Electronic Journal of Statistics
Vol. 8 (2014) 1301–1308
ISSN: 1935-7524
DOI: 10.1214/14-EJS906

Discussion of "Dynamic treatment regimes: Technical challenges and applications"*

Jesse Y. Hsu

Department of Biostatistics and Epidemiology, Perelman School of Medicine, University of Pennsylvania
618 Blockley Hall, 423 Guardian Drive, Philadelphia, PA 19104, U.S.A.
e-mail: hsu9@mail.med.upenn.edu

Dylan S. Small

Department of Statistics, The Wharton School, University of Pennsylvania
400 Jon M. Huntsman Hall, 3730 Walnut Street, Philadelphia, PA 19104, U.S.A.
e-mail: dsmall@wharton.upenn.edu

Received May 2014.

1. Introduction

We congratulate Drs. Laber, Lizotte, Qian, Pelham and Murphy on an outstanding review of dynamic treatment regimes (DTRs) [4]. This group has done pioneering work in advancing the theory and applications of DTRs. In a DTR, the treatment type and level are repeatedly adjusted according to an individual's need. An important part of designing DTRs is to choose tailoring variables. A tailoring variable is a variable that is used to decide how to adjust the treatment. Laber et al. say that the current state of the art for choosing tailoring variable is to construct low-dimensional summaries of information that is available at the time of treatment on a subject's status, where the summaries are motivated by clinical judgment, exploratory analyses and convenience. Laber et al. state that an important open problem is the development of formal feature extraction and construction techniques for choosing tailoring variables for DTRs. In this discussion, we discuss a method for choosing tailoring variables when the data available is a simple randomized clinical trial in which the treatment regimes after the initial randomization are set by protocol. This includes trials in which the treatment regime is changed over time based on a patient's status in a set way specified by the protocol. In such trials, at a given time past the initial treatment assignment, there is no variation in how a patient is treated at that time given a patient's status and past treatment history, i.e., every patient who has the same status and treatment history at time t past initial treatment assignment receives the treatment at time t that is specified

*Main article 10.1214/14-EJS920.

by the protocol. Because of this lack of variation in treatment regimes after the initial treatment assignment, such simple randomized clinical trials cannot be used to evaluate DTRs. However, we will show that such simple randomized clinical trials may provide useful data for designing DTRs to evaluate in a later Sequential, Multiple, Assignment Randomized Trial (SMART).

2. Defining a tailoring variable using a post-treatment variable

We consider a simple randomized clinical trial in which R is the randomized treatment (assumed to be binary) and Y is the long run outcome of interest which is measured some time after the treatment. We are interested in choosing tailoring variables which are measured sometime between the administration of treatment R and measurement of the long run outcome Y. Can we learn whether a post-treatment variable D could be useful as a tailoring variable in a DTR? The post-treatment variable D would be a good tailoring variable if the value of D would provide information about whether the current treatment regime is working and should be maintained, or is not working and should be altered. The post-treatment variable D being a good tailoring variable has some connection to D being an effect modifier. We say that a post-randomization variable D is an effect modifier if $E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = d) \neq E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = d')$ for $d \neq d'$, where Y_i^r and D_i^r are potential outcomes for subject i under treatment level r for Y and D. For example, considering binary D, that is, $d \in \{1, 0\}$, one type of effect modification by D is

$$E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 1) > E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 0). \tag{1}$$

This effect modification by D suggests that if $D^{r=1} = 0$ (i.e., $D = 0$ when $R = 1$), it may be beneficial to modify the treatment strategy since the treatment's expected benefit is less than if $D^{r=1} = 1$ (i.e., $D = 1$ when $R = 1$). Note that even if (1) holds and D is an effect modifier, using D as a tailoring variable to modify the treatment strategy might not work because the modification may make things worse. Whether D is an effective tailoring variable needs to be tested out in a SMART trial or an observational study in which there is variation in treatment regimes; a simple randomized clinical trial in which the only time point t at which two patients with the same status and treatment history are given different treatments is the initial treatment assignment can be used to generate hypotheses about what might be good tailoring variables for a DTR but cannot be used to validate a DTR.

We now discuss a method for estimating whether a post-treatment variable D is an effect modifier and needed assumptions. This continues work that Tom Ten Have was doing before his untimely death as discussed in Section 3.3 of Small, Joffe, Lynch, Roy and Localio (2014) [5]. For example, Ten Have (2007) [7] reported proof of concept results that for treating suicide attempters in a way that prevents subsequent depression six months later, the effect of a cognitive behavioral therapy intervention was modified by early hopelessness (hopelessness three months after the therapy had begun) or early suicide ideation (suicide ideation three months after therapy had begun).

One approach to estimating post-randomization effect modification is to use structural mean models as in Faerber, Joffe, Zhang, Brown, Beck and Ten Have (2014) [1] and Stephens, Keele and Joffe (2014) [6]; see also Section 3.3 of Small et al. (2014) [5]. Here, we develop a principal stratification (Frangakis and Rubin, 2002) approach [2]. We consider a binary D. The potential outcomes of D, $D^{r=1}$ and $D^{r=0}$, create four principal strata: (1) 'always-high' consist of subjects who regardless of treatment assignment always have the high value of the post-randomization variable ($D^{r=1} = 1, D^{r=0} = 1$); (2) 'never-high' consist of subjects who regardless of treatment assignment never have the high value of the post randomization variable ($D^{r=1} = 0, D^{r=0} = 0$); (3) 'treatment positively affected' consist of subjects who have the high value of the post-randomization variable when assigned to the treatment but have the low value when assigned to the control ($D^{r=1} = 1, D^{r=0} = 0$); (4) 'treatment negatively affected' consist of subjects who have the low value of the post-randomization variable when assigned to the treatment but have a high value of the post-randomization variable when assigned to the control ($D^{r=1} = 0, D^{r=0} = 1$). We make a monotonicity assumption that the treatment never negatively affects the post-randomization variable, $D^{r=1} \geq D^{r=0}$ so that there are no treatment negatively affected subjects.

Let p_A, p_N, and p_T denote the proportion of always-high, never-high, and treatment positively affected, respectively, such that $p_A + p_N + p_T = 1$:

$$\begin{cases} p_A = \Pr(D_i^{r=1} = 1, D_i^{r=0} = 1) = \int \Pr(D_i^{r=1} = 1, D_i^{r=0} = 1 \mid \mathbf{X}) \Pr(\mathbf{X}) dx \\ p_N = \Pr(D_i^{r=1} = 0, D_i^{r=0} = 0) = \int \Pr(D_i^{r=1} = 0, D_i^{r=0} = 0 \mid \mathbf{X}) \Pr(\mathbf{X}) dx \\ p_T = \Pr(D_i^{r=1} = 1, D_i^{r=0} = 0) = 1 - p_A - p_N \end{cases}.$$

We consider the following model for potential outcomes,

$$\begin{aligned} Y_i^{r=1} = Y_i^{r=0} &+ \beta_1 I(D_i^{r=1} = 1, D_i^{r=0} = 1) \\ &+ \beta_2 I(D_i^{r=1} = 0, D_i^{r=0} = 0) + \beta_3 I(D_i^{r=1} = 1, D_i^{r=0} = 0), \end{aligned} \quad (2)$$

where $Y_i^{r=0}$ is an iid random variable. Denote a matrix of p pre-treatment variables by $\mathbf{X}$. We assume that the association of $\mathbf{X}$ with potential outcomes is linear and the same for each principal stratum:

$$\begin{aligned} E(Y_i^{r=0} \mid D_i^{r=1}, D_i^{r=0}, \mathbf{X}_i) &= \alpha_1 I(D_i^{r=1} = 1, D_i^{r=0} = 1) \\ &+ \alpha_2 I(D_i^{r=1} = 0, D_i^{r=0} = 0) + \alpha_3 I(D_i^{r=1} = 1, D_i^{r=0} = 0) + \boldsymbol{\gamma}^T \mathbf{X}_i. \end{aligned} \quad (3)$$

We can rewrite (1) as

$$\begin{aligned} & E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 1) - E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 0) > 0 \\ \Rightarrow\ & E\{\beta_1 I(D_i^{r=1} = 1, D_i^{r=0} = 1) + \beta_3 I(D_i^{r=1} = 1, D_i^{r=0} = 0) \mid D_i^{r=1} = 1\} \\ & - E\{\beta_2 I(D_i^{r=1} = 0, D_i^{r=0} = 0) \mid D_i^{r=1} = 0\} > 0 \\ \Rightarrow\ & \beta_1 \Pr(D_i^{r=1} = 1, D_i^{r=0} = 1 \mid D_i^{r=1} = 1) + \beta_3 \Pr(D_i^{r=1} = 1, D_i^{r=0} = 0 \mid D_i^{r=1} = 1) \end{aligned}$$

$$
\begin{aligned}
&- \beta_2 \Pr(D_i^{r=1} = 0, D_i^{r=0} = 0 \mid D_i^{r=1} = 0) > 0 \\
\Rightarrow\ & \frac{\beta_1 \Pr(D_i^{r=1} = 1, D_i^{r=0} = 1) + \beta_3 \Pr(D_i^{r=1} = 1, D_i^{r=0} = 0)}{\Pr(D_i^{r=1} = 1)} \\
&- \beta_2 \frac{\Pr(D_i^{r=1} = 0, D_i^{r=0} = 0)}{\Pr(D_i^{r=1} = 0)} > 0 \\
\Rightarrow\ & \frac{\beta_1 \Pr(D_i^{r=1} = 1, D_i^{r=0} = 1) + \beta_3 \Pr(D_i^{r=1} = 1, D_i^{r=0} = 0)}{\Pr(D_i^{r=1} = 1, D_i^{r=0} = 1) + \Pr(D_i^{r=1} = 1, D_i^{r=0} = 0)} \\
&- \beta_2 \frac{\Pr(D_i^{r=1} = 0, D_i^{r=0} = 0)}{\Pr(D_i^{r=1} = 0, D_i^{r=0} = 0)} > 0 \\
\Rightarrow\ & \frac{\beta_1 p_A + \beta_3 p_T}{p_A + p_T} - \beta_2 > 0. \qquad (4)
\end{aligned}
$$

If (4) holds, D is a post-treatment effect modifier and accordingly might be a good tailoring variable.

To identify the parameters in (4), we will assume the stable unit treatment value assumption (SUTVA) under which we can write observed outcome Y as $Y = (1 - R)Y^{r=0} + RY^{r=1}$. The model for the observed outcome given R, D, and $\mathbf{X}$ is then

$$
\begin{aligned}
&E\left(Y_i \mid R_i, D_i, \mathbf{X}_i\right) = E\left\{(1 - R_i)Y_i^{r=0} + R_i Y_i^{r=1} \mid R_i, D_i, \mathbf{X}_i\right\} \\
&= E\left\{ \begin{array}{l} (1 - R_i)Y_i^{r=0} + R_i\{Y_i^{r=0} + \beta_1 I(D_i^{r=1} = 1, D_i^{r=0} = 1)\} \\ +R_i\{\beta_2 I(D_i^{r=1} = 0, D_i^{r=0} = 0) + \beta_3 I(D_i^{r=1} = 1, D_i^{r=0} = 0)\} \end{array} \middle| \, R_i, D_i, \mathbf{X}_i \right\} \\
&= E\left\{ \begin{array}{l} \alpha_1 I(D_i^{r=1} = 1, D_i^{r=0} = 1) + \alpha_2 I(D_i^{r=1} = 0, D_i^{r=0} = 0) \\ +\alpha_3 I(D_i^{r=1} = 1, D_i^{r=0} = 0) + \boldsymbol{\gamma}^T \mathbf{X}_i \\ +\beta_1 R_i I(D_i^{r=1} = 1, D_i^{r=0} = 1) + \beta_2 R_i I(D_i^{r=1} = 0, D_i^{r=0} = 0) \\ +\beta_3 R_i I(D_i^{r=1} = 1, D_i^{r=0} = 0) \end{array} \middle| \, R_i, D_i, \mathbf{X}_i \right\} \\
&= E\left\{ \begin{array}{l} (\alpha_1 + \beta_1 R_i) I(D_i^{r=1} = 1, D_i^{r=0} = 1) \\ +(\alpha_2 + \beta_2 R_i) I(D_i^{r=1} = 0, D_i^{r=0} = 0) \\ +(\alpha_3 + \beta_3 R_i) I(D_i^{r=1} = 1, D_i^{r=0} = 0) \\ +\boldsymbol{\gamma}^T \mathbf{X}_i \end{array} \middle| \, R_i, D_i, \mathbf{X}_i \right\} \\
&= E\left\{ \begin{array}{l} I(R_i = 1, D_i = 1)\{(\alpha_1 + \beta_1 R_i) a_1(\mathbf{X}_i)\} \\ +I(R_i = 1, D_i = 1)[(\alpha_3 + \beta_3 R_i)\{1 - a_1(\mathbf{X}_i)\}] \\ +I(R_i = 1, D_i = 0)(\alpha_2 + \beta_2 R_i) \\ +I(R_i = 0, D_i = 0)\{(\alpha_2 + \beta_2 R_i) a_2(\mathbf{X}_i)\} \\ +I(R_i = 0, D_i = 0)[(\alpha_3 + \beta_3 R_i)\{1 - a_2(\mathbf{X}_i)\}] \\ +I(R_i = 0, D_i = 1)(\alpha_1 + \beta_1 R_i) \\ +\boldsymbol{\gamma}^T \mathbf{X}_i \end{array} \middle| \, R_i, D_i, \mathbf{X}_i \right\},
\end{aligned}
$$

$$= E\left\{\begin{array}{l} R_iD_i\{(\alpha_1+\beta_1R_i)a_1(\mathbf{X}_i)\} \\ +R_iD_i[(\alpha_3+\beta_3R_i)\{1-a_1(\mathbf{X}_i)\}] \\ +R_i(1-D_i)(\alpha_2+\beta_2R_i) \\ +(1-R_i)(1-D_i)\{(\alpha_2+\beta_2R_i)a_2(\mathbf{X}_i)\} \\ +(1-R_i)(1-D_i)[(\alpha_3+\beta_3R_i)\{1-a_2(\mathbf{X}_i)\}] \\ +(1-R_i)D_i(\alpha_1+\beta_1R_i) \\ +\boldsymbol{\gamma}^T\mathbf{X}_i \end{array}\middle| \; R_i, D_i, \mathbf{X}_i \right\}$$

where $a_1(\mathbf{X}_i) = p_{A|\mathbf{X}_i}/(p_{A|\mathbf{X}_i} + p_{T|\mathbf{X}_i})$ and $a_2(\mathbf{X}_i) = p_{N|\mathbf{X}_i}/(p_{N|\mathbf{X}_i} + p_{T|\mathbf{X}_i})$.

Furthermore,

$$E\left(Y_i \mid R_i, D_i, \mathbf{X}_i\right)$$

$$= E\left\{\begin{array}{l} (\alpha_1-\alpha_3)R_iD_ia_1(\mathbf{X}_i) + (\beta_1-\beta_3)R_i^2D_ia_1(\mathbf{X}_i) + \alpha_3R_iD_i \\ +\beta_3R_i^2D_i + \alpha_2R_i + \beta_2R_i^2 - \alpha_2R_iD_i - \beta_2R_i^2D_i + \alpha_3 + \beta_3R_i \\ +(\alpha_2-\alpha_3)a_2(\mathbf{X}_i) + (\beta_2-\beta_3)R_ia_2(\mathbf{X}_i) - \alpha_3R_i - \beta_3R_i \\ -(\alpha_2-\alpha_3)R_ia_2(\mathbf{X}_i) - (\beta_2-\beta_3)R_ia_2(\mathbf{X}_i) - \alpha_3D_i - \beta_3R_iD_i \\ -(\alpha_2-\alpha_3)D_ia_2(\mathbf{X}_i) - (\beta_2-\beta_3)R_iD_ia_2(\mathbf{X}_i) + \alpha_3R_iD_i \\ +\beta_3R_i^2D_i + (\alpha_2-\alpha_3)R_iD_ia_2(\mathbf{X}_i) + (\beta_2-\beta_3)R_i^2D_ia_2(\mathbf{X}_i) \\ +\alpha_1D_i + \beta_1R_iD_i - \alpha_1R_iD_i - \beta_1R_i^2D_i + \boldsymbol{\gamma}^T\mathbf{X}_i \end{array}\middle| \; R_i, D_i, \mathbf{X}_i \right\}$$

$$= E\left\{\begin{array}{l} \boldsymbol{\gamma}^T\mathbf{X}_i + \alpha_3 + (\alpha_2-\alpha_3+\beta_2)R_i + (\alpha_1-\alpha_3)D_i \\ +(\alpha_3-\alpha_1+\alpha_3-\alpha_2+\beta_3-\beta_2)R_iD_i \\ +(\alpha_1-\alpha_3+\beta_1-\beta_3)R_iD_ia_1(\mathbf{X}_i) \\ +(\alpha_2-\alpha_3)(1-R_i)(1-D_i)a_2(\mathbf{X}_i) \end{array}\middle| \; R_i, D_i, \mathbf{X}_i \right\}. \tag{5}$$

Table 1 on page 1306 shows the correspondence between the coefficients in the regression (5) and the parameters $\{\boldsymbol{\gamma}, \alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3\}$. The parameters $\{\boldsymbol{\gamma}, \alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3\}$ can be solved for as long as we know the coefficients in the regression (5). In order for the coefficients in the regression (5) to be identified, we need that $a_1(\mathbf{X})$ and $a_2(\mathbf{X})$ vary with $\mathbf{X}$; otherwise we will have collinearity and be unable to estimate the regression coefficients.

2.1. Estimation

In this section, we describe how to estimate the parameters $\{\boldsymbol{\gamma}, \alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3\}$ using the observed data.

First, we estimate $a_1(\mathbf{X})$ and $a_2(\mathbf{X})$ by fitting a multinomial logistic regression model with parameters $\boldsymbol{\theta}$ for probability of principal strata ($\mathcal{P}$) given $\mathbf{X}$; i.e., $p_{A|\mathbf{X}}$, $p_{N|\mathbf{X}}$, and $p_{T|\mathbf{X}}$. We can use the EM algorithm to estimate $\boldsymbol{\theta}$. In the E-step, we estimate $\Pr(\mathcal{P} \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)$, where $\widehat{\boldsymbol{\theta}}^*$ are the previous step estimates for $\boldsymbol{\theta}$. In the M-step, we fit a weighted multinomial logistic regression model with augmented data that for each original data point contains observations for each possible principal stratum for that data point and weights equal to $\Pr(\mathcal{P} \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)$ divided by the sum of all $\Pr(\mathcal{P} \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)$ that are consistent with that data point (Ibrahim, 1990) [3]. For example, if $D_i = 1, R_i = 1$, then the possible principal strata are always-high and treatment positively affected. In augmented data, we have two observations for data point i, one with

TABLE 1
Parameters for least squares regression in Equation (5)

Variable	Coefficient
$\mathbf{X}$	$\boldsymbol{\gamma}$
Intercept	α_3
R	$\alpha_2 - \alpha_3 + \beta_2$
D	$\alpha_1 - \alpha_3$
$R \times D$	$\alpha_3 - \alpha_1 + \alpha_3 - \alpha_2 + \beta_3 - \beta_2$
$R \times D \times a_1(\mathbf{X})$	$\alpha_1 - \alpha_3 + \beta_1 - \beta_3$
$(1-R) \times (1-D) \times a_2(\mathbf{X})$	$\alpha_2 - \alpha_3$

$\mathcal{P}$ = always-high (A) and one with $\mathcal{P}$ = treatment positively affected (T), and weights equal to $\Pr(A \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)/\{\Pr(A \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)+\Pr(T \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)\}$ and $\Pr(T \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)/\{\Pr(A \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*) + \Pr(T \mid \mathbf{X}, R, D, \widehat{\boldsymbol{\theta}}^*)\}$ respectively.

Second, to estimate the parameters $\{\boldsymbol{\gamma}, \alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3\}$, we run least squares regression of Y on intercept, $\mathbf{X}$, R, D, $R \times D$, $R \times D \times a_1(\mathbf{X})$, $(1-R) \times (1-D) \times a_2(\mathbf{X})$, plugging in the estimated values of $a_1(\mathbf{X})$ and $a_2(\mathbf{X})$ from the previous step, and then solve for $\{\boldsymbol{\gamma}, \alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3\}$ using Table 1 on page 1306. As noted above, in order to for the regression coefficients on on $\mathbf{X}$, R, D, $R \times D$, $R \times D \times a_1(\mathbf{X})$, $(1-R) \times (1-D) \times a_2(\mathbf{X})$ to be identified, we need that $a_1(\mathbf{X})$ and $a_2(\mathbf{X})$ vary with $\mathbf{X}$; otherwise we will have collinearity.

Based on the estimates of β_1, β_2, β_3, $a_1(\mathbf{X})$ and $a_2(\mathbf{X})$, we can evaluate whether inequality (1) holds; if inequality (1) holds, this suggests that D might be a good tailoring variable for a DTR. Also if the reverse of inequality (1) holds, i.e., $E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 1) > E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 0)$, this also suggests that D might be a good tailoring variable where now $D^{r=1} = 1$ would indicate that the treatment might not be working well and we might want to adjust it. If $E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 1) - E(Y_i^{r=1} - Y_i^{r=0} \mid D_i^{r=1} = 0)$ is close to zero, this suggests that D might not be a good tailoring variable.

2.2. Simulation studies

In this section, we conduct a simulation study to demonstrate the idea of estimating whether a post-treatment variable from a simple randomized clinical trial is an effect modifier, i.e., we estimate the parameters of interest in (4) from the observed data. There are n subjects, $i = 1, \ldots, n$. We consider the following setup for a simulation study. Let X_i be a continuous pre-treatment variable with $X_i \sim U(0,1)$ and $\mathcal{P}_i \in \{A, N, T\}$ be principal strata with $\mathcal{P}_i \sim \text{Multi}(n, \mathbf{p}_i)$ where $\mathbf{p}_i = [p_{A|X_i}, p_{N|X_i}, p_{T|X_i}]^T$ are

$$\begin{cases} p_{A|X_i} = \dfrac{\exp(\theta_{A0} + \theta_{A1} X_i)}{\exp(\theta_{A0} + \theta_{A1} X_i) + \exp(\theta_{N0} + \theta_{N1} X_i) + 1} \\ p_{N|X_i} = \dfrac{\exp(\theta_{N0} + \theta_{N1} X_i)}{\exp(\theta_{A0} + \theta_{A1} X_i) + \exp(\theta_{N0} + \theta_{N1} X_i) + 1} \\ p_{T|X_i} = 1 - p_{A|X_i} - p_{N|X_i} \end{cases}, \tag{6}$$

TABLE 2
Simulation studies for estimating the parameters of interest using observed data from 100 simulated data sets of size 10,000

Parameter	Truth	Mean Estimate	Standard Deviation
α_1	0.75	0.7478	0.0418
α_2	0.5	0.5045	0.1970
α_3	1	0.9982	0.1608
β_1	0.2	0.1857	0.2042
β_2	0.1	0.0943	0.1773
β_3	0.3	0.3190	0.1820
γ	1	1.0016	0.0647
p_A	1/3	0.3325	0.0059
p_N	1/3	0.3305	0.0068
p_T	1/3	0.3370	0.0092
$\frac{\beta_1 p_A + \beta_3 p_T}{p_A + p_T} - \beta_2$	0.15	0.1583	0.2660

and $\{\theta_{A0}, \theta_{A1}, \theta_{N0}, \theta_{N1}\}$ are set to be $\{-0.5, 1, -0.5, 1\}$. We consider a binary treatment assignment R_i following a Bernoulli distribution with probability 0.5. A binary post-treatment variable D_i is determined based on the principal strata $\mathcal{P}_i$ and the treatment assignment R_i such that

$$D_i = \begin{cases} 1 & \text{if } \{\mathcal{P}_i = A\} \text{ or } \{\mathcal{P}_i = T \text{ and } R_i = 1\} \\ 0 & \text{if otherwise} \end{cases}. \tag{7}$$

The potential outcomes $Y_i^{r=0}$ and $Y_i^{r=1}$ are iid random variables such that

$$\begin{cases} Y_i^{r=0} = \alpha_1 I(\mathcal{P}_i = A) + \alpha_2 I(\mathcal{P}_i = N) + \alpha_3 I(\mathcal{P}_i = T) + \gamma X_i + \epsilon_i \\ Y_i^{r=1} = Y_i^{r=0} + \beta_1 I(\mathcal{P}_i = A) + \beta_2 I(\mathcal{P}_i = N) + \beta_3 I(\mathcal{P}_i = T) \end{cases},$$

where $\epsilon_i \sim N(0,1)$ and $\{\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3, \gamma\} = \{0.75, 0.5, 1, 0.2, 0.1, 0.3, 1\}$. The observed outcome is then $Y_i = R_i Y_i^{r=1} + (1 - R_i) Y_i^{r=0}$. We use the observed data $\{Y_i, R_i, D_i, X_i\}$ to estimate the parameters of interest $\{\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3, \gamma\}$. The sample size for each simulated data set is $n = 10,000$. Table 2 on page 1307 shows averaged estimates for the parameters of interest and their standard errors from 100 simulated data sets. The estimates of the parameters are approximately unbiased. D is estimated to be an effect modifier satisfying (1); this suggests that D might be a good tailoring variable for a DTR.

3. Conclusion

Drs. Laber, Lizotte, Qian, Pelham and Murphy have provided a valuable review of DTRs. In our discussion, we have considered the issue raised by Laber et al. of how to choose tailoring variables. We have considered the setting of a simple randomized clinical trial in which at times t past initial treatment assignment,

every patient who has the same status and treatment history receives the treatment at time t that is specified by the protocol. Such trials cannot be used to evaluate DTRs because there is no variation in treatment regimes beyond the initial treatment assignment, but such trials can suggest potential tailoring variables for designing DTRs to be evaluated in a later SMART. We have provided a method for estimating whether a post-treatment variable is an effect modifier; a post-treatment effect modifier might be a good tailoring variable. The key assumptions of our method are that there are pre-treatment covariates $\mathbf{X}$ which interact with the treatment in affecting the post-treatment variable ($a_1(\mathbf{X})$ and $a_2(\mathbf{X})$ vary with $\mathbf{X}$) but do not interact with treatment assignment in affecting the outcome (the association of $\mathbf{X}$ with potential outcomes is the same for each principal stratum), and also that there are no subjects whose post-treatment variable is negatively affected by being treated. In this paper, we have considered one binary post-treatment variable D as a potential tailoring variable; it would be useful to extend the method to consider continuous and multiple post-treatment variables in future work.

Acknowledgements

We are grateful to the Editor for the invitation to discuss the work "Dynamic treatment regimes: technical challenges and applications" by Laber et al. This study was supported by grant SES-1260782 from the Measurement, Methodology and Statistics Program of the U.S. National Science Foundation.

References

[1] Faerber, J. A., Joffe, M. M., Zhang, R., Brown, G. K., Beck, A. T. and Ten Have, T. R. A causal model for post-randomization effect modification of randomized interventions. Working paper.
[2] Frangakis, C. E. and Rubin, D. B. (2002). Principal stratification in causal inference. *Biometrics*, 58, 21–29. MR1891039
[3] Ibrahim, J. G. (1990). Incomplete data in generalized linear models. *Journal of the American Statistical Association*, 85, 765–769.
[4] Laber, E. B., Lizotte, D. J., Qian, M., Pelham, W. E. and Murphy, S. A. (2014). Dynamic treatment regimes: technical challenges and applications. *Electronic Journal of Statistics*, in press.
[5] Small, D. S., Joffe, M. M., Lynch, K. G., Roy, J. A. and Localio, R. (2014). Tom Ten Have's contributions to causal inference and biostatistics: review and future research directions. *Statistics in Medicine*, in press.
[6] Stephens, A., Keele, L. and Joffe, M. M. Estimating posttreatment effect modification with generalized structural mean models. Working paper.
[7] Ten Have, T. R. (2007). The Other Translational Research: Adapting mental health research to the heterogeneity of community (A statisticians view). Award lecture for the 2007 Harvard Award in Psychiatric Epidemiology and Biostatistics.