Statistics
in Medicine

Research Article

Received 19 June 2012, Accepted 21 August 2013 Published online 20 September 2013 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.5987

A global logrank test for adaptive treatment strategies based on observational studies

Zhiguo Li,[a*†] Marcia Valenstein,[b,c] Paul Pfeiffer[b,c] and Dara Ganoczy[c]

In studying adaptive treatment strategies, a natural question that is of paramount interest is whether there is any significant difference among all possible treatment strategies. When the outcome variable of interest is time-to-event, we propose an inverse probability weighted logrank test for testing the equivalence of a fixed set of pre-specified adaptive treatment strategies based on data from an observational study. The weights take into account both the possible selection bias in an observational study and the fact that the same subject may be consistent with more than one treatment strategy. The asymptotic distribution of the weighted logrank statistic under the null hypothesis is obtained. We show that, in an observational study where the treatment selection probabilities need to be estimated, the estimation of these probabilities does not have an effect on the asymptotic distribution of the weighted logrank statistic, as long as the estimation of the parameters in the models for these probabilities is $\sqrt{n}$-consistent. Finite sample performance of the test is assessed via a simulation study. We also show in the simulation that the test can be pretty robust to misspecification of the models for the probabilities of treatment selection. The method is applied to analyze data on antidepressant adherence time from an observational database maintained at the Department of Veterans Affairs' Serious Mental Illness Treatment Research and Evaluation Center. Copyright © 2013 John Wiley & Sons, Ltd.

Keywords: adaptive treatment strategy; observational study; survival outcome; weighted logrank test

1. Introduction

In the process of treating chronic diseases or conditions, it is common that the treatment needs to be adjusted over time according to accruing patient information. Typically, treatments are adjusted according to the efficacy, side effects, burden, and compliance, with current and previous treatments. For example, in treating attention deficit hyperactivity disorder, patients may start with low-dose medication and continue the medication if the response to the initial treatment is good; otherwise, the physician may increase the dosage of the medication or switch the patient to behavioral therapy. The aim of adaptively adjusting treatment is to optimize patient outcomes. Therefore, it is of interest to know if different rules to specify treatment over time result in different outcomes and, if so, then what the best rule is. Formally, a specific rule to specify treatment at each decision point over time according to all available historical information up to that point is called an adaptive treatment strategy (abbreviated treatment strategy or strategy), also called a dynamic treatment regime or a treatment policy [1–4]. In the aforementioned example, one possible strategy is to start with low-dose medication and continue low-dose medication if patient responds well; otherwise, switch to behavioral therapy. In most of the cases, there are only two decision points as in the attention deficit hyperactivity disorder example, but sometimes there may be more than two decision points, for which an example (in treating depression) can be found in [5]. In that case, a treatment strategy involves multiple decisions over time. It

[a] *Department of Biostatistics and Bioinformatics, Duke University, Durham, NC, U.S.A.*
[b] *Department of Psychiatry, University of Michigan, Ann Arbor, MI, U.S.A.*
[c] *SMITREC, HSR&D Center of Excellence, Department of Veterans Affairs, Ann Arbor, MI, U.S.A.*
**Correspondence to: Zhiguo Li, Department of Biostatistics and Bioinformatics, Duke University, Durham, NC, U.S.A.*
[†] *E-mail: zhiguo.li@duke.edu*

Copyright © 2013 John Wiley & Sons, Ltd.

is worth mentioning that adaptive treatment strategies can be loosely thought of as adaptive treatment sequences. However, even when treated with the same adaptive treatment strategy, different subjects may follow different treatment sequences, depending on pre-treatment characteristics and intermediate responses to treatments. For example, in the treatment strategy given earlier, some patients (responders to the initial low-dose medication) are continuously treated with low-dose medication, while other patients (nonresponders to the initial low-dose medication) are initially treated with low-dose medication and then switch to behavioral therapy, although both groups of patients are consistent with the same adaptive treatment strategy. Hence, there is a drastic difference between an adaptive treatment strategy and a pre-determined non-adaptive treatment sequence. Our focus here is on adaptive treatment strategies, and methods for dealing with non-adaptive treatment sequences can be found, for example, in [6, 7].

We are motivated by an observational study in antidepressant (AD) adherence in treating depression in veterans in the U.S.A. In this study, data come from an observational database maintained at the Department of Veteran Affairs' Serious Mental Illness Treatment Research and Evaluation Center in Ann Arbor, Michigan. This database has records of depressed veterans filling their prescribed ADs over time, as well as patient demographic and comorbidity information. For any patient, the AD could be continued, switched, or discontinued over a certain period of time. Here, the decision about treatment switch is made with a less structured conversation between the provider and the patient about symptom levels and role functioning instead of a standardized scale. It is well known that in treating depression, one challenge is medication adherence, which can be limited by side effects or insufficient efficacy of common ADs. Based on data from this database, an initial research question was whether there is any difference in adherence time of the seven commonly prescribed ADs, where the adherence time is defined as the time from initiating the AD to switching to another AD or discontinuation of the AD. A simple analysis based on the data using Cox proportional hazards model showed that there is no significant difference among the ADs. A further question raised by the physicians is whether there is any difference in adherence time for different 'combinations' of the first ADs and the possible second ADs. Specifically, does it matter which AD to start with, and if it is decided that the first AD is switched, does it matter which AD to switch to? This is a typical problem of testing for the equivalence of all the adaptive treatment strategies. Note that this is not a problem of comparing non-adaptive treatment sequences because the decision of switching treatment is based on intermediate outcomes after initial treatment. Denote A_1 to be the initial AD, and denote A_2 to be the second AD if there is a switch. Here, a strategy is defined as follows: start with A_1 and switch to A_2 if it is decided that the initial AD needs to be switched; otherwise, stay on A_1. Here, the adherence time is time from initiating the first AD to the time of the second AD switch. The standard logrank test is not applicable to testing the equivalence of all the strategies based on observational data because of the following: (i) the treatments are not randomized, and hence, selection bias resulting from the possibility that those who have worse outcomes may be more likely to choose certain ADs is likely to exist; and (ii) a subject starting with, for example, AD A, who did not have a switch of AD or discontinued the AD (in the study period), is consistent with any treatment strategy with initial treatment A, which makes different comparison groups dependent. Therefore, we need appropriate methods to adjust for the selection bias and to take into account the dependence among strategies.

When the outcome variable is time-to-event, in order to test for the equivalence of all possible treatment strategies based on data from observational studies, we generalize the standard logrank test (for multiple groups) by using inverse probability weighting [8], where the weights can be regarded as the inverse of the probability of being consistent with a certain strategy. If there are two decision points, then this weight involves two probabilities, that is, the probability of selecting the first treatment and the conditional probability of selecting the second treatment given the first treatment. This weight function takes into account both the selection bias and the dependence among different groups mentioned earlier. Our weight function is very similar to the weight function used in [9–11] for analyzing data from two-stage stratified samples. However, in observational studies, the selection probabilities in the weight function are unknown and thus have to be estimated from the observed data. We provide large sample distribution for the test statistic to justify the usage of the weighted logrank test. As we will see, the estimation of the probabilities does not have an effect on the asymptotic distribution of the weighted logrank statistic, as long as the estimation of the parameters in the models for these probabilities is $\sqrt{n}$-consistent. Finally, we argue that a doubly robust procedure to deal with the misspecification of the models for the selection probabilities is practically not very helpful in our case, and we explore the robustness of the approach in a sensitivity analysis in simulation.

Copyright © 2013 John Wiley & Sons, Ltd.

In order to compare different treatment strategies, investigators may also conduct randomized controlled clinical trials, and in this scenario, the so-called sequential multiple assignment randomized trial (SMART) has been the topic of many research articles. For example, [5, 12–15] considered the design and sample size issues, and [16–21] focused on inference for treatment strategies based on data from SMARTs, especially two-stage randomized trials, which is a special SMART with only two decision points. Particularly, [21] considered the test for the equivalence of multiple treatment strategies, but it is not applicable to observational studies. There is not much existing work as we are aware of in the setting of observational studies. In [22], Murphy and collaborators proposed methodology for estimating the mean response for a treatment strategy based on observational studies, but the outcome is assumed to be uncensored. In [23, 24], the authors considered inference for observational studies where the outcome can be censored but the focus is on estimation of the optimal treatment strategies. Our work is motivated by an important practical problem raised by psychiatrists as described earlier. This work is novel because there is no existing methodology that can find if there is any difference in a censored survival outcome among multiple adaptive treatment strategies in the observational study setting. We propose the appropriate methodology and establish the necessary theory for the statistical inference. A well-known advantage of observational studies is their large sample sizes compared with clinical trials. For example, in the VA database described earlier, there are records of 100,517 veterans taking AD medication. With large sample sizes, we have sufficient power to detect any meaningful differences among the strategies. As the number of treatment strategies may be large (usually at least four), testing such a hypothesis is likely to be too ambitious in clinical trials, but usually, this should not be a problem for observational studies.

The rest of the article is arranged as follows. In Section 2, we introduce the notation and assumptions. In Section 3, we describe the weighted logrank test for observational studies and its asymptotic properties under the null hypothesis. We present results from a simulation study in Section 4, and the proposed method is applied to analyze the AD adherence data in Section 5. Finally, we conclude with a discussion in Section 6. The asymptotic distribution of the test statistic is derived in the Appendix.

2. Notation and assumptions

For simplicity, we assume that there are only two decision points. Denote the baseline covariate vector as X_1. At the beginning, a subject chooses a treatment, denoted by A_1, from J possible treatments $1, 2, \ldots,$ and J. If at some time point after the initial treatment and before τ, where τ is the study duration, it is decided that the treatment is switched, the second treatment, denoted by A_2, can be chosen from any of K possible treatments $1, 2, \ldots,$ and K. Denote all the observed data upon this time point as X_2, which may include X_1, A_1, and some measurements about response, side effects, and so on, after A_1 is initiated. In this setting, there are JK possible treatment strategies. We define the following strategy as strategy 'jk': start with treatment $A_1 = j$, and if it is decided to switch treatment, then switch to $A_2 = k$, for $1 \leqslant j \leqslant J$, $1 \leqslant k \leqslant K$. Figure 1 illustrates all treatment strategies in the case where $J = K = 2$. From the figure, we note that the group of subjects who do not have a switch of treatment are consistent with two different treatment strategies.

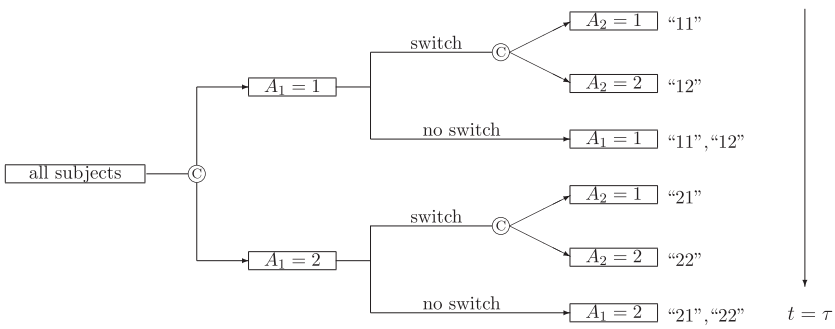

Figure 1. Illustration of treatment strategies in an observational study. C: choosing treatments; switch: decide to switch treatment before time τ; no switch: no switch of treatment before τ. The quoted data on the right are the strategies that the sequences are consistent with.

Copyright © 2013 John Wiley & Sons, Ltd.

The hypothesis we test is on counterfactual times to event [25], which we now define. Let T_j^* be the time-to-event if the subject is treated with initial treatment $A_1 = j$, and let S_j be the time to treatment switch under initial treatment $A_1 = j$. Denote T_{jk}^* to be the time from switching treatment to the time of event if the subject is initially treated with $A_1 = j$ and then switches to $A_2 = k$, which is only defined if $R_j^* = I\left(S_j < T_j^*\right) = 1$. Denote T_{jk} to be counterfactual time-to-event if the subject follows strategy jk, for $1 \leqslant j \leqslant J$ and $1 \leqslant k \leqslant K$. The variable T_{jk} is related to the aforementioned variables by the relationship $T_{jk} = \left(1 - R_j^*\right) T_j^* + R_j^* \left(S_j + T_{jk}^*\right)$. All observed data without censoring consist of $(X_1,\ A_1,\ R^*,\ R^*A_2,\ R^*X_2,\ R^*S,\ T)$, where $R^* = R_{A_1}^*$, $S = S_{A_1}$, and $T = \left(1 - R_{A_1}^*\right) T_{A_1}^* + R_{A_1}^* \left(S_{A_1} + T_{A_1A_2}^*\right)$. If the data are subject to random censoring and C is the censoring time, then all the observed data are $(X_1,\ A_1,\ X_2,\ R,\ RA_2,\ RX_2,\ V,\ U)$, where $U = T \wedge C$, $V = S \wedge U$, and $R = I\ (S \leqslant U)$ is the indicator for switching treatment. Suppose the sample size is n. Then the observed data for all n subjects are $(X_{1i},\ A_{1i},\ X_{2i},\ R_i,\ R_iA_{2i},\ R_iX_{2i},\ V_i,\ U_i)$, $1 \leqslant i \leqslant n$, which are n i.i.d. copies of $(X_1,\ A_1,\ X_2,\ R,\ RA_2,\ RX_2,\ V,\ U)$. As usual, define the counting process $N(t) = I(T \leqslant t, T \leqslant C)$ and the 'at-risk' process $Y(t) = I(T \geqslant t, C \geqslant t)$. Denote $\Lambda_{jk}(t)$ to be the cumulative hazard function of T_{jk}, $1 \leqslant j \leqslant J, 1 \leqslant k \leqslant K$. Then the hypothesis to be tested is

$$H_0 : \Lambda_{jk}(t) \equiv \Lambda(t), 1 \leqslant j \leqslant J,\ 1 \leqslant k \leqslant K,$$

where $\Lambda(t)$ is the common cumulative hazard function under H_0.

Suppose the probability of choosing the initial treatment $A_1 = j$ depends only on observed baseline covariate vector X_1 and is $P\ (A_1 = j|X_1) = p_j\ (X_1)$, $1 \leqslant j \leqslant J$. Suppose that the probability of choosing second treatment $A_2 = k$ depends only on X_2 and is denoted by $P\ (A_2 = k|X_2, R = 1) = q_k\ (X_2)$, $1 \leqslant k \leqslant K$. Note that X_2 contains A_1, and hence, the choice of second-stage treatment may depend on the first treatment. We assume that $p_j(X_1) > 0$ and $q_k(X_2) > 0$ for all X_1 and X_2, for $1 \leqslant j \leqslant J$ and $1 \leqslant k \leqslant K$. We make the consistency assumption [26] that for a subject who actually followed treatment strategy jk, the observed time-to-event of this subject is equal to its counterfactual time-to-event under strategy jk, for all strategies. We also assume that there are no unmeasured confounders, which means that the treatment selection at each stage is independent of the potential times to event conditioning on the covariate history and treatment history [26]. Finally, as usual, the censoring time C is assumed to be independent of the time-to-event T.

3. Weighted logrank test for observational studies

The standard logrank test for multiple groups is described in [27]. Suppose there are $p + 1$ groups to be tested, indexed by $j = 0, 1, \ldots, p$. The ith individual in the jth group is indexed by ji, $1 \leqslant i \leqslant n_j$, and $0 \leqslant j \leqslant p$. The total sample size is $n = \sum_{j=0}^{p} n_j$. The null hypothesis is $\tilde{H}_0 : \Lambda_j(u) \equiv \Lambda(u)$, $0 \leqslant j \leqslant p$, where Λ_j is the cumulative hazard function for the time-to-event for the jth group, and Λ is the common cumulative hazard function under $\tilde{H}_0$. Let $Y_{ji}(t)$ and $N_{ji}(t)$ be the at-risk process and the counting process for the time-to-event of interest, respectively, for the ith subject in the jth group. Let $k.$ denote summation over subjects in the kth group, and let the double dots $(\cdot\cdot)$ denote summation over all subjects. Define

$$d_j = \sum_{k=0}^{p} \int_0^{\tau} \left\{\delta_{jk} - \frac{Y_{j.}(u)}{Y_{..}(u)}\right\} dN_{k.}(u),\ 0 \leqslant j \leqslant p,$$

where $\delta_{jk} = I(j = k)$, and let $d = (d_1, \ldots, d_p)^T$. Let $\hat{D}$ be the estimate of covariance matrix of d (see [27], page 171 for details). The logrank statistic for testing $\tilde{H}_0$ is then defined as $U = d^T \hat{D}^{-1} d$. It can be shown that the asymptotic distribution of U under $\tilde{H}_0$ is χ_p^2.

For testing the equivalence of all treatment strategies based on observational data, our test is based on a similar statistic as mentioned earlier, but we use inverse probability weighting [8] to adjust for potential selection bias. Weights in a simpler form are also needed for SMARTs (see, for example, [15, 21]). If the probabilities $p_j(X_1)$ and $q_k(X_2)$ for treatment selection are known, the weight function used for strategy jk is defined as

$$W_{jk} = \frac{I\ (A_1 = j)}{p_j\ (X_1)} \left\{1 - R + \frac{I(A_2 = k)}{q_k(X_2)} R\right\}, \tag{1}$$

Copyright © 2013 John Wiley & Sons, Ltd.

for $1 \leqslant j \leqslant J,\ 1 \leqslant k \leqslant K$. For a given value of R, the weight is the inverse probability that the subject is consistent with strategy jk. In practice, the probabilities $p_j(X)$ and $q_k(X)$ are unknown in observational studies, so we estimate them first, usually based on some parametric models $p_j(X_1) = p_j(X_1, \beta)$ and $q_k(X_2) = q_k(X_2, \eta)$, by estimating the unknown (finite dimensional) parameters $\theta = \left(\beta^T, \eta^T\right)^T$. Denoting $\hat{\theta}$ to be the estimate of θ, we use the following estimated weight:

$$W_{jk}(\hat{\theta}) = \frac{I\left(A_1 = j\right)}{p_j(X_1, \hat{\beta})} \left\{1 - R + \frac{I\left(A_2 = k\right)}{q_k\left(X_2, \hat{\eta}\right)} R\right\},$$

for strategy jk, for $1 \leqslant j \leqslant J$ and $1 \leqslant k \leqslant K$. By using this weight function, we account not only for the potential selection bias but also the fact that a single subject may be consistent with two or more than two strategies. If this happens, then for this subject, the weight function is nonzero for all strategies that this subject is consistent with. This weight function is constant and does not depend on the time. A time-dependent weight that is more efficient was proposed by [17] and has been used by other authors thereafter. This weight function is defined as

$$W_{jk}(t; \hat{\theta}) = \frac{I\left(A_1 = j\right)}{p_j(X_1, \hat{\beta})} \left\{1 - R(t) + \frac{I\left(A_2 = k\right)}{q_k\left(X_2, \hat{\eta}\right)} R(t)\right\},$$

where $R(t) = I(S \leqslant \min(t, T, C))$. The reason that this weight is more efficient is discussed in [15, 17]. For simplicity of presentation, in this section, we assume that the constant weight is used.

It is necessary to point out that, although we assume that there are two decision points (two stages), our method generalizes to multiple stages straightforwardly, as also pointed out by [15]. If there are multiple stages, the only difference is in the weight function. Instead of the inverse of the multiplication of two probabilities, when there are multiple stages, the weight function would be the inverse of the multiplication of multiple probabilities, one for the selection of treatment in each stage. We focus on the two-stage case because in most practical situations only two stages are considered.

Let $\delta_{jk,j'k'} = I(j = j', k = k')$. Similar as the d_js defined earlier, for strategy jk, we define

$$L_{jk} \quad = \quad \sum_{j'=1}^{J} \sum_{k'=1}^{K} \int_0^{\tau} \left[\delta_{jk,j'k'} - \frac{\hat{\bar{Y}}_{Wjk}(u)}{\hat{\bar{Y}}_W(u)}\right] d\hat{\bar{N}}_{Wj'k'}(u),$$

where $\hat{\bar{Y}}_{Wjk}(u) = \sum_{i=1}^{n} W_{jk,i}(\hat{\theta}) Y_i(u)/n$, $\hat{\bar{N}}_{Wjk}(u) = \sum_{i=1}^{n} W_{jk,i}(\hat{\theta}) N_i(u)/n$, for $1 \leqslant j \leqslant J$ and $1 \leqslant k \leqslant K$, and $\hat{\bar{Y}}_W(u) = \sum_{j=1}^{J} \sum_{k=1}^{K} \hat{\bar{Y}}_{Wjk}(u)$. Let

$$L = (L_{11}, \ldots, L_{1K}, L_{21}, \ldots, L_{JK-1})^T.$$

To state the asymptotic distribution for L, we need some additional notation. Define $M(t) = N(t) - \int_0^t Y(u) d\Lambda(u)$. Denote θ_0 to be the true value of θ. Let $l = (l_{11}, \ldots, l_{1K}, l_{21}, \ldots, l_{JK-1})^T$, where

$$l_{jk} = \sum_{j'=1}^{J} \sum_{k'=1}^{K} \left(\delta_{jk,j'k'} - \frac{1}{JK}\right) W_{j'k'}(\theta_0) \int_0^{\tau} dM(u).$$

The following theorem gives the asymptotic distribution of $\sqrt{n}L$ under H_0. The proof of the theorem is deferred to the Appendix.

Theorem 1

Assume that the probabilities $p_j(X_1, \beta)$ and $q_k\left(X_2, \eta\right)$ are bounded away from 0, for $1 \leqslant j \leqslant J$ and $1 \leqslant k \leqslant K$. Also assume that $\sqrt{n}\left(\hat{\theta} - \theta_0\right) = O_p(1)$. Then $\sqrt{n}L \to_d N(0, \Sigma)$ under H_0 , as $n \to \infty$, where $\Sigma = \mathrm{var}(l)$.

From the proof of the theorem, it is easy to see that, if the true weights $W_{jk}(\theta_0)$ were used in constructing L, the asymptotic distribution of $\sqrt{n}L$ under H_0 would still be $N(0, \Sigma)$. Hence, the estimation

Copyright © 2013 John Wiley & Sons, Ltd.

of θ does not affect the asymptotic behavior of L, as long as $\sqrt{n}\left(\hat{\theta}-\theta_0\right)=O_p(1)$. In order to construct the test statistic, we need to estimate the covariance matrix Σ. Denote

$$\hat{l}_{jk,i}=\sum_{j'=1}^{J}\sum_{k'=1}^{K}\left(\delta_{jk,j'k'}-\frac{1}{JK}\right)W_{j'k',i}(\hat{\theta})\int_0^{\tau}\left\{dN_i(u)-Y_i(u)d\hat{\Lambda}(u)\right\},$$

where $\hat{\Lambda}(t)$ is obtained (under H_0) from all subjects by the usual Nelson–Aalen estimator. Let $\hat{l}_i=\left(\hat{l}_{11,i},\ldots,\hat{l}_{1K,i},\hat{l}_{21,i},\ldots,\hat{l}_{JK-1,i}\right)^T$, $1\leqslant i\leqslant n$. Then the covariance matrix Σ can be consistently estimated by $\hat{\Sigma}=\sum_{i=1}^{n}\hat{l}_i\hat{l}_i^T/n$. Now, the test statistic for H_0 is defined as

$$T=L^T\hat{\Sigma}^{-1}L.$$

By Theorem 1, under H_0, nT converges to χ^2_{JK-1} in distribution as $n\to\infty$. Under significance level α, we reject H_0 when $T>\chi^2_{JK-1}(\alpha)$, where $\chi^2_{JK-1}(\alpha)$ is the upper $100\times\alpha$ quantile of the χ^2_{JK-1} distribution. Note that the test of the equivalence of any subset of all the treatment strategies (for example, two strategies) follows directly from the same idea.

Double robustness is often considered in similar circumstances when the assumed model for the estimated probabilities may not be correct [28]. In a missing data problem, a doubly robust estimator is asymptotically unbiased when either the assumed model for the mean outcome is correct or the model for the missing probability is correct. In other words, the statistic is unbiased even when the assumed probability model is incorrect as long as the assumed model for the mean outcome is correct. A similar property also holds for our inverse probability weighted logrank statistic described earlier. However, here, what we need is not the correct model for the mean outcome (of the failure time T); instead, we need the correct model for the mean of a complicated quantity involving a stochastic integral with respect to the counting process of the failure time (the details of which are omitted here). As the mean of the this complicated quantity does not have a simple and intuitive interpretation as the mean outcome of interest, it is difficult to have a correct model for this quantity. Consequently, it is not very helpful to investigate the double robustness of the aforementioned weighted logrank statistic. Nevertheless, we can perform a sensitivity analysis to assess the robustness of the inverse probability weighted logrank test to the estimated weights. In the sensitivity analysis, we intentionally leave out some of the important covariates in fitting the probabilities of treatment selection and check the resulting significance level of the test procedure to see the extent to which it is affected when these covariates are left out.

4. Simulation

We carry out a Monte Carlo simulation study to assess the finite sample performance of the weighted logrank test. For simplicity, we assume $J=K=2$.

Assume that the baseline covariate vector is two dimensional, $X_1=(X_{11},X_{12})^T$, X_{11} follows a Bernoulli distribution with success probability 0.3, and X_{12} follows a $N(0,1)$ distribution truncated at -1 and 1 to make it bounded. Given X_1, the probability of choosing $A_1=1$ as the initial treatment is $P(A_1=1|X_1)=e^{\beta^TX_1}/\left(1+e^{\beta^TX_1}\right)$, where $\beta=(0.5,0.5)^T$. If it is decided that the initial treatment is switched ($R=1$), then the probability of choosing k as the second treatment is $P(A_2=1|X_2,R=1)=e^{\eta^TX_1}/\left(1+e^{\eta^TX_1}\right)$, where $\eta=(1,-0.8)$. We first generate the counterfactual times to event $\left(S_j,T_j^*,T_{j1}^*,T_{j2}^*\right)$, $j=1,2$. As these variables are likely to be dependent, we generate them from a Frank copula model [29] with association parameter 0.5 and 0.7 for $j=1,2$, respectively, which makes the components positively associated. The marginal distributions of all the components are exponential distributions with parameters assuring that H_0 holds. The dependence of these times on X_1 is enforced by a proportional hazards model. The censoring time C is assumed to be independent of all other variables and be uniform in $(0,100)$. Assuming $\tau=35$, these parameters generate about 30% censored subjects. The sample sizes we consider are $n=100$, 500, 1000, 2000, and 10,000, which are reasonable in observational studies.

Under these scenarios, we generate observed data from the aforementioned generative model. Then we fit linear logistic models for $p_j(X_1)$ and $q_k(X_2)$ using covariates X_1 and X_2, respectively, obtain the fitted values of these functions, and calculate the (estimated) weights, including both the constant

Copyright © 2013 John Wiley & Sons, Ltd.

Table I. Empirical significance levels of the weighted logrank test and unweighted logrank test in observational studies in simulation.

Method	$n = 100$	$n = 500$	$n = 1000$	$n = 2000$	$n = 10,000$
Weighted	0.038	0.042	0.047	0.045	0.044
Unweighted	0.243	0.478	0.882	0.990	1
MIS1*	0.031	0.059	0.041	0.040	0.055
MIS2†	0.037	0.039	0.045	0.059	0.064

*Using the weighted method but leaving the first covariate out in fitting the probabilities of selecting treatments.
†Using the weighted method but leaving the second covariate out in fitting the probabilities of selecting treatments.

Table II. Comparison of the empirical covariance matrix of the weighted logrank statistic $L = (L_{11}, L_{12}, L_{21})^T$ with the mean of its estimated covariance matrices in the simulation ($n = 500$).

Empirical covariance matrix			Mean of the estimated covariance matrices		
0.602	−0.035	−0.347	0.641	−0.064	−0.332
−0.035	0.545	−0.268	−0.064	0.583	−0.303
−0.347	−0.268	0.604	−0.332	−0.303	0.651

weights and time-dependent weights, for all strategies for all subjects. Finally, L and $\hat{\Sigma}$ are calculated to form the test statistics. This procedure is repeated 1000 times, and the observed significance level, that is, the percent of times H_0 is rejected, is obtained. In the simulation, we also compare the results of the weighted logrank test with the test without weight (equivalently, assume $p_j(X_1) \equiv 1$ and $q_k(X_2) \equiv 1$ for all j and k in the weight functions. In addition, we check the robustness of the test procedure by intentionally leaving out one of the covariates X_{11} and X_{12} in fitting the logit models to estimate the probabilities of treatment selection. We obtain empirical significance levels when the covariate is left out.

The results of the simulation are shown in Tables I and II. We see that the empirical significance levels of the weighted logrank test are close to the desired 5% level regardless of the sample size. The components of the empirical covariance matrix of the statistic L are also close to the mean of the estimated covariance matrix. However, if the logrank test is used without the inverse probability weighting, the observed significance level can be far away from the desired level. The last two rows of Table I show that the empirical significance level of the test is not affected much if one of the two covariates X_{11} and X_{12} is left out in fitting the logit models for the probabilities of treatment selection. Results of additional simulation, which are not shown here, indicate that the empirical significance level is affected more, especially in very large samples, when X_{12} is left out in the model if X_{12} has a wider range of values, for example, when it has a $N(0, 1)$ distribution truncated at -2 and 2. Therefore, empirically, we see that our procedure is robust at least when important covariates that are left out have relatively narrow ranges. The results of using time-dependent weights, which are not shown here, indicate that the observed significance levels are very close to those with constant weights. Finally, we obtained simulation results in a case where the other assumptions are similar but $P(A_2 = 1|X_2, R = 1) = e^{\eta X_{11}} / \left(1 + e^{\eta X_{11}}\right)$, that is, the selection of the second treatment only depends on X_{11}. The results are very similar to those in Table I and are thus omitted.

5. Analysis of antidepressant adherence data

The database at the Department of Veteran Affairs' Serious Mental Illness Treatment Research and Evaluation Center has records of depressed veterans filling their prescribed ADs over time, which include the dates of AD fill, the amount, and AD names. According to these data, we obtain for each subject the adherence time, that is, time on drug before the second AD switch or discontinuation of the second AD. As guidelines usually recommend treatment with ADs for at least 6 months, the study duration is set to be $\tau = 6$. Specifically, the adherence time is calculated as follows. If a patient had a second AD switch in the 6-month period, the adherence time is time from initiating the first AD to the time of the second AD switch. For a patient who only had one AD switch in the 6-month period, the adherence time is time from initiating the first AD to the time of discontinuation of the second AD if the second AD was discontinued before 6 months, and it is censored at 6 months if the second AD was continued beyond 6 months. For a

Copyright © 2013 John Wiley & Sons, Ltd.

Table III. Choices of the first and second ADs (if there was a switch) among the 100,517 patients.

		Second AD							
		Bupropion	Citalopram	Fluoxetine	Mirtazapine	Paraxetine	Sertraline	Venlafaxin	NA
First AD	Bupropion	0	574	215	174	472	159	154	10839
	Citalopram	1338	0	603	791	685	571	479	32231
	Fluoxetine	480	706	0	261	392	184	156	13781
	Mirtazapine	199	417	115	0	378	127	96	6118
	Paraxetine	240	289	88	132	0	112	103	6615
	Sertraline	434	529	124	256	270	0	161	13043
	Venlafaxin	216	202	67	124	79	66	0	4672

NA, no switch of treatment.

patient who did not have AD switch in the 6-month period, the adherence time is time from initiating the first AD to the time of discontinuation of the first AD if the first AD was discontinued before 6 months, and it is censored at 6 months if the first AD was continued beyond the 6 months. In addition to this, the adherence time may be censored by death if the patient died in 6 months after initiating the first AD. We are interested in knowing if there is any difference in adherence time if the patient is treated by different treatment strategies, which are defined in Section 1. All of the patients started with one of seven ADs, which include bupropion, citalopram, fluoxetine, mirtazapine, paraxetine, sertraline, and venlafaxin. For patients who had at least one switch within 6 months, most of them switched to one of the remaining six ADs other than the first one, but a small proportion of them switched to one of three ADs that are not one of the seven initial ADs. However, as the proportion of subjects who switched to the other three ADs is very small (less than 1%), they are excluded in our analysis. Table III lists the number of patients who started each of the seven ADs and the number of patients who switched to each of the remaining six ADs or did not have a switch. The total number of subjects included in our analysis is 100,517. There are a total number of $7 \times 6 = 42$ strategies to be compared. The database also includes the following demographic variables: date of birth, gender, race, ethnicity, and data for comorbidities including major depression, personality disorder, alcohol abuse, and drug abuse.

Based on these data, we first fit a multicategory linear logistic model for the choice of the first AD, including all the demographic and comorbidity variables as covariates, based on which we calculate the fitted probabilities of choosing each of the seven ADs as the initial drug. Next, we fit a second linear logistic model for the choice of the second AD for those who switched AD in the 6-month period, including all the aforementioned variables as well as the first AD as covariates. Based on this, the fitted probabilities for choosing each of the seven ADs as the second AD are calculated. Then we calculate the weights associated with each of the 42 strategies for each subject. Finally, we obtain the weighted logrank test statistic. The test statistic is 7.55, which yields a p value close to 1 if compared with a χ^2 distribution with 41 degrees of freedom. We conclude that there is no significant difference among the 42 treatment strategies. An analysis using time-dependent weights yields similar results. There are two possibilities for this: (i) it really does not matter which treatment to start with and which treatment to switch to even for those who decided to switch treatment; or (ii) it does matter which treatment to switch to for those who decided to switch treatment, but as most of subjects did not switch treatment, it does not matter which treatment strategy a subject takes in the population level.

Finally, we make a cautionary note in making the conclusion. We showed by simulation that our test is robust to misspecification of the treatment selection probabilities to some extent, but we cannot exclude the possibility that the test is not robust when the probabilities are highly misspecified. This may happen when covariates that influence the selection probabilities with wide ranges are not included in the models for these probabilities. For example, the depression score at visit was not collected in our database and thus was not included in the probability model. As the depression score is likely a variable that influences the decision on AD selection and does not have a narrow range, exclusion of it may lead to erroneous conclusion. Fortunately, as our result is highly insignificant, it is unlikely that missing some covariates can make such a big difference that the result becomes significant. In practice, what one can do is to try the best to include all possible variables that potentially influence the treatment selection. As the goal is to fit the probabilities for the observed subjects as well as possible, as in the application of the propensity score method, overfitting is not a problem here. Moreover, physicians play an important role in this

Copyright © 2013 John Wiley & Sons, Ltd.

aspect as they usually have a good idea about the potential factors that affect treatment selection. If it is ascertained that all important variables are included, then we are confident in making the conclusion. Otherwise, caution should be taken, especially when the result is borderline significant or insignificant.

6. Discussion

When faced with multiple adaptive treatment strategies, a natural question is whether all the strategies are equivalent. Motivated by a real example in psychiatric studies, we present a weighted logrank test for testing the equivalence of all possible adaptive treatment strategies in an observational study. The probabilities in the weight function need to be estimated from the observed data. However, the estimation of these probabilities does not affect the asymptotic behavior of the weighted logrank statistic, as long as estimation of the parameters in the parametric models for these probabilities is $\sqrt{n}$-consistent. A doubly robust procedure is not very helpful in this scenario because we need to correctly specify the mean of a stochastic integral, which is difficult because it does not have a simple and intuitive interpretation. In a simulation study, we show that the test is pretty robust when important covariates are left out in fitting the probabilities of treatment selection when the range of these covariates is relatively narrow.

Appropriate test of all treatment strategies requires the positivity assumption described in Section 2. When the positivity assumption does not hold, that is, the data for some of the treatment strategies are sparse, care needs to be taken when using the proposed approach to compare all the strategies. The estimation of weights associated with strategies with sparse data are likely to be unstable. In this case, one may want to exclude those strategies with sparse data in the test.

In our study design, we assume the treatment is switched in the second stage if 'it is decided that the treatment needs to be switched'. This formulation is quite general and includes various possibilities. There are many reasons to decide that the treatment needs to be switched, which include nonresponse (efficacy), serious side effects, high burden, and so on. Also, it does not matter who makes the decision for treatment switch. For example, it could be made by the physician or the patient or by both. Moreover, 'switch' can be 'switch to a different treatment' or 'stay on the same treatment but with dose or intensity change'. Our method applies to all types of observational studies, as long as we construct the appropriate weight functions for the particular design. A difference between SMARTs and observational studies is in the timing of treatment switch. In SMARTs, the treatment switch usually occurs at prespecified time points (e.g., at 3 months after initial treatment when the response nor nonresponse status is checked) or can only occur at finite time points (e.g., during scheduled visits), while in observational studies the treatment may be switched at any time. This difference does not affect the construction of the weight function, as from (1) the weight only depends on the treatment switch indicator but not the timing of switch.

If the test yields an insignificant result as in the earlier AD adherence example, then there is no necessity for further exploration. However, if the test yields a significant result, then the most interesting question in the next step is to find the best strategy or a set of best strategies. The statistical selection method [30] can be used for this purpose. The selection based on (weighted) logrank statistic is not straightforward, and this will be explored in future research.

Appendix

Here, we provide a sketch of the derivation of asymptotic distribution of the proposed weighted logrank statistic given in Theorem 1. In the following, we always assume that H_0 holds, and assume constant weights are used for simplicity of presentation.

Define

$$r_{jk,n}(u,\theta) = \sum_{i=1}^{n} W_{jk,i}(\theta) Y_i(u) \Big/ \left\{ \sum_{j=1}^{J} \sum_{k=1}^{K} \sum_{i=1}^{n} W_{jk,i}(\theta) Y_i(u) \right\},$$

for $1 \leqslant j \leqslant J$ and $1 \leqslant k \leqslant K$. At first, by simple algebra, we have

$$\sqrt{n} L_{jk} = \frac{1}{\sqrt{n}} \sum_{i=1}^{n} \sum_{j'=1}^{J} \sum_{k'=1}^{K} W_{j'k',i}(\hat{\theta}) \int_0^{\tau} \left\{ \delta_{jk,j'k'} - r_{jk,n}(u,\hat{\theta}) \right\} dM_i(u).$$

Copyright © 2013 John Wiley & Sons, Ltd.

Denote $\dot{W}(\theta) = \partial W(\theta)/\partial\theta$ and $\dot{r}_{jk,n}(u,\theta) = \partial r_{jk,n}(u,\theta)/\partial\theta$. By Taylor expansion, we have, for some θ^* between $\hat{\theta}$ and θ_0,

$$\begin{aligned}
&\frac{1}{\sqrt{n}}\sum_{i=1}^{n} W_{j'k',i}(\hat{\theta})\int_0^\tau \left\{\delta_{jk,j'k'} - r_{jk,n}(u,\hat{\theta})\right\} dM_i(u)\\
=\ &\frac{1}{\sqrt{n}}\sum_{i=1}^{n} W_{j'k',i}(\theta_0)\int_0^\tau \left\{\delta_{jk,j'k'} - r_{jk,n}(u,\theta_0)\right\} dM_i(u)\\
&+\sqrt{n}(\hat{\theta}-\theta_0)^T A_n + \sqrt{n}(\hat{\theta}-\theta_0)^T B_n + \sqrt{n}(\hat{\theta}-\theta_0)^T C_n \sqrt{n}(\hat{\theta}-\theta_0), \qquad (2)
\end{aligned}$$

where

$$A_n = \frac{1}{n}\sum_{i=1}^{n} \dot{W}_{j'k',i}(\theta_0)\int_0^\tau \left\{\delta_{jk,j'k'} - r_{jk,n}(u,\theta_0)\right\} dM_i(u)$$

$$B_n = \frac{1}{n}\sum_{i=1}^{n} W_{j'k',i}(\theta_0)\int_0^\tau \left\{\delta_{jk,j'k'} - \dot{r}_{jk,n}(u,\theta_0)\right\} dM_i(u)$$

and

$$C_n = \frac{1}{n^{3/2}}\left.\frac{\partial^2\left[\sum_{i=1}^n W_{j'k',i}(\theta)\int_0^\tau \left\{\delta_{jk,j'k'} - r_{jk,n}(u,\theta)\right\} dM_i(u)\right]}{\partial\theta\partial\theta^T}\right|_{\theta=\theta^*}.$$

It is easy to show that $\sup_{u\in[0,\tau]}|r_{jk,n}(u,\theta_0) - 1/JK| = o_p(1)$. From this, it follows that

$$A_n = \frac{1}{n}\sum_{i=1}^{n} \dot{W}_{j'k',i}(\theta_0)\int_0^\tau \left(\delta_{jk,j'k'} - \frac{1}{JK}\right) dM_i(u) + o_p(1). \qquad (3)$$

By the proof of Theorem 1 in [17], we have

$$E_\theta W_{j'k'}(\theta)\int_0^\tau \xi(u)dM(u) \equiv 0 \qquad (4)$$

for any deterministic function $\xi(u)$, where E_θ means the expectation is taken assuming the true parameter is θ. Therefore,

$$\begin{aligned}
E_\theta \dot{W}_{j'k'}(\theta)\int_0^\tau \left(\delta_{jk,j'k'} - \frac{1}{JK}\right) dM(u) &= \frac{\partial}{\partial\theta}E_\theta W_{j'k'}(\theta)\int_0^\tau \left(\delta_{jk,j'k'} - \frac{1}{JK}\right) dM(u) \qquad (5)\\
&= 0.
\end{aligned}$$

This, combined with (3) and the law of large numbers, yields $A_n = o_p(1)$. Similar as the preceding equation, we can show that $\sup_{u\in[0,\tau]}|\dot{r}_{jk,n}(u,\theta_0) - \dot{r}_{jk,0}(u,\theta_0)| = o_p(1)$ for a deterministic function $\dot{r}_{jk,0}(u,\theta_0)$. Consequently,

$$\begin{aligned}
B_n &= \frac{1}{n}\sum_{i=1}^{n} W_{j'k',i}(\theta_0)\int_0^\tau \left\{\delta_{jk,j'k'} - \dot{r}_{jk,0}(u,\theta_0)\right\} dM_i(u) + o_p(1)\\
&= o_p(1),
\end{aligned}$$

by the law of large numbers. In a similar manner, we can show that

$$\frac{1}{n}\left.\frac{\partial^2\left[\sum_{i=1}^n W_{j'k',i}(\theta)\int_0^\tau \left\{\delta_{jk,j'k'} - r_{jk,n}(u,\theta)\right\} dM_i(u)\right]}{\partial\theta\partial\theta^T}\right|_{\theta=\theta^*} = o_p(1)$$

Copyright © 2013 John Wiley & Sons, Ltd.

and hence $C_n = o_p\left(n^{-1/2}\right)$. As we assume $\sqrt{n}(\hat{\theta} - \theta_0) = O_p(1)$, by the preceding results and (2), we obtain

$$\sqrt{n}L_{jk} = \frac{1}{\sqrt{n}}\sum_{i=1}^{n} W_{j'k',i}(\theta_0)\int_0^\tau \left\{\delta_{jk,j'k'} - r_{jk,n}(u,\theta_0)\right\} dM_i(u) + o_p(1). \tag{6}$$

Now, we define the classes of functions

$$\mathcal{F}_{jk,j'k'} = \left\{W_{j'k'}(\theta_0)\int_0^\tau \left\{\delta_{jk,j'k'} - \xi(u)\right\} dM(u) :\ \xi(u) \in \Phi\right\},$$

where

$$\Phi = \{\xi_1(u)/\xi_2(u) :\ \xi_1(u) \text{ and } \xi_2(u) \text{ are increasing functions on } [0,\tau]\}.$$

It is shown in [15] that $\mathcal{F}_{jk,j'k'}$ is a Donsker class [31]. Hence, by $\sup_{u\in[0,\tau]}|r_{jk,n}(u,\theta_0) - 1/JK| = o_p(1)$ and the equicontinuity property of empirical processes [31], we have

$$\begin{aligned}&\frac{1}{\sqrt{n}}\sum_{i=1}^{n} W_{j'k',i}(\theta_0)\int_0^\tau \left\{\delta_{jk,j'k'} - r_{jk,n}(u,\theta_0)\right\} dM_i(u)\\ =&\frac{1}{\sqrt{n}}\sum_{i=1}^{n} W_{j'k',i}(\theta_0)\int_0^\tau \left(\delta_{jk,j'k'} - \frac{1}{JK}\right) dM_i(u) + o_p(1).\end{aligned}$$

By this and (6), it follows that

$$\sqrt{n}L_{jk} = \sum_{j'=1}^{J}\sum_{k'=1}^{K}\frac{1}{\sqrt{n}}\sum_{i=1}^{n}\left(\delta_{jk,j'k'} - \frac{1}{JK}\right) W_{j'k',i}(\theta_0)\int_0^\tau dM_i(u) + o_p(1),$$

for all $1 \leqslant j \leqslant J$ and $1 \leqslant k \leqslant K$, and hence

$$\sqrt{n}L = \frac{1}{\sqrt{n}}\sum_{i=1}^{n}\begin{bmatrix} \sum_{j'=1}^{J}\sum_{k'=1}^{K}\left(\delta_{11,j'k'} - \frac{1}{JK}\right) W_{j'k',i}(\theta_0)\int_0^\tau dM_i(u) \\ \sum_{j'=1}^{J}\sum_{k'=1}^{K}\left(\delta_{12,j'k'} - \frac{1}{JK}\right) W_{j'k',i}(\theta_0)\int_0^\tau dM_i(u) \\ \vdots \\ \sum_{j'=1}^{J}\sum_{k'=1}^{K}\left(\delta_{J(K-1),j'k'} - \frac{1}{JK}\right) W_{j'k',i}(\theta_0)\int_0^\tau dM_i(u) \end{bmatrix} + o_p(1).$$

This completes the proof of the theorem.

Remark 1
From this proof, we can see that the fact that the estimation of θ does not have an effect on the asymptotic distribution on $\sqrt{n}L$ is due to (5), which is a result of (4). If (5) did not hold, then we would have an asymptotic representation $\sqrt{n}L = \sum_{i=1}^{n} l_i/\sqrt{n} + H\sqrt{n}\left(\hat{\theta} - \theta_0\right)$ for some nonzero deterministic matrix H, where $l = \left(l_{11}, l_{12}, \ldots, l_{J(K-1)}\right)^T$, and the asymptotic distribution of $\sqrt{n}L$ would then be influenced by that of $\sqrt{n}\left(\hat{\theta} - \theta_0\right)$.

Acknowledgements

We thank three anonymous referees for their helpful comments that improved this article.

References

1. Robins JM. A new approach to causal inference in mortality studies with sustained exposure periods—application to control of the healthy worker survivor effect. *Computers and Mathematics with Applications* 1986; **14**:1393–1512.
2. Lavori PW, Dawson R, Roth AJ. Flexible treatment strategies in chronic disease: clinical and research implications. *Biological Psychiatry* 2000; **48**:605–614. DOI: 10.1016/S0006-3223(00)00946-X.
3. Murphy SA. An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine* 2005; **24**:1455–1481. DOI: 10.1002/sim.2022.

Copyright © 2013 John Wiley & Sons, Ltd.

4. Murphy SA, Lynch KG, Oslin D, McKay JR, TenHave T. Developing adaptive treatment strategies in substance abuse research. *Drug and Alcohol Dependence* 2007; **88**(2):S24–S30. DOI: 10.1016/j.drugalcdep.2006.09.008.
5. Rush AJ, Fava M, Wisniewski SR, Lavori PW, Trivedi MH, Sackeim HA, Thase ME, Nierenberg AA, Quitkin FM, Kashner TM. Sequenced treatment alternatives to relieve depression (STAR*D): rationale and design. *Controlled Clinical Trials* 2004; **25**:119–142. DOI: 10.1016/S0197-2456(03)00112-0.
6. Robins JM, Hernan MA, Brumback B. Marginal structural models and causal inference in epidemiology. *Epidemiology* 2000; **11**(5):550–560.
7. Hernan 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.
8. Robins JM, Rotnitzky A, Zhao LP. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association* 1994; **89**:846–866.
9. Breslow NE, Wellner JA. Weighted likelihood for semiparametric models and two-phase stratified samples, with application to Cox regression. *Scandinavian Journal of Statistics* 2007; **34**:86–102.
10. Li Z, Gilbert P, Nan B. Weighted likelihood method for grouped survival data in case–cohort studies with application to HIV vaccine trials. *Biometrics* 2008; **64**:1247–1255.
11. Li Z, Nan B. Relative risk regression for current status data in case–cohort studies. *Canadian Journal of Statistics* 2011; **39**:557–577.
12. Oetting AI, Levy JA, Weiss RD, Murphy SA. Statistical methodology for a SMART design in the development of adaptive treatment strategies. In *Causality and Psychopathology: Finding the Determinants of Disorders and their Cures*, Shrout PE (ed.). American Psychiatric Publishing, Inc: Arlington VA, 2007.
13. Feng W, Wahed SA. A supremum log rank test for comparing adaptive treatment strategies and corresponding sample size formula. *Biometrika* 2008; **95**(3):695–707. DOI: 10.1093/biomet/asn025.
14. Feng W, Wahed SA. Sample size for two-stage studies with maintenance therapy. *Statistics in Medicine* 2009; **28**:2028–2041. DOI: 10.1002/sim.3593.
15. Li Z, Murphy SA. Sample size formulae for two-stage randomized trials with survival outcomes. *Biometrika* 2011; **98**(3):503–518. DOI: 10.1093/biomet/asr019.
16. Lunceford JK, Davidian M, Tsiatis AA. Estimation of survival distributions of treatment strategies in two-stage randomization designs in clinical trials. *Biometrics* 2002; **58**:48–57.
17. Guo X, Tsiatis AA. A weighted risk set estimator for survival distributions in two-stage randomization designs with censored survival data. *The International Journal of Biostatistics* 2005; **1**(1):1–15. DOI: 10.2202/1557-4679.1000.
18. Wahed SA, Tsiatis AA. Semiparametric efficient estimation of survival distribution in two-stage randomization designs in clinical trials with censored data. *Biometrika* 2006; **93**:163–177. DOI: 10.1093/biomet/93.1.163.
19. Lokhnygina Y, Helterbrand JD. Cox regression methods for two-stage randomization designs. *Biometrics* 2007; **63**:422–428. DOI: 10.1111/j.1541-0420.2006.00707.x.
20. Miyahara S, Wahed SA. Weighted Kaplan–Meier estimators for two-stage treatment regimes. *Statistics in Medicine* 2010; **29**:2581–2591. DOI: 10.1002/sim.4020.
21. Kidwell KM, Wahed SA. Weighted log-rank statistic to compare shared-path adaptive treatment strategies. *Biostatistics* 2013; **14**(2):299–312. DOI: 10.1093/biostatistics/kxs042.
22. Murphy SA, van der Laan MJ, Robins JM, CPPRG. Marginal mean models for dynamic regimes. *Journal of the American Statistical Association* 2001; **96**:1410–1423.
23. Robins JM, Orellana L, Rotnitzky A. Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in Medicine* 2008; **27**:4678–4721.
24. Orellana L, Rotnitzky A, Robins JM. Dynamic regime marginal structural mean models for estimation of optimal dynamic treatment regimes, part I: main content. *The International Journal of Biostatistics* 2010; **6**(2).
25. Holland PW. Statistics and causal inference. *Journal of the American Statistical Association* 1986; **81**:945–960.
26. Robins JM. Causal inference from complex longitudinal data. In *Latent Variable Modelling and Applications to Causality*, Vol. 120, Berkane M (ed.), Lecture Notes in Statistics. Springer: New York, 1997; 69–117.
27. Kalbfleisch JD, Prentice RL. *The Statistical Analysis of Failure Time Data*. John Wiley: New York, 2002.
28. Tsiatis AA. *Semiparametric Theory and Missing Data*. Springer: New York, 2006.
29. Roger BN. *An Introduction to Copulas*. Springer-Verlag: New York, 1998.
30. Bechhofer RE, Santner TJ, Goldsman DM. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. Wiley: New York, 1995.
31. van der Vaart AW, Wellner JA. *Weak Convergence and Empirical Processes*. Springer-Verlag: New York, 1996.

Copyright © 2013 John Wiley & Sons, Ltd.