include an overall event proportion for all patients belonging to the interval I_j, which we denote by T_{1j}, and an uncensored event proportion for all patients completing infusion in I_j, which we denote by T_{2j}. In Table 2, we present the results of the estimated event proportions for the two naive estimators and estimators proposed in this article. As part of a sensitivity analysis we considered three different estimators that make different assumptions about the confounding relationship—$\hat{\mu}_{jn}^{(0)}$ assumes no confounding is present, $\hat{\mu}_{jn}^{(1)}$ assumes confounding is present only through baseline covariates, and $\hat{\mu}_{jn}^{(2)}$ assumes the possibility of time-dependent confounding; thus, this last estimator includes the time-dependent covariate related to enzyme level in addition to the baseline covariates. We immediately note that for the first interval T_{11} is much larger than T_{21}. This can be explained by the higher event proportion among censored patients compared to uncensored patients, i.e., 0.189 compared to 0.061, and by the fact that 89 patients were censored at or before 16 hours while only 61 patients completed physician-recommended infusion at 16 hours. As might be expected under these conditions, the proposed estimator $\hat{\mu}_{jn}$ is greater than the naive uncensored event proportions, T_{2j}, for all $j = 1, \ldots, k$.

In Table 2, we also give the estimated standard errors for $\hat{\mu}_{jn}$ derived in Section 3. We note that the three estimators $\hat{\mu}_{jn}^{(0)}$, $\hat{\mu}_{jn}^{(1)}$, and $\hat{\mu}_{jn}^{(2)}$ are in good agreement which is consistent with the intuition of the study investigators that no strong measured confounders exist in the data set. Based on the results from Table 2, there is a strong suggestion that infusing patients for more than 16 hours does not improve the event rate and could possibly be detrimental to the patients.

Verifying the conjecture that assumption (3) holds, that is, that there are no unmeasured confounders, is a much more delicate issue since this assumption is inherently nonidentifiable from the observed data. To obtain indirect evidence of the validity of this assumption we analyzed the data from the placebo group. It may be reasonable to assume that placebo has no effect on outcome. That is, the counterfactual response at time t for an arbitrary individual in our population receiving placebo, namely Y_t^{*P}, is independent of t for $t \leq C^P$, where C^P represents the time that this individual, if continuously treated with placebo, would have a treatment-terminating event. One may view this assumption as a causal null hypothesis. Under such an assumption, the counterfactual mean response $\mu_j^P = \mathrm{E}(Y_{t_j \wedge C^P}^{*P})$ for the treatment-duration policy t_j would be the same for all j. Consequently, if an analysis of the placebo patients showed that the estimates for μ_j^P were sufficiently different, then we might conclude that there are unmeasured confounders not properly accounted for. We estimated μ_j^P in the placebo group using the identical methods and variables used in the group treated with eptifibatide. The results, summarized in Table 3, did not show evidence of a dose (treatment-duration) response relationship. We emphasize that this is not proof of *no* unmeasured confounders but only suggestive that this may be a reasonable assumption.

5. Simulation Results

In this section, we investigate the small sample properties of our estimator and explore the sensitivity of our estimator to the assumption of no unmeasured confounders. These simulations assume that patients are assigned to one of a finite number of treatment-duration policies at values $\mathbf{t} = (t_1, \ldots, t_4)$. Although the proposed method allows for time-dependent covariates, for simplicity, we only consider time-independent covariates.

In the first simulation, we let $\mathbf{t} = (15, 20, 25, 30)$. We considered a single covariate, Z_1, following a standard normal distribution, for each individual. We then generate a

Table 2
Estimated event proportions for the ESPRIT trial. T_{1j} denotes the overall event proportion for all patients belonging to interval I_j, T_{2j} denotes the uncensored event rate for patients belonging to I_j, $\hat{\mu}_{jn}^{(0)}$ is the proposed estimator assuming no confounding is present, $\hat{\mu}_{jn}^{(1)}$ is the proposed estimator assuming confounding is present through baseline factors only, and $\hat{\mu}_{jn}^{(2)}$ is the proposed estimator assuming time-dependent confounding. Standard errors are presented in parentheses.

j	t_j (hours)	T_{1j}	T_{2j}	$\hat{\mu}_{jn}^{(0)}$	$\hat{\mu}_{jn}^{(1)}$	$\hat{\mu}_{jn}^{(2)}$
1	16	0.140	0.018	0.047 (0.021)	0.040 (0.016)	0.044 (0.018)
2	18	0.047	0.046	0.065 (0.010)	0.066 (0.010)	0.070 (0.011)
3	20	0.068	0.068	0.079 (0.017)	0.078 (0.017)	0.078 (0.017)
4	22	0.050	0.050	0.071 (0.024)	0.071 (0.024)	0.067 (0.022)
5	24	0.124	0.124	0.116 (0.027)	0.121 (0.035)	0.109 (0.032)

Table 3

Estimated event proportions in the placebo arm. T_{1j} denotes the overall event rate for all patients belonging to interval I_j, T_{2j} denotes the uncensored event rate for patients belonging to I_j, $\hat{\mu}_{jn}^{(0)}$ is the proposed estimator assuming no confounding is present, $\hat{\mu}_{jn}^{(1)}$ is the proposed estimator assuming confounding is present through baseline factors only, and $\hat{\mu}_{jn}^{(2)}$ is the proposed estimator assuming time-dependent confounding. Standard errors are presented in parentheses.

j	t_j (hours)	T_{1j}	T_{2j}	$\hat{\mu}_{jn}^{(0)}$	$\hat{\mu}_{jn}^{(1)}$	$\hat{\mu}_{jn}^{(2)}$
1	16	0.187	0.097	0.106 (0.036)	0.116 (0.040)	0.131 (0.045)
2	18	0.073	0.071	0.083 (0.012)	0.083 (0.012)	0.091 (0.013)
3	20	0.116	0.116	0.125 (0.022)	0.125 (0.022)	0.126 (0.022)
4	22	0.070	0.070	0.081 (0.026)	0.079 (0.026)	0.079 (0.026)
5	24	0.187	0.187	0.191 (0.033)	0.170 (0.031)	0.124 (0.024)

treatment-censoring random variable C as $\exp\{\rho(Z)\}$ random variable, where

$$\rho(Z) = 0.005 \exp(\varphi_1 Z_1), \quad \varphi_1 = -2.$$

The treatment-duration data are simulated according to the following algorithm, which is consistent with the assumptions made: Start by letting $m = 1$.

1. If $C < t_m$, then define $U = C$ and $\Delta = 0$.
2. For $C \geq t_m$, generate a Bernoulli random variable Q_m, the indicator variable for stopping treatment at time t_m, with probability $\lambda_m(Z)$, where
$$\text{logit}\{\lambda_m(Z)\} = \alpha_m + \beta_1 Z_1.$$
3. If $Q_m = 1$, then assign $U = t_m$ and $\Delta = 1$; if $Q_m = 0$ and $m < k$, then increment m to $m + 1$ and go to Step 1.
4. If $U = t_j$ and $\Delta = 1$, then generate the corresponding response Y as a Bernoulli random variable with probability π, where
$$\text{logit}(\pi) = \eta_j + \zeta_1 Z_1,$$
whereas, if $U = C$ and $\Delta = 0$, then generate the corresponding response Y as a Bernoulli random variable with probability π, where
$$\text{logit}(\pi) = \min_{\{a: t_a \geq C\}} \eta_a + \zeta_1 Z_1 + \upsilon.$$

Note that the parameter υ is only present when $\Delta = 0$, thus inducing a dependence of censoring of treatment infusion on the probability of outcome. The population parameter of interest μ_j is difficult to evaluate analytically; thus, we approximated its value by simulation. Using the above algorithm, we forced treatment duration to be stopped at time t_j, if not already censored by replacing Step 2 with "$Q_m = 0$ for $m = 1, \ldots, j-1$ and $Q_j = 1$," generating the outcome Y 100,000 times, and then taking the sample average.

The chosen values of the parameters are as follows: $\alpha = (-1.2, -0.75, 0)$, $\beta_1 = -0.5$, $\eta = (-5, -4, -4, -2)$, $\zeta_1 = -2$, $\upsilon = 2$. For each data set, nominal 95% Wald confidence intervals were constructed using $\hat{\mu}_{jn}$, its standard error from Section 3, and a critical value of 1.96. Then, the empirical coverage probability (ECP) is calculated as the number of MC data sets where the true μ_j falls within the Wald confidence interval divided by the total number of MC data sets.

Table 4 presents simulation results for estimating mean response as a function of treatment-duration policy t_j, $j = 1, \ldots, 4$ when the no unmeasured confounders assumption is true. We use a sample size of $n = 1000$ and generate 1000 Monte Carlo data sets. Monte Carlo bias is <2% in every case and <1% in most cases; interval coverage is approximately the nominal level. The naive estimators, T_{1j} and T_{2j}, are biased and do not possess the correct coverage probability.

Next, we investigate the sensitivity of our estimator to deviations from the "no unmeasured confounders" assumption. For this, we include an additional covariate Z_2 to represent a potential unmeasured confounder. To induce additional confounding, we include an additional term $\beta_2 Z_2$ in the treatment choice model in Step 2 of the algorithm and an additional term

Table 4

Simulation summary of mean response for 1000 *Monte Carlo data sets when treatment-duration data are discrete. T_{1j} is the average Y for $U_i \in I_j$ and T_{2j} is the average Y for $(U_i \in I_j, \Delta_i = 1)$. $\hat{\mu}_{jn}$ is our estimator assuming confounding is present through Z_1. ECP is defined as the empirical coverage probability. Estimated standard errors are given in parentheses.*

t_j	μ_j	$\hat{\mu}_{jn}$	T_{1j}	T_{2j}	ECP $\hat{\mu}_{jn}$	ECP T_{1j}	ECP T_{2j}
t_1	0.055	0.056 (0.010)	0.118	0.045	0.949	0.029	0.808
t_2	0.091	0.091 (0.015)	0.099	0.068	0.947	0.927	0.703
t_3	0.100	0.099 (0.016)	0.064	0.046	0.949	0.441	0.099
t_4	0.151	0.151 (0.027)	0.063	0.052	0.938	0.005	0.000

Table 5

Simulation summary of mean response for 1000 *Monte Carlo data sets when the no unmeasured confounder assumption is violated. $\hat{\mu}_{jn}$ assumes confounding is present only through Z_1. ECPs are given in parentheses.*

		$\zeta_2 = -1$		$\zeta_2 = -0.5$	
t_j	β_2	μ_j	$\hat{\mu}_{jn}$(ECP)	μ_j	$\hat{\mu}_{jn}$(ECP)
	−0.5		0.094 (0.961)		0.087 (0.970)
t_1	−0.25	0.090	0.092 (0.974)	0.084	0.085 (0.961)
	0		0.090 (0.958)		0.085 (0.955)
	−0.5		0.126 (0.957)		0.118 (0.961)
t_2	−0.25	0.125	0.126 (0.966)	0.117	0.117 (0.963)
	0		0.126 (0.955)		0.117 (0.958)
	−0.5		0.130 (0.907)		0.126 (0.942)
t_3	−0.25	0.134	0.133 (0.924)	0.130	0.127 (0.949)
	0		0.137 (0.973)		0.129 (0.954)
	−0.5		0.152 (0.767)		0.152 (0.877)
t_4	−0.25	0.172	0.161 (0.871)	0.166	0.157 (0.917)
	0		0.173 (0.944)		0.164 (0.954)

$\zeta_2 Z_2$ in the prognostic model in Step 4 of the algorithm. We fix $\beta_1 = -0.5$, $\zeta_1 = -2$, and let the other parameters, α, η, υ, be the same as in the first simulation. Different values of β_2 and ζ_2 are considered to study the effect of the strength of confounding of the unmeasured variable Z_2 on the resulting estimator.

The estimator $\hat{\mu}_{jn}$ is computed using only the covariate Z_1. We again use a sample size of $n = 1000$ and 1000 Monte Carlo data sets. The simulation results are displayed in Table 5.

As expected, our estimator performs well when $\beta_2 = 0$ (no unmeasured confounders); however, even with modest confounding of the unmeasured covariate Z_2, the estimator performed relatively well.

6. Discussion

We have discussed methods to estimate the mean response as a function of treatment-duration policy, possibly right-censored, in an observational study. The methods use the theory of causal inference for time-dependent treatments that was developed by Robins and his colleagues based on representation of the problem in terms of potential outcomes and inverse probability weighted methods. Simulation studies show that the proposed estimator performs reliably well in realistic sample sizes when the assumptions underlying the theory hold.

As in most observational studies, the key assumption that allows us to derive estimators of causal parameters using the observed data is that of no unmeasured confounders. This assumption is, unfortunately, also the most difficult to verify. In order that this assumption be plausible we must have some degree of confidence that all important information that may affect the treatment decision process be captured in the database. For that reason, it is important during the design stage of such an observational study that discussions with investigators be carried out to identify all the factors that they believe would influence their treatment decisions and all possible effort be made to capture such information. In addition, various sensitivity analyses should be conducted such as those described in Section 4.

We derived methods where treatment-duration data are assumed to occur at finitely many time points—$t_1, \ldots, t_k$—in the population. When this is untrue, i.e., treatment duration occurs along a continuum of time, then, we suggest partitioning treatment duration (U) into intervals and estimating the mean response at the midpoint of the interval, as we did in the ESPRIT analysis. This approach seemed to work reasonably well in simulation scenarios (not shown here) analogous to those presented in Section 5.

Acknowledgements

This work was supported in part by NSF grant DMS 9819895, NIH grants R01-AI31789 and R01-CA51962, and NIEHS grant T32-ES07018. The authors are grateful to the associate editor, two reviewers, and Marie Davidian for providing helpful suggestions on earlier versions of the manuscript and Anne Hellkamp from the Duke Clinical Research Institute for providing the data used in the example.

Résumé

Après qu'un traitement a été trouvé efficace dans une étude clinique, l'attention se porte souvent sur l'influence de la durée du traitement sur le résultat. Une telle analyse facilite l'établissement de recommandations sur la durée la plus bénéfique. Dans beaucoup d'études, la durée du traitement est laissée, dans certaines limites, à l'appréciation des investigateurs. Il arrive souvent que le traitement doive être arrêté prématurément à cause d'un événement indésirable, auquel cas la recommandation d'une durée de traitement consiste à traiter les patients jusqu'à la première des deux dates suivantes :une date prédéfinie ou la date de survenue d'un événement indésirable. L'évaluation de la réponse moyenne pour un choix particulier d'une durée de traitement à partir de données d'observation est difficile en raison de la censure et parce qu'il n'est pas raisonnable de supposer que les différentes stratégies ont un pronostic similaire. Nous proposons un estimateur de la réponse moyenne comme fonction de la stratégie choisie pour la durée de traitement. La méthode fait appel à des résultats potentiels et s'appuie sur des hypothèses permettant une estimateur cohérent de la réponse moyenne. Cet estimateur est évalué par le biais de simulations et est appliqué à l'essai ESPRIT coordonné par le Centre médical de la Duke University.

References

Agresti, A. (1990). *Categorical Data Analysis*. New York: John Wiley.

Blight, B. (1970). Estimation from a censored sample for the exponential family. *Biometrika* **57,** 389–395.

Carroll, R. J., Ruppert, D., and Stefanski, L. A. (1995). *Measurement Error in Nonlinear Models*. Boca Raton, Florida: Chapman & Hall/CRC.

Cassel, C. M., Särndal, C. E., and Wretman, J. H. (1983). Some uses of statistical models in connection with the nonresponse problem. In *Incomplete Data in Sample Surveys*, W. G. Madow and I. Olkin (eds), Volume 3, 143–160. New York: Academic Press.

Hernan, M. A., Robins, J. M., and Brumback, B. (2000). Marginal structural models to estimate the causal effect of zidovudine on the survival of HIV-positive men. *Epidemiology* **11,** 561–570.

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

Huber, P. J. (1964). Robust estimation of a location parameter. *Annals of Mathematical Statistics* **35,** 73–101.

Murphy, S. A., van der Laan, M. J., and Robins, J. M. (2001). Marginal mean models for dynamic regimes. *Journal of the American Statistical Association* **96,** 1410–1423.

Robins, J. M. (1997). Causal inference from complex longitudinal data. In *Latent Variable Modeling and Applications to Causality. Lecture Notes in Statistics,* Volume 120, M. Berkane (ed), 69–117. New York: Springer Verlag.

Robins, J. M., Rotnizky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* **89,** 846–866.

Robins, J. M., Hernan, M. A., and Brumback, B. (2000). Marginal structural models and causal inference in epidemiology. *Epidemiology* **11,** 550–560.

Rosenbaum, P. R. (1987). Model-based direct adjustment. *Journal of the American Statistical Association* **82,** 387–394.

Rosenbaum, P. R. and Rubin, D. B. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika* **70,** 41–55.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* **6,** 688–701.

Satten, G. A., Datta, S., and Robins, J. M. (2001). Estimating the marginal survival function in the presence of time dependent covariates. *Statistics and Probability Letters* **54,** 397–403.

Turnbull, B. W. (1974). Nonparametric estimation of a survivorship function with doubly censored data. *Journal of the American Statistical Association* **69,** 169–173.

Turnbull, B. W. (1976). The empirical distribution with arbitrarily grouped, censored and truncated data. *Journal of the Royal Statistical Society, Series B* **38,** 290–295.

Received January 2002. *Revised November* 2003.
Accepted November 2003.

Appendix

Proof of Formula (7)

We now argue that the policy of infusing for t_j units of time is a special case of a dynamic treatment regime as defined by Murphy, van der Laan, and Robins (2001) and that the key result given by (4) follows directly from the general theory developed for estimating the mean response for such dynamic treatment regimes. Using notation consistent with Murphy et al. (2001), denote the level of treatment at time t as a_t. For our problem a_t is the indicator of treatment continuation at time t; that is, $a_t = 1$ if treatment infusion is continuing at time t and $a_t = 0$ if the infusion was stopped prior to time t for $t \le t_k$. The history of treatment up to and including time t is denoted by $\bar{a}_t = \{a_u, 0 \le u \le t\}$. The treatment assignment at time t is stochastic and is denoted by the random variable A_t and the history of treatment assignment up to and including time t is denoted by the stochastic process $\bar{A}_t = \{A_u,\, 0 \le u \le t\}$. We also define the history of treatment assignment up to and not including time t by $\bar{A}_{t^-} = \{A_u, 0 \le u < t\}$. Other variables available at time t are denoted by L_t and the history up to and including time t by $\bar{L}_t = \{L_u, 0 \le u \le t\}$. For our problem $L_t = \{I(C > t),\ Z(t)I(C > t)\}$; hence, $\bar{L}_t$ denotes whether a treatment-terminating event has occurred prior to time t or not, the time C of the treatment-terminating event if it occurred prior to time t, and the covariate history through the minimum of t and C.

A dynamic treatment regime is a rule that dictates the level of treatment at time t as a function of $\bar{L}_t$ and in Murphy et al. (2001) is denoted by the rule $\bar{d}$ which assigns treatment $d_t(\bar{L}_t)$ at time t for $t \le t_k$. In our problem, the treatment-duration policy "infuse for t_j units of time or until a treatment-terminating event" is an example of a dynamic treatment regime which we will denote by $\bar{d}^j$, where, in our notation,

$$d_t^j(\bar{L}_t) = I(t_j > t, C > t), \quad t \le t_k.$$

Also, in our problem, treatment is terminated immediately upon the occurrence of a treatment-terminating event; otherwise, treatment is terminated by physician discretion at one of the finite set of times $t_1, \ldots, t_k$. Moreover, once treatment is terminated, it will not be continued at some later time. Therefore, in terms of our notation, when $t \ne \{t_1, \ldots, t_k\}$

$$P(A_t = 0 \,|\, \bar{L}_t, \bar{A}_{t^-}) = I(U < t) + I(U = t, \Delta = 0). \quad \text{(A.1)}$$

Equation (A.1) reflects the fact that, at times $t \ne \{t_1, \ldots, t_k\}$, treatment decisions, as a function of past covariate-treatment history, are deterministic. Whereas, for t_j, $j = 1, \ldots, k$

$$P\big(A_{t_j} = 0 \,|\, \bar{L}_{t_j}, \bar{A}_{t_j^-}\big) = \lambda_j(\bar{Z}_j) I(U \ge t_j), \quad \text{(A.2)}$$

where $\lambda_j(\bar{Z}_j)$ is defined by (3).

Under assumption (3) of no unmeasured confounders, we use equation (4.5) of Murphy et al. (2001) to deduce that

$$\mathrm{E}\left((Y - \mu_j) \left[\prod_{t \le t_k} \frac{I\big\{A_t = d_t^j\big(\bar{L}_{t^-}\big)\big\}}{\pi_t(A_t \,|\, \bar{L}_t, \bar{A}_{t^-})} \right] \right) = 0, \quad \text{(A.3)}$$

where μ_j denotes the mean response for the dynamic treatment regime $\bar{d}^j$ and

$$\pi_t(a_t \,|\, \bar{l}_t, \bar{a}_{t^-}) = P(A_t = a_t \,|\, \bar{L}_t = \bar{l}_t, \bar{A}_{t^-} = \bar{a}_{t^-}).$$

By definition of the dynamic treatment regime $\bar{d}^j$,

$$\prod_{t \le t_k} I\big\{A_t = d_t^j\big(\bar{L}_{t^-}\big)\big\} = I(U = t_j, \Delta = 1) + I(U < t_j, \Delta = 0).$$

Because of (A.1) and (A.2), when ($U = t_j$, $\Delta = 1$)

$$\pi_t(A_t \,|\, \bar{L}_t, \bar{A}_{t^-}) = 1 \quad \text{if } t \ne \{t_1, \ldots, t_j\},$$

$$\begin{aligned} \pi_{t_m}\big(A_{t_m} \,\big|\, \bar{L}_{t_m}, \bar{A}_{t_m^-}\big) &= P\big(A_{t_m} = 1 \,\big|\, \bar{L}_{t_m}, \bar{A}_{t_m^-}\big) \\ &= 1 - \lambda_m(\bar{Z}_m) \quad \text{if } t_m = \{t_1, \ldots, t_{j-1}\}, \end{aligned}$$

$$\pi_{t_j}\left(A_{t_j} \,\middle|\, \bar{L}_{t_j}, \bar{A}_{t_j^-}\right) = P\left(A_{t_j} = 0 \,\middle|\, \bar{L}_{t_j}, \bar{A}_{t_j^-}\right) = \lambda_j(\bar{Z}_j),$$

and when $(U < t_j,\, \Delta = 0)$

$$\begin{aligned}
\pi_t(A_t \,|\, \bar{L}_t, \bar{A}_{t^-}) &= 1 \quad \text{if } t \neq \left\{t_1, \ldots, t_{[U]}\right\}, \\
\pi_{t_m}\left(A_{t_m} \,\middle|\, \bar{L}_{t_m}, \bar{A}_{t_m^-}\right) &= P\left(A_{t_m} = 1 \,\middle|\, \bar{L}_{t_m}, \bar{A}_{t_m^-}\right) \\
&= 1 - \lambda_m(\bar{Z}_m) \quad \text{if } t_m = \left\{t_1, \ldots, t_{[U]}\right\}.
\end{aligned}$$

The results above lead us to the conclusion that the statistic in (A.3)

$$\begin{aligned}
&(Y - \mu_j)\left[\prod_{t \le t_k} \frac{I\left\{A_t = d_t^j\left(\bar{L}_{t^-}\right)\right\}}{\pi_t(A_t \,|\, \bar{L}_t, \bar{A}_{t^-})}\right] \\
&\quad = (Y - \mu_j)\left\{\frac{I(U = t_j, \Delta = 1)}{f_j(\bar{Z}_j)} + \frac{I(U < t_j, \Delta = 0)}{K_{[U]}\left(\bar{Z}_{[U]}\right)}\right\},
\end{aligned}$$

which gives the desired result in (4).