Semiparametric inference in observational duration-response studies, with duration possibly right-censored

By BRENT A. JOHNSON

Department of Biostatistics, University of North Carolina,

Chapel Hill, North Carolina 27599, U.S.A.

brentj@email.unc.edu

and ANASTASIOS A. TSIATIS

Department of Statistics, North Carolina State University,

Raleigh, North Carolina 27695, U.S.A.

tsiatis@stat.ncsu.edu

Summary

Once treatment is found to be effective in clinical studies, attention often focuses on optimum or efficacious treatment delivery. In treatment duration-response studies, the optimum treatment delivery refers to the treatment length that optimises the mean response. In many studies, the treatment length is often left to the discretion of an attending investigator or physician but may be abruptly terminated because of treatment-terminating events. Thus, a recommended treatment length often delineates a 'treatment duration policy' which prescribes that treatment be given for a specified length of time or until a treatment-terminating event occurs, whichever comes first.

Estimating a functional relationship between the response and a treatment duration policy, continuously in time, is the focus of this paper.

Some key words: Confounding; Infusion trial; Missing data; Survival analysis.

1. Introduction

Once one treatment is found to be superior in a randomised trial of competing treatments, attention often focuses on determining the optimum treatment dose or length that minimises or maximises study endpoint. While subjects are randomly assigned to the treatment or placebo arm, the actual dose or amount of treatment a subject receives is not randomised. In studies where treatment is given continuously over time, several factors often determine a subject's actual treatment level. First, a patient may experience an adverse event which prematurely terminates the treatment process. Secondly, the decision to stop or continue treatment, when an adverse event has not already occurred, is left to the patient's physician. The ESPRIT, Enhanced Suppression of the Platelet IIb/IIIa Receptor with Integrilin Therapy, infusion study, which motivated this research and discussed in § 5, is one such study where this design was implemented. We note that the realisation of such a composite event may also be one of the adverse events that would necessitate treatment termination. Patients were initially randomised to receive either the experimental treatment regimen or placebo regimen. Then, once it was determined that the treatment was indeed effective, investigators focused their interest on an optimal infusion length. This investigation is difficult for the two reasons stated above, namely, that patients who experienced an infusion-terminating event did not complete the treatment process

and patients were not randomised to infusion lengths.

Define a treatment duration policy for t units of time as a recommendation to treat for t units of time or until a treatment-terminating event occurs, whichever comes first. In order to conceptualise properly a duration-response relationship in our particular setting, we introduce the ideas of potential random variables (Rubin, 1974). We define the potential random variable C as the time at which a randomly selected individual from a population would have a treatment-terminating event, if continuously treated, and by $Y_C^\star$ the response of the individual if this occurs. We also define the individual's potential response $Y_t^\star$ as the response if treatment were terminated at time t, for $t < C$. 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 $Y_{t\wedge C}^\star$, where $t \wedge C$ denotes the minimum of t and C. This may also be written as

$$Y_{t\wedge C}^\star = Y_t^\star I(C > t) + Y_C^\star I(C \leq t). \tag{1}$$

Johnson & Tsiatis (2004) have shown how to estimate consistently the population mean response for this treatment duration policy, $E(Y_{t\wedge C}^\star)$, when treatment duration can take on only a finite number of values, $t_1, \ldots, t_K$. In truth, treatment duration in such studies is a continuous random variable. Hence, in order to apply the methods of Johnson & Tsiatis (2004), the data had to be discretised in an ad hoc fashion. Here we make the more realistic assumption that treatment duration is a continuous random variable and consider the duration-response relationship as a continuous function in time

through a finite number of parameters,

$$E(Y^{\star}_{t\wedge C}) = m(t, \beta), \tag{2}$$

for $\beta = (\beta_1, \ldots, \beta_p)^{\mathrm{T}}$. Keep in mind that our model $m(t, \beta)$ is only a function of t but is indexing a treatment duration policy $t \wedge C$, that is, the policy which treats for t units of time or until a treatment-terminating event occurs, whichever comes first.

2. Methods and Assumptions

The collection of random variables $\{C, Y^{\star}_{C}, Y^{\star}_{t}, t < C\}$ defined above are referred to as potential random variables, or counterfactuals, because, contrary to fact, they may not actually be observed. In contrast, for a randomly selected individual from our population, the observable random variables are given by $\{Y, U, \Delta\}$, where Y denotes the observed response, U denotes the actual continuously-varying, treatment duration, and Δ is an indicator variable such that $\Delta = 1$ when treatment duration was stopped by choice and $\Delta = 0$ when treatment duration was stopped because of treatment-terminating events. In our motivating example, the response Y is a binary indicator, but the methods proposed are applicable for continuous as well as discrete response variables.

We assume that the observed response Y may be written as the following function of potential outcomes $Y^{\star}_{t}$, for $t \leq C$:

$$Y = Y^{\star}_{U} = \Delta Y^{\star}_{U} + (1 - \Delta) Y^{\star}_{C}. \tag{3}$$

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 $Z^H(t)$ to be the history of covariate information up to and including time t, that is

$$Z^H(t) = \{Z(u), u \leq t\}.$$

Define the observable cause-specific hazard function to be

$$\lambda\{t, Z^H(t)\} = \lim_{h \to 0} h^{-1} \text{pr}\{t \leq U < t + h, \Delta = 1 | U \geq t, Z^H(t)\}. \qquad (4)$$

A key assumption that allows us to estimate consistently the parameter β from a random sample of observed data is given by

$$\lambda\{t, Z^H(t)\} =$$
$$\lim_{h \to 0} h^{-1} \text{pr}\left\{t \leq U < t + h, \Delta = 1 | U \geq t, Z^H(t), CI(U \geq t), Y_x^\star, t \leq x \leq C\right\} \qquad (5)$$

In words, assumption (5) implies that, given that a patient has continuously received treatment up to and including time t without a treatment-censoring event, and given the patient's covariate history up to and including time t, the decision to terminate or continue a patient's treatment at time t does not depend on future prognosis. Such an assumption is plausible if information about an individual up until time t, which may be prognostic and which an investigator may use to make decisions on treatment duration, is captured in the data $Z^H(t)$. In the epidemiological literature, assumption (5) is referred to as 'no unmeasured confounder.'

3. The Estimator and Its Properties

3·1. *Motivation*

We now address how to estimate the p-dimensional parameter vector β, where $m(t, \beta)$ is the hypothesised mean model in (2), from a sample of observed data, $\{Y_i, U_i, \Delta_i, Z_i^H(U_i)\}, i = 1, \ldots, n$. As a result of the censoring of

treatment duration when a patient experiences an adverse event and because of the potential confounding, it is not obvious how one can derive an estimator for β directly from the data in an observational study. In contrast, if we had designed an experiment where treatment duration policy was assigned to patients at random, then unbiased estimating equations and estimators for β can be derived easily. This hypothetical randomised experiment is an example of what Murphy et al. (2001) refer to as a random dynamic treatment regime. Murphy et al. (2001) derived the Radon-Nikodym derivative of the distribution of the data from this hypothetical randomised study with respect to the distribution of the data from the observational study under assumption (5) of no unmeasured confounder. Therefore, the strategy is to derive unbiased estimating equations for β under the hypothetical randomised study, which can be done easily, and then to weight these estimating equations by the Radon-Nikodym derivative to obtain unbiased estimating equations for the observational study. We now give the details of this approach.

Let the random variable T denote the assigned treatment duration policy in our hypothetical study. In our ideal experiment, we would randomise individuals to treatment duration policy $T = t$ according to some probability density $h(t)$. The density $h(t)$ is that of a continuous random variable which is known by design and has support on the interval $[\tau_l, \tau_u]$ chosen to correspond to a region of treatment duration actually used in the observational study. If an individual was randomised to policy $T = t$, then this individual would receive treatment for t units of time or until a treatment-terminating event, whichever came first. The data from such an experiment can be summarised as $\{Y_i, T_i, \Delta_i, (1 - \Delta_i)C_i\}, i = 1, \ldots n$, where, for the ith individual in our

sample, Y_i denotes the observed response, T_i the treatment duration policy, Δ_i is the indicator variable for which, as before, $\Delta_i = 0$ corresponds to the case where treatment is stopped prematurely because of a treatment-terminating event and C_i is the time to the treatment-terminating event, observed only if $\Delta_i = 0$.

For such a design, the parameter β can be estimated using generalised estimating equations defined by Liang & Zeger (1986) as

$$\sum_{i=1}^{n} \psi^*(Y_i, T_i, \beta) = 0, \tag{6}$$

where $\psi^*(y, t, \beta) = w(t, \beta)\{y - m(t, \beta)\}$, and $w(t, \beta)$ is a p-dimensional vector of functions of time t and the parameters β. The optimal generalised estimating equation is given by

$$\psi^\dagger(y, t, \beta) = m_\beta(t, \beta)V^{-1}(t)\{y - m(t, \beta)\},$$

where $m_\beta(t, \beta)$ is the p-dimensional gradient of $m(t, \beta)$ with respect to β and $V(t) = \mathrm{var}(Y|T = t)$.

Unlike in the hypothetical randomised study described above, in the observational study we can never observe T_i if $\Delta_i = 0$. Consequently, we now consider how we would estimate the parameter β from the idealised randomised study if we only observed data $(Y_i, U_i, \Delta_i), i = 1, \ldots, n$, where $U_i = \min(T_i, C_i)$.

We first note that the estimating function, divided by n, i.e. $n^{-1}\sum_{i=1}^{n} \psi^*(Y_i, T_i, \beta)$, given in (6), is an unbiased estimator for $E_R\{\psi^*(Y, T, \beta)\}$, where we use $E_R(\cdot)$ to denote expectation under this hypothetical randomised study. The function $\psi^*(Y, T, \beta)$ is referred to as an unbiased estimating function as its

expectation, under the truth, $\beta = \beta_0$, is equal to zero and can be used to derive a consistent and asymptotically normal estimator of β by solving the estimating equation (6). The expectation $E_R\{\psi^*(Y,T,\beta)\}$ can be written as

$$E_R\{\Delta\psi^*(Y,T,\beta)+(1-\Delta)\psi^*(Y,T,\beta)\} = E_R\{\Delta\psi^*(Y,U,\beta)+(1-\Delta)\psi^*(Y_C^\star,T,\beta)\}. \tag{7}$$

As a consequence of randomisation, T is independent of the potential outcomes C and $Y_C^\star$. Consequently, the expectation of the second term of the sum in the right-hand side of equation (7) can be derived as

$$\begin{aligned}
&E_R\{(1-\Delta)\psi^*(Y_C^\star,T,\beta)\} \\
&= E_R[E_R\{(1-\Delta)\psi^*(Y_C^\star,T,\beta)|\Delta,C,Y_C^\star\}] \\
&= E_R[(1-\Delta)E_R\{\psi^*(Y_C^\star,T,\beta)|T>C,C,Y_C^\star\}] \\
&= E_R\left\{\frac{(1-\Delta)}{H(C)}\int_C^\infty \psi^*(Y_C^\star,u,\beta)h(u)du\right\} \\
&= E_R\left\{\frac{(1-\Delta)}{H(U)}\int_U^\infty \psi^*(Y,u,\beta)h(u)du\right\},
\end{aligned} \tag{8}$$

where $H(x) = \int_x^\infty h(u)du$ is the survival function $\text{pr}(T \geq x)$. For notational convenience, we take the integral from U to ∞ in equation (8). Keep in mind, however, that the density $h(u)$ only has support from $[\tau_l, \tau_u]$ and hence the integral is actually restricted from $\min(U, \tau_l)$ to τ_u.

Replacing the second term of the sum in the right-hand side of (7) by the expression in (8), we deduce that the expectation of the estimating function $\psi^*(Y,T,\beta)$ is equal to the expectation of

$$\Delta\psi^*(Y,U,\beta) + \frac{(1-\Delta)}{H(U)}\int_U^\infty \psi^*(Y,u,\beta)h(u)du. \tag{9}$$

Consequently, (9) is an unbiased estimating function which can be used to derive a consistent, asymptotically normal estimator of β from the ideal

randomised study, using the data $(Y_i, U_i, \Delta_i), i = 1, \ldots, n$, by solving the estimating equation

$$\sum_{i=1}^{n} \left\{ \Delta_i \psi^*(Y_i, U_i, \beta) + \frac{(1-\Delta_i)}{H(U_i)} \int_{U_i}^{\infty} \psi^*(Y_i, u, \beta) h(u) du \right\} = 0. \qquad (10)$$

The randomised experiment described above is an example of a random dynamic treatment regime as defined by Murphy et al. (2001). Of course, in the observational study, individuals were not randomised by design. However, because of the assumption of no unmeasured confounder, given by (5), we can use the theory developed by Murphy et al. (2001), which provides a Radon-Nikodym derivative, to connect the probability distribution under the idealised randomised study, denoted by $\mathrm{pr}_R(\cdot)$, to the probability distributions for the actual observational study, denoted by $\mathrm{pr}(\cdot)$. Under some regularity conditions, including that

$$\mathrm{pr}[\lambda\{t, Z^H(t)\} > 0 \text{ for all } t \text{ such that } h(t) > 0] = 1, \qquad (11)$$

Lemma 4.1 of Murphy et al. (2001) can be used to deduce that the distribution of (Y, U, Δ) under $\mathrm{pr}_R(\cdot)$ is absolutely continuous with respect to the distribution of (Y, U, Δ) under $\mathrm{pr}(\cdot)$, and a version of the Radon-Nikodym derivative is

$$E\left[\Delta \frac{h(U)}{f\{U, Z^H(U)\}} + (1-\Delta) \frac{H(U)}{K\{U, Z^H(U)\}} | Y = y, U = u, \Delta = \delta \right], \qquad (12)$$

where $K\{t, Z^H(t)\} = \exp\left[-\Lambda\{t, Z^H(t)\}\right]$, $\Lambda\{t, Z^H(t)\} = \int_0^t \lambda\{u, Z^H(u)\} du$ and $f\{t, Z^H(t)\} = \lambda\{t, Z^H(t)\} K\{t, Z^H(t)\}$. In the Appendix we show how the Radon-Nikodym derivative given by equation (12) falls out from the Murphy et al. (2001) theory.

Since (12) is the Radon-Nikodym derivative, this implies that

$$
\begin{aligned}
& E_R\left\{\Delta\psi^*(Y,U,\beta)+\frac{(1-\Delta)}{H(U)}\int_U^\infty \psi^*(Y,u,\beta)h(u)du\right\} \\
& = E\left(\left[\Delta\frac{h(U)}{f\{U,Z^H(U)\}}+(1-\Delta)\frac{H(U)}{K\{U,Z^H(U)\}}\right]\right. \\
& \left.\times\left\{\Delta\psi^*(Y,U,\beta)+\frac{(1-\Delta)}{H(U)}\int_U^\infty \psi^*(Y,u,\beta)h(u)du\right\}\right) \\
& = E\left[\Delta\psi^*(Y,U,\beta)\frac{h(U)}{f\{U,Z^H(U)\}}\right. \\
& \left.+\frac{(1-\Delta)}{K\{U,Z^H(U)\}}\int_U^\infty \psi^*(Y,u,\beta)h(u)du\right]. \qquad (13)
\end{aligned}
$$

Letting D denote $\{U,\Delta,Z^H(U)\}$, we define

$$
\begin{aligned}
\psi(Y,D,\beta) &= \Delta\psi^*(Y,U,\beta)\frac{h(U)}{f\{U,Z^H(U)\}} \\
&+ \frac{(1-\Delta)}{K\{U,Z^H(U)\}}\int_U^\infty \psi^*(Y,u,\beta)h(u)du. \qquad (14)
\end{aligned}
$$

By the series of arguments given above we have shown that $E\{\psi(Y,D,\beta)\}$ is equal to the expectation of the unbiased, under $\mathrm{pr}_R(\cdot)$, estimating function given in equation (6), namely $\psi^*(Y,T,\beta)$. Consequently, if the hazard function $\lambda\{t,Z^H(t)\}$ were known to us, then $\psi(Y,D,\beta)$ is an unbiased, under $\mathrm{pr}(\cdot)$, estimating function. Hence, a consistent asymptotically normal estimator for β could be derived by solving the estimating equation

$$\sum_{i=1}^n \psi(Y_i,D_i,\beta)=0. \qquad (15)$$

Since the estimator above depends on $\lambda\{t,Z^H(t)\}$, which is unknown, it must be estimated. A popular and flexible model is the proportional hazards model of Cox (1972), in which

$$\lambda\{t,Z^H(t)\}=\lambda_0(t)\exp\{\gamma^{\mathrm{T}}Z^H(t)\}. \qquad (16)$$

The estimating function $\psi(Y, D, \beta)$ will have mean zero for arbitrary $h(t)$ as long as assumption (11) is satisfied, although we expect the choice of $h(t)$ to have an impact on the efficiency of the resulting estimator. In fact, since estimating equations can be defined up to a proportionality constant, the function $h(t)$ can be any positive function with finite integral $\int_{\tau_l}^{\tau_u} h(u)du$. In conjunction with (16), consider the convenient choice for $h(t)$, i.e. $h(t) = \lambda_0(t)$. Using the definition of $\lambda\{t, Z^H(t)\}$ defined in (16), $H(t) = \Lambda_0(t)$, and letting $\mathcal{T}$ denote the interval $[\tau_l, \tau_u]$, we define our proposed estimating function as

$$\begin{aligned}\psi(Y, D, \beta, \gamma) = & \\ & \frac{\Delta\psi^*(Y, U, \beta)I(U \in \mathcal{T})}{g\{U, Z^H(U), \gamma\}} \\ & + \frac{(1-\Delta)I(U \in \mathcal{T})}{K\{U, Z^H(U), \gamma\}} \int_{\mathcal{T}\cap\{U<t\}} \psi^*(Y, t, \beta)d\Lambda_0(t), \end{aligned} \tag{17}$$

where $g\{t, Z^H(t), \gamma\} = \exp\{\gamma^{\mathrm{T}} Z^H(t)\}K\{t, Z^H(t)\}$. For notational convenience, define

$$\Upsilon(\beta) = \int_{\mathcal{T}\cap\{U<s\}} \psi^*(Y, s, \beta)d\Lambda_0(s),$$

and $\Upsilon_i(\beta)$ as $\Upsilon(\beta)$ with (Y, U) replaced with (Y_i, U_i). Our parameterisation of $\lambda\{t, Z^H(t)\}$ and our choice for $h(t)$ allows us to derive uniformly consistent estimators for $g\{t, Z^H(t)\}$, $K\{t, Z^H(t)\}$, and $\Upsilon(\beta)$ using standard counting process methodology (Andersen et al., 1993, VII.2).

3·2. *Asymptotic properties*

Our proposed estimator is the solution $\hat{\beta}_n$ to

$$\sum_{i=1}^{n} \psi(Y_i, D_i, \beta, \gamma) = 0,$$

with $g\{t, Z^H, \gamma_0\}$, $K\{t, Z^H(t), \gamma_0\}$ and $\Upsilon_i(\beta, \gamma_0)$ replaced by their estimated quantities, $g\{t, Z^H(t), \hat{\gamma}_n\}$, $K\{t, Z^H(t), \hat{\gamma}_n\}$ and $\Upsilon_i(\beta, \hat{\gamma}_n)$, respectively, where

$$\begin{aligned} \Upsilon_i(\beta, \gamma) &= \int_{\mathcal{T}\cap\{U<s\}} \psi^*(Y_i, s, \beta) d\hat{\Lambda}_0(s, \gamma) \\ &= \int_{\mathcal{T}\cap\{U<s\}} \frac{\psi^*(Y, s, \beta) dN(s)}{\sum_{k=1}^n \exp\{\gamma^t Z_k^H(s)\} R_k(s)}, \end{aligned} \tag{18}$$

$N_i(t) = I(U_i \le t, \Delta_i = 1)$, $N(t) = \sum_{i=1}^n N_i(t)$ and $R_i(t) = I(U_i \ge t)$.

Here, we also show that our estimator $\hat{\beta}_n$ is consistent and show that $n^{1/2}(\hat{\beta}_n - \beta_0)$ is asymptotically normal under the proposed conditions. Define

$$S^{(r)}(t, \gamma) = n^{-1} \sum_{k=1}^n R_k(t) e^{\gamma^{\mathrm{T}} Z_k^H(t)} Z_k^{H\otimes r}(t), \;\; s^{(r)}(t, \gamma) = E\{S^{(r)}(t, \gamma)\},$$

$$\bar{Z}^H(t, \gamma) = \frac{S^{(1)}(t,\gamma)}{S^{(0)}(t,\gamma)}, \;\; \mu(t, \gamma) = \frac{s^{(1)}(t,\gamma)}{s^{(0)}(t,\gamma)},$$

for $r = 0, 1, 2$, where, for a column vector a, $a^{\otimes 0} = 1$, $a^{\otimes 1} = a$ and $a^{\otimes 2} = aa^{\mathrm{T}}$. Let $\hat{\gamma}_n$ be the maximum partial likelihood estimator (Cox, 1972) so that

$$n^{1/2}(\hat{\gamma}_n - \gamma_0) = n^{-1/2} \sum_{i=1}^n \Sigma_\gamma^{-1} \int_0^\infty \left\{ Z_i^H(u) - \mu(u, \gamma_0) \right\} dM_i(u) + o_p(1), \tag{19}$$

where $M_i(t) = N_i(t) - \int_0^t \exp\{\gamma^{\mathrm{T}} Z_i^H(u)\} R_i(u) d\Lambda_0(u)$ and Σ_γ is

$$\int_0^\infty E\left[\left\{ Z^H(u) - \mu(u, \gamma_0) \right\} \left\{ Z^H(u) - \mu(u, \gamma_0) \right\}^{\mathrm{T}} \exp\{\gamma_0^{\mathrm{T}} Z^H(u)\} R(u) \right] d\Lambda_0(u).$$

Let $\theta = (\beta^{\mathrm{T}}, \gamma^{\mathrm{T}})^{\mathrm{T}}$ be the $p+q$ column vector of all the parameters and define $J(\theta)$ as the sum of three $p \times q$ matrices, i.e.

$$J(\theta) = P(\theta) + Q(\theta) + G(\theta),$$

where

$$
\begin{aligned}
P(\theta) &= -\int_0^{\tau_u} E\{a(u,\theta)\}\mu(u,\gamma)d\Lambda_0(u) \\
Q(\theta) &= E\{b(\theta)\} + \int_0^{\tau_u} E\{c(u,\theta)\}\mu(u,\gamma)d\Lambda_0(u) \\
G(\theta) &= -\int_{\mathcal{T}} E\{d(u,\theta)\}\mu(u,\gamma)d\Lambda_0(u) \\
a(u,\theta) &= \psi(Y,D,\beta,\gamma)\exp\{\gamma^{\mathrm{T}}Z^H(u)\}I(U \geq u) \\
b(\theta) &= -\frac{\Delta\psi^*(Y,U,\beta)I(U \in \mathcal{T})Z^H(U)}{g\{U,Z^H(U),\gamma\}} \\
c(u,\theta) &= \psi(Y,D,\beta,\gamma)\exp\{\gamma^{\mathrm{T}}Z^H(u)\}Z^{HT}(u)I(U \geq u) \\
d(u,\theta) &= \frac{(1-\Delta)\psi^*(Y,u,\beta)I(U < u)}{K\{U,Z^H(U),\gamma\}}.
\end{aligned}
$$

Then the ith influence function is given by

$$
\begin{aligned}
&h(Y_i,D_i,\theta_0) = \\
&\quad A^{-1}(\theta_0)\left\{\int_0^{\tau_u}\left[J(\theta_0)\Sigma_\gamma^{-1}\{Z_i^H(u)-\mu(u,\gamma_0)\}+\right.\right. \\
&\quad \left.\left.\frac{E\{a(u,\theta_0)+d(u,\theta_0)\}}{s^{(0)}(u,\gamma_0)}\right]dM_i(u)+\psi(Y_i,D_i,\beta_0,\gamma_0)\right\},
\end{aligned}
$$

where $A(\theta_0)$ is the $p \times p$ matrix of $E\left\{(-\partial/\partial\beta)\,\psi(Y_i,D_i,\beta,\gamma)\right\}$, with β and γ evaluated at their true values. A consistent estimator for the asymptotic variance of $n^{1/2}(\hat{\beta}_n - \beta_0)$ is given by

$$
n^{-1}\sum_{i=1}^{n}\hat{h}(Y_i,D_i,\hat{\theta}_n)\hat{h}^{\mathrm{T}}(Y_i,D_i,\hat{\theta}_n),
$$

where all expectations are estimated by their respective sample mean estimators and $h(Y_i,D_i,\hat{\theta}_n)$ is estimated using methods for estimating martingale residuals (Fleming & Harrington, 1991, §4.5). An outline of the derivation of the asymptotic variance is given in the appendix.

4. Numerical Studies

We conducted several numerical studies to evaluate the small sample properties of our estimator, some of which we show here. Our goal is to simulate data for a Bernoulli response using a linear logistic model, i.e.

$$m(t, \beta) = F\{\beta_0 + \beta_1(t - t_\mu)\},$$

where $F(t) = (1 + e^{-t})^{-1}$ and t_μ is a centring constant.

To simulate the data for this model, we begin by independently simulating the treatment censoring random variable C and a potential response threshold W following distributions described below. Then we propose that $Y_t^\star$ follow the rule

$$Y_t^\star = \begin{cases} 1 & \text{if } t < \min(C, W) \\ 0 & \text{if } W \leq t < C, \end{cases}$$

and $Y_C^* = 0$. Now, we want

$$\text{pr}(Y_{t \wedge C}^\star = 1) = \text{pr}\{\min(C, W) > t\} = \text{pr}(C > t) P(W > t)$$

to follow the linear logistic model $m(t, \beta)$. If we take $\text{pr}(W > t) = \{\text{pr}(C > t)\}^\theta$, then some simple algebra shows that

$$\text{pr}(C > t) = \left[F\{\beta_0 + \beta_1(t - t_\mu)\}\right]^{1/(\theta+1)}.$$

Such a distribution is easily generated. A similar strategy is employed for the random variable W.

Continuing, we simulate our confounding random variable Z as normal with conditional mean $\alpha_0 + \alpha_1 \min(C, W)$ and conditional variance σ_z^2. Next, we simulate the potential treatment duration T as Weibull with shape r and hazard

$$\lambda(t; Z) = r t^{r-1} \rho^r \exp(\gamma Z).$$

Finally, if $T < C$, then $U = T$, $\Delta = 1$ and $Y = Y_T^\star$. If $T \geq C$, then $U = C$, $\Delta = 0$ and $Y = Y_C^\star = 0$. We repeat this process for n subjects.

While some parameters are chosen to reflect aspects of the dataset in our example, others are chosen for convenience and simplicity of exposition. In the simulation scenarios to follow, $r = 5$, $\rho = 0.04$, and $\gamma = -0.5$ were chosen to yield a density for T with mode around 25 hours and an inverse relationship between the variables Z and T, i.e. large Z implies small $\lambda(t; Z)$. For the distribution of Z, we chose the parameters $\alpha_0 = 0$ and $\sigma_z = 0.25$, but allowed α_1 to vary in our simulation scenarios to reflect different degrees of confounding. Finally, we chose $\theta = 3$, $t_\mu = 25$, $\beta_0 = 0$ and $\beta_1 = -1$ in our linear logistic model.

[Table 1 about here]

[Table 2 about here]

In Table 1, we see that the bias is minimal and that the Monte Carlo standard errors match well with the standard error estimates. In Table 2, we also see that our estimator covers the true value at the nominal level. We compare our estimators to two maximum likelihood estimators, $\tilde{\beta}$ and β^*. In our simulations, $\tilde{\beta}$ is the usual estimator from a logistic regression model that relates the binary response Y to the observed treatment duration U on all subjects while β^* only includes those observations from the uncensored subjects. Here, we only use data up to time $\tau_u = 27$ for all estimators. The results show that $\tilde{\beta}$ performs poorly on all accounts while the uncensored estimator β^* performs poorly some of the time.

5. Example

We now apply our methods to data from the ESPRIT infusion trial from Duke University Medical Center. The main objective of ESPRIT was to compare eptifibatide, i.e. Integrilin, therapy to placebo on the basis of the composite binary endpoint Y of death, myocardial infarction or urgent target vessel revascularisation within 30 days treatment initiation. The random variable U denotes the observed infusion length and Δ is an indicator taking the value one when the attending physician stopped the infusion process and equal to zero when the infusion was stopped by an infusion-terminating event, such as abrupt closure, no reflow, or coronary thrombosis. We also included the following potential confounders in our analysis: diabetes (0/1), percutaneous transluminal coronary angioplasty (PTCA,0/1), angina (0/1), heparin (0/1) and weight, in kilograms, which is consistent with earlier findings.

[Table 3 about here]

We also include the two naive estimators presented in § 4, $\tilde{\beta}$ and β^*, obtained from logistic regression using all the data and using only the uncensored data, respectively. We only present the results for a linear model in Table 3. As we can see, the linear term is not significant which leads us to believe that patients will have about the same probability of the endpoint regardless of the length of infusion. These results confirm earlier suggestions that longer infusion lengths do not improve prognosis.

[Figure 1 about here]

Figure 1 shows that there appears to be some increase in the expected probability of endpoint as infusion length increases, this relationship is nei-

ther strong nor significant. We overfitted our data with higher-order polynomial models simply to show that our method can identify nonlinear relationships when they exist and can potentially match our earlier results closely. The crosses in Figure 1 denote the policy estimates found in § 4 of Johnson & Tsiatis (2004) at infusion lengths equal to 16, 18, 20, 22, and 24 hours.

Acknowledgement

We would like to acknowledge and thank an anonymous reviewer for the innovative suggestion to cast this problem using a hypothetical randomised study together with the Radon-Nikodym derivative to motivate our unbiased estimating equation, thereby making the entire exposition simpler and more elegant. This work was supported in part by grants from the U. S. National Science Foundation, National Institutes of Health, and National Institutes for the Environmental Health Sciences.

Appendix

Technical details

Deriving the Radon-Nikodym derivative. In our problem, when treatment is discontinued for any patient, it is never reinitiated. Consequently, treatment assignment can be characterised by the counting process $A(t) = I(U \leq t, \Delta = 1)$, where we let $A(t) = 0$ if a subject is still treated at time t or a treatment-terminating event has occurred, and $A(t) = 1$ if treatment has been discontinued by time t. This was also referred to as $N(t)$ is § 3.3. The infinitesimal $dA(t) = 1$ denotes whether the treatment was discontinued in

the interval $[t, t + dt)$ and zero otherwise.

If we define a filtration for a single observation as the increasing sigma algebra

$$\mathcal{F}(t) = \sigma\{A(x), I(U \leq x, \Delta = 0),\ x \leq t, Z^H(t)\},$$

then the intensity process of the counting process for the observational data $\tilde{\lambda}(t)$, defined heuristically as $\tilde{\lambda}(t)dt = \text{pr}\{dA(t) = 1|\mathcal{F}(t^-)\}$, is equal to $\lambda\{t, Z^H(t)\}I(U \geq t)$, where $\lambda\{t, Z^H(t)\}$ is defined by (4).

The hypothetical randomised study introduced in § 3 is an example of what Murphy et al. (2001) refer to as a random dynamic treatment regime; that is, where the treatment assignment at time t can be made at random with probability that depends on previous treatment assignments and other variables $S(t)$. In our hypothetical randomised experiment, a patient's treatment has to be terminated with certainty if they experience a treatment-terminating event. Therefore, at time t, a patient will receive treatment only if he/she has received continuous treatment up until time t without a treatment-terminating event. The intensity process is given by $\tilde{\lambda}_R(t) = \lambda_R(t)I(U \geq t)$, where $\lambda_R(t) = h(t)/H(t)$ is the hazard function for the random variable T, the assigned treatment duration policy, in the hypothetical randomised study. As a result of randomisation, the intensity is not affected by the introduction of the covariate history $Z^H(t)$ in the filtration for the hypothetical randomised study.

Under assumption (5) of no unmeasured confounders, also referred to as the sequential randomisation assumption, and assumption (11), Murphy et al. (2001) showed that the Radon-Nikodym derivative of the distribution of (Y, U, Δ) for the hypothetical randomised study, i.e. the random dynamic

treatment regime, with respect to the distribution of (Y, U, Δ) for the observational study is given by

$$E\left[\prod_{0\le t\le\tau_u}\left\{\frac{\tilde{\lambda}_R(t)}{\tilde{\lambda}(t)}\right\}^{dA(t)}\frac{\exp\{-\tilde{\Lambda}_R(\tau_u)\}}{\exp\{-\tilde{\Lambda}(\tau_u)\}}\Big|Y=y,U=u,\Delta=\delta\right], \qquad \text{(A1)}$$

where $\tilde{\Lambda}(\tau_u)=\int_0^{\tau_u}\tilde{\lambda}(x)dx$.

We note that (A1) is the conditional expectation, with respect to (Y, U, Δ), of the contribution of a single observation to the partial likelihood ratio for the distribution of the counting process $A(t)$ for the randomised study with respect to the distribution of $A(t)$ for the observational study as defined in § II of Andersen et al. (1993). Strictly speaking, Murphy et al. (2001) assumed that treatment decisions can be made only at discrete times and the partial likelihood ratio derived in their paper is that for a discrete time process. This is generalised to the continuous time process using the product limit representation as given by Jacod's formulas (Andersen et al., 1993, Corollary II.7.3).

Since

$$\left[\prod_{0\le t\le\tau_u}\left\{\frac{\tilde{\lambda}_R(t)}{\tilde{\lambda}(t)}\right\}^{dA(t)}\right]=\Delta\left[\frac{h(U)/H(U)}{\lambda\{U,Z^H(U)\}}\right],$$

$\exp\{-\tilde{\Lambda}_R(\tau_u)\}=H(U)$ and $\exp\{-\tilde{\Lambda}(\tau_u)\}=K\{U,Z^H(U)\}$, we obtain that the expression for the partial likelihood ratio in (A1) is equal to

$$\begin{aligned}&\frac{\Delta\lambda_R(U)H(U)+(1-\Delta)H(U)}{\Delta\lambda\{U,Z^H(U)\}K\{U,Z^H(U)\}+(1-\Delta)K\{U,Z^H(U)\}}\\ =\;&\Delta\frac{h(U)}{f\{U,Z^H(U)\}}+(1-\Delta)\frac{H(U)}{K\{U,Z^H(U)\}},\end{aligned} \qquad \text{(A2)}$$

and the conditional expectation of (A2) given $(Y=y,U=u,\Delta=\delta)$ is the Radon-Nikodym derivative given in equation (12).

Asymptotic properties of $\hat{\beta}_n$. For simplicity, we outline the asymptotic variance calculations of our estimator when $Z^H(U) = Z$; that is, the covariate history includes time-independent confounders only. The calculations of the asymptotic variance for the more general case are similar but notationally more cumbersome. The details for both cases are given in B. A. Johnson's unpublished 2003 Ph. D. Thesis from the North Carolina State University.

Recall that our estimator satisfies the system of p estimating equations

$$\sum_{i=1}^{n} \psi(Y_i, D_i, \hat{\beta}_n, \hat{\gamma}_n) = 0,$$

with $\psi(Y_i, D_i, \beta, \gamma)$ defined in (17). Under the assumption that $Z^H(U) = Z$, we have

$$\begin{aligned} g(U_i, Z_i, \gamma_0) &= \exp\{\gamma^{\mathrm{T}} Z_i - \Lambda_0(U_i) \exp(\gamma^{\mathrm{T}} Z_i)\} \\ K(U_i, Z_i, \gamma_0) &= \exp\{\Lambda_0(U_i) \exp(\gamma^{\mathrm{T}} Z_i)\}. \end{aligned}$$

Since we assumed that $(\partial/\partial\beta)\, \psi^*(Y, D, \beta)$ is bounded by a integrable random variable in neighbourhoods of β_0, it is straightforward to show that our estimator minus the estimand may be written as

$$(\hat{\beta}_n - \beta_0) = n^{-1} \sum_{i=1}^{n} A^{-1}(\beta_0) \psi(Y_i, D_i, \beta_0, \hat{\gamma}_n) + O_p(1), \tag{A3}$$

where

$$\begin{aligned} A_n(\beta) &= -n^{-1} \sum_{i=1}^{n} \frac{\partial}{\partial \beta} \psi(Y_i, D_i, \beta, \hat{\gamma}_n) \\ A(\beta) &= \lim_{n \to \infty} A_n(\beta). \end{aligned}$$

Note that $\psi(Y_i, D_i, \beta_0, \hat{\gamma}_n)$ in (A3) has three estimated quantities, namely $\hat{\gamma}_n$, $\hat{\Lambda}_0(U_i)$ and $\hat{\Upsilon}_i(\beta, \hat{\gamma}_n)$, defined in (18). We proceed by adding and subtracting

$\psi(\cdot)$ with each of these estimated quantities replaced by its true value. We define the following shorthand for the calculations below: $\hat{\Lambda}_0 = \hat{\Lambda}_0(U_i, \hat{\gamma}_n)$ and $\hat{\Upsilon}_i = \hat{\Upsilon}_i(\beta_0, \hat{\gamma}_n)$. We also use the abbreviated estimating function when $\psi(Y_i, D_i, \beta_0)$ is a function of the true values γ_0, $\Lambda_0(U_i)$ and $\Upsilon_i(\beta_0)$. Now, it is important to express the following three differences as sums of independent and identically distributed random variables plus remainder terms, r_n, where $n^{1/2} r_n = o_p(1)$:

$$\sum_{i=1}^{n} \left\{ \psi(Y_i, D_i, \beta_0, \hat{\gamma}_n, \hat{\Lambda}_0, \hat{\Upsilon}_i) - \psi(Y_i, D_i, \beta_0, \gamma_0, \hat{\Lambda}_0, \hat{\Upsilon}_i) \right\} \tag{A4}$$

$$\sum_{i=1}^{n} \left\{ \psi(Y_i, D_i, \beta_0, \gamma_0, \hat{\Lambda}_0, \hat{\Upsilon}_i) - \psi(Y_i, D_i, \beta_0, \gamma_0, \Lambda_0, \hat{\Upsilon}_i) \right\} \tag{A5}$$

$$\sum_{i=1}^{n} \left\{ \psi(Y_i, D_i, \beta_0, \gamma_0, \Lambda_0, \hat{\Upsilon}_i\} - \psi\{Y_i, D_i, \beta_0, \gamma_0, \Lambda_0, \Upsilon_i) \right\} \tag{A6}$$

In (A4), we are only considering the $\hat{\gamma}_n$ in the definition of $g\{t, Z^H(t), \hat{\gamma}_n\}$ and $K\{t, Z^H(t), \hat{\gamma}_n\}$ that does not come about through the definition of $\hat{\Lambda}_0(\cdot, \hat{\gamma}_n)$ or $\Upsilon_i(\beta, \hat{\gamma}_n)$. Therefore, by a first-order Taylor-series expansion and (19), (A4) may be rewritten as

$$n^{-1} \sum_{i=1}^{n} \int_0^{\infty} \left[Q(\theta_0) \Sigma_{\gamma}^{-1} \{ Z_i - \mu(u, \gamma_0) \} \right] dM_i(u) + r_n,$$

where $\theta = (\beta^{\mathrm{T}}, \gamma^{\mathrm{T}})^{\mathrm{T}}$ and

$$\begin{aligned} b_i(\theta_0) &= \psi(Y_i, D_i, \beta_0, \gamma_0) \left[\left\{ \Lambda_0(U_i) \exp(\gamma_0^{\mathrm{T}} Z_i) - \Delta_i \right\} Z_i^{\mathrm{T}} \right] \\ \bar{b}(\theta_0) &= n^{-1} \sum_{i=1}^{n} b_i(\theta_0) \\ Q(\theta_0) &= E\{b(\theta_0)\}. \end{aligned}$$

Similarly, a Taylor-series expansion of $\hat{\Lambda}_0(U_i, \hat{\gamma}_n)$ around $\Lambda_0(U_i)$ in (A5) yields

$$n^{-1} \sum_{i=1}^{n} \psi(Y_i, D_i, \beta_0) \exp(\gamma_0^{\mathrm{T}} Z_i) \{ \hat{\Lambda}_0(U_i, \hat{\gamma}_n) - \Lambda_0(U_i) \} + r_n. \tag{A7}$$

Using standard counting process methodology (Fleming & Harrington, 1991), the Weak Law of Large Numbers and Lenglart's Inequality (Lenglart, 1977), we show that (A7) equals

$$n^{-1}\sum_{i=1}^{n}\int_0^\infty \left[P(\theta_0)\Sigma_\gamma^{-1}\{Z_i-\mu(u,\gamma_0)\}+\frac{E\{a(u,\theta_0)\}}{s^{(0)}(u,\gamma_0)}\right]dM_j(u)+r_n,$$

where

$$\begin{aligned} a(u,\theta_0) &= \psi(Y,D,\theta_0)\exp(\gamma_0^{\mathrm{T}}Z)I(U>u,U\in\mathcal{T}) \\ P(\theta_0) &= \int_0^L E\{a(u,\theta_0)\}\mu(u,\gamma_0)d\Lambda_0(u). \end{aligned}$$

Finally, (A6) may be rewritten as

$$\sum_{i=1}^{n}\frac{(1-\Delta_i)}{K(U_i,Z_i)}\{\Upsilon_i(\beta_0,\hat\gamma_n)-\Upsilon_i(\beta_0)\}= \sum_{i=1}^{n}\frac{(1-\Delta_i)}{K(U_i,Z_i)}\int_{\mathcal{T}\cap\{U_i<u\}}\psi^*(Y_i,u,\beta_0)\{d\hat\Lambda_0(u,\hat\gamma_n)-d\Lambda_0(u)\}. \quad \text{(A8)}$$

Similarly, (A8) may be written as

$$n^{-1}\sum_{i=1}^{n}\int_0^\infty\left[\frac{E\{d(u,\theta_0)\}}{s^{(0)}(u,\gamma_0)}-G(\theta_0)V^{-1}\{Z_i-\mu(u,\gamma_0)\}\right]dM_i(u)+r_n,$$

where

$$\begin{aligned} d(u,\theta_0) &= \frac{(1-\Delta)\psi^*(Y,t,\beta_0)I(U<u)}{K\{U,Z,\Lambda_0(U),\gamma_0\}} \\ G(\theta_0) &= \int_{\mathcal{T}}E\{d(u,\theta_0)\}\mu(u,\gamma_0)d\Lambda_0(u). \end{aligned}$$

Then

$$n^{1/2}(\hat\beta_n-\beta_0)=n^{-1/2}\sum_{i=1}^{n}h(Y_i,D_i,\theta_0)+o_p(1),$$

where the ith influence function is

$$h(Y_i, D_i, \theta_0) = A^{-1}(\beta_0)\Bigg(\int_0^\infty \Bigg[J(\theta_0)\Sigma_\gamma^{-1}\{Z_i - \mu(u, \gamma_0)\} + \frac{E\{a(u, \theta_0) + d(u, \theta_0)\}}{s^{(0)}(u, \gamma_0)} \Bigg] dM_i(u) + \psi(Y_i, D_i, \beta_0, \gamma_0) \Bigg),$$

where

$$J(\theta_0) = P(\theta_0) + Q(\theta_0) + G(\theta_0).$$

A consistent estimator of the asymptotic variance is given by

$$n^{-1} \sum_{i=1}^{n} \hat{h}(Y_i, D_i, \hat{\theta}_n)\hat{h}^{\mathrm{T}}(Y_i, D_i, \hat{\theta}_n),$$

where all expectations are evaluated by their respective sample averages, which matches precisely the estimator given in § 3 with $Z^H(U)$ replaced by Z.

References

Andersen, P. K., Borgan, O., Gill, R. D. & Keiding, N. (1993). *Statistical Models Based on Counting Processes*. New York: Springer-Verlag.

Cox, D. R. (1972). Regression models and life tables (with Discussion). *J. R. Statist. Soc.* B **34**, 187–220.

Fleming, T. A. & Harrington, D. P. (1991). *Counting Processes and Survival Analyses*. New York: Wiley.

Johnson, B. A. & Tsiatis, A. A. (2004). Estimating mean response as a function of treatment duration in an observational study, where duration may be informatively censored. *Biometrics* **60**, 315–23.

Lenglart, E. (1977). Relation de domination entre deux processus. *Ann. Inst. Henri Poincaré* **13**, 171–9.

Liang, K.-Y. & Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika* **73**, 13–22.

Murphy, S. A., van der Laan, M. J. & Robins, J. M. (2001). Marginal mean models for dynamic regimes. *J. Am. Statist. Assoc.* **96**, 1410–23.

Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* **6**, 688–701.

Table 1. *Numerical study for Bernoulli response using 1000 Monte Carlo datasets.* $\hat{\beta}$ *is our estimator, SSE is the Monte Carlo standard error, and SEE is the average standard error across Monte Carlo datasets. The true parameters are* $\beta_0 = 0$ *and* $\beta_1 = -1$*. Sample size is 250.*

	α_1	$\hat{\beta}$	SSE	SEE
	0	0.00	0.26	0.26
β_0	0.25	-0.01	0.26	0.26
	0.5	0.01	0.27	0.27
	0	-1.04	0.20	0.19
β_1	0.25	-1.03	0.19	0.18
	0.5	-1.04	0.20	0.19

Table 2. *Numerical study for Bernoulli response using 1000 Monte Carlo datasets.* $\hat{\beta}$ *is our estimator,* $\tilde{\beta}$ *is maximum likelihood estimator from logistic regression using all the subjects, while* β^* *is the maximum likelihood estimator using just those uncensored subjects. Empirical coverage probabilities are given in parentheses. The true parameters are* $\beta_0 = 0$ *and* $\beta_1 = -1$*. Sample size is 250.*

	α_1	$\hat{\beta}$	$\tilde{\beta}$	β^*
β_0	0	0.00 (0.95)	-2.26 (0.00)	-1.60 (0.15)
	0.25	-0.01 (0.96)	-2.26 (0.00)	-1.47 (0.21)
	0.5	0.01 (0.96)	-2.35 (0.00)	-1.48 (0.26)
β_1	0	-1.04 (0.94)	-0.85 (0.73)	-1.01 (0.93)
	0.25	-1.03 (0.94)	-0.87 (0.76)	-0.95 (0.90)
	0.5	-1.04 (0.95)	-0.94 (0.89)	-0.97 (0.91)

Table 3. *Duration-response summary for ESPRIT trial data. Time is recorded in hours;* $\tau_u = 24$ *and* $t_\mu = 20$. *Estimates are shown with estimated standard errors given in parentheses.*

	$\hat{\beta}$	$\tilde{\beta}$	β^*
β_0	-2.54 (0.13)	-2.86 (0.16)	-2.78 (0.16)
β_1	0.12 (0.07)	-0.09 (0.02)	0.16 (0.07)

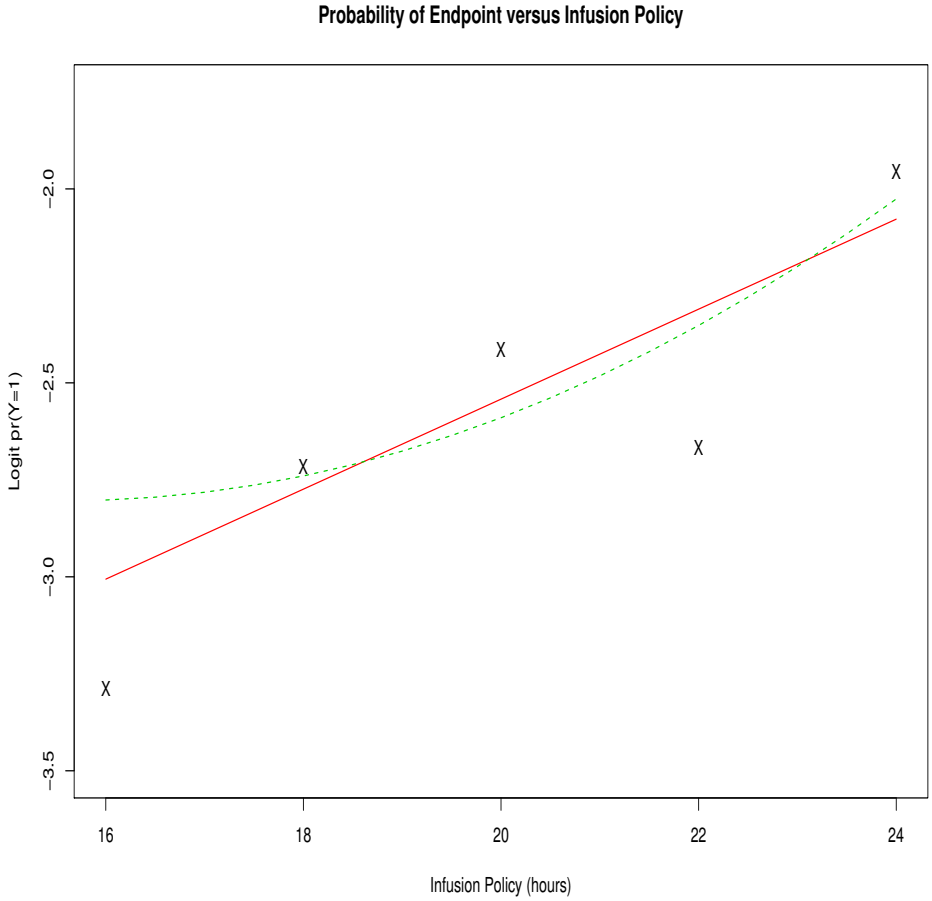

Fig. 1. Summary of ESPRIT infusion trial data. Fitted curves are shown for a linear fit, solid line, and a quadratic fit, dashed line. Infusion policy point estimates are displayed, as crosses, for a discretisation analysis at time equal to 16, 18, 20, 22 and 24 hours.