University of Michigan School of Public Health

The University of Michigan Department of Biostatistics Working
Paper Series

Year 2004

Paper 34

Causal Inference in Hybrid Intervention Trials Involving Treatment Choice

Qi Long* Rod Little[†]

Xihong Lin[‡]

*University of Michigan, qlong@umich.edu

[†]University of Michigan, rlittle@umich.edu

[‡]University of Michigan, xlin@umich.edu

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<http://biostats.bepress.com/umichbiostat/paper34>

Copyright ©2004 by the authors.

Causal Inference in Hybrid Intervention Trials Involving Treatment Choice

Qi Long, Rod Little, and Xihong Lin

Abstract

Randomized allocation of treatments is a cornerstone of experimental design, but has drawbacks when a limited set of individuals are willing to be randomized, or the act of randomization undermines the success of the treatment. Choice-based experimental designs allow a subset of the participants to choose their treatments. We discuss here causal inferences for experimental designs where some participants are randomly allocated to treatments and others receive their treatment preference. This paper was motivated by the “Women Take Pride” (WTP) study (Janevic et al., 2001), a doubly randomized preference trial (DRPT) to assess behavioral interventions for women with heart disease. We propose a model that allows us to estimate the causal effects in the subpopulations defined by treatment preferences and the preference effects for a DRPT, and develop an EM Algorithm to compute maximum likelihood estimates of the model parameters. The method is illustrated by analyzing treatment compliance of the WTP data. Our results show that there were strong preference effects in the WTP study, that is, women assigned to their preferred treatment were more likely to comply. We also expand these methods to handle a broader class of designs, and discuss alternative designs from the perspective of the strength of assumptions required to make causal inferences.

Causal Inference in Hybrid Intervention Trials Involving Treatment Choice

Qi Long, Roderick J. Little, and Xihong Lin*

March 10, 2004

Abstract

Randomized allocation of treatments is a cornerstone of experimental design, but has drawbacks when a limited set of individuals are willing to be randomized, or the act of randomization undermines the success of the treatment. Choice-based experimental designs allow a subset of the participants to choose their treatments. We discuss here causal inferences for experimental designs where some participants are randomly allocated to treatments and others receive their treatment preference. This paper was motivated by the “Women Take Pride” (WTP) study (Janevic et al., 2001), a doubly randomized preference trial (DRPT) to assess behavioral interventions for women with heart disease. We propose a model that allows us to estimate the causal effects in the subpopulations defined by treatment preferences and the preference effects for a DRPT, and develop an EM algorithm to compute maximum likelihood estimates of the model parameters. The method is illustrated by analyzing treatment compliance of the WTP data. Our results show that there were strong preference effects in the WTP study, that is, women assigned to their preferred treatment were more likely to comply. We also expand these methods to handle a broader class of designs, and discuss alternative designs from the perspective of the strength of assumptions required to make causal inferences.

KEY WORDS: Clinical Trials; Doubly Randomized Preference Trials; EM algorithm; Partially Randomized Preference Trials; Randomization; Selection Bias.

*Qi Long is a PhD student, Roderick J. Little is Professor and Xihong Lin is Professor, Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109. This research is supported by the National Cancer Institute grant R01CA76404. We would like to thank Noreen Clark for providing the WTP data and useful comments.

1 Introduction

Randomized assignment of subjects to treatments is a cornerstone of good experimental design. With full compliance and no missing data, randomization allows valid estimates of average treatment effects under minimal assumptions, by avoiding selection bias and ensuring that the observed mean for each treatment is an unbiased estimate of the overall mean if all individuals in the population had received that treatment (e.g., Little and Rubin, 2001).

However, it is also well known that randomization does not solve all problems in experiments involving alternative treatments. It is not always ethically feasible, as in medical trials when the principle of equipoise is not widely accepted, or in assessments of potentially harmful environmental effects that are unlikely to be beneficial. Inferences are only possible for the subset of subjects willing to be randomized, potentially excluding a significant fraction of the population, including subjects with strong treatment preferences. A behavioral treatment, in which strong motivation on behalf of the participant is required and treatment assignment can not be blinded, may be more successful if subjects are allowed to choose, rather than are randomized to a treatment. Random assignment to a treatment perceived to be inferior may lead to issues of noncompliance and missing data that undermine the randomization and complicate causal inferences.

An alternative to randomization is to simply allow participants to choose their treatments. When people are allowed to choose their treatment in a trial, they are more likely to participate and fully comply with the protocol, and the trial may more realistically measure the outcome of the treatment in the population of interest. On the other hand, this approach has all the well known problems of observational studies, in that treatment comparisons are obscured by the confounding effects of selection. Few experimentalists would accept the notion that the advantages of these designs compensate for their major weaknesses.

A natural question is whether there are hybrid designs between the extremes of randomization or choice that improve on either design. One candidate is Zelen's (1990) randomized consent design, which reverses the usual order of consent and randomization, by randomizing prior to consent and predicating the consent process on the randomized treatment. In a study with two alternative treatments (say A and B), participants are randomized to treatment (say A or B). Those randomized to A are asked if they are willing to receive A, after a discussion of the two treatments. If the participant agrees, A is given. If not, then B is given. The same procedure is followed in the other arm, with roles of A and B reversed. Zelen (1990) shows that this design allows for valid tests of the null hypothesis of no treatment effect, and can be more powerful than a randomized design

restricted to participants willing to be randomized. However, the ethics of not describing the initial randomization to the participant have been questioned (Ellenberg, 1992). See Altman et al. (1995) for more discussion of this design.

We consider here other hybrid designs that combine features of randomization and patient choice of treatments (Lambert and Wood, 2000). Perhaps the simplest approach is to simply ask the participants' treatment preferences in the context of a conventional fully randomized trial (Torgerson et al., 1996), and use that information as a covariate or effect-modifier. A more radical approach is the partially randomized preference trial (PRPT), where participants who are willing to be randomized to the treatments are randomized, and those that are not are assigned to their preferred treatment (Brewin and Bradley, 1989). A variation on this theme is the doubly-randomized preference trial (DRPT), where participants are randomized into a "randomization arm", within which treatments are randomized, and a "preference arm", within which participants get to choose their treatments. Versions of the DRPT are described by Rücker (1989), Wennberg et al. (1993) and Janevic et al. (2003). It seems plausible that in PRPT's and DRPT's, the additional information on participants who get to choose their treatments might usefully supplement the information from the participants who are randomized, although as discussed in Section 4 this often requires modeling assumptions. Overall, hybrid designs enable us to estimate the preference effects and causal effects in subpopulations defined by treatment preferences, which can not be estimated in a completely randomized trial.

We consider causal inference for these designs within the framework of potential outcomes, also known as Rubin's Causal model (Holland, 1986). Originally formalized by Neyman (1923) in the context of randomized experiments, this framework was generalized and extended by Rubin (1974, 1977, 1978) to nonrandomized studies. The key underlying idea is that causal estimands are comparisons of the potential outcomes that would have been observed under different exposures of the same units to treatments at a particular place and time. Robins (1986, 1987) extended Rubin's "point treatment" potential outcome framework to evaluate direct and indirect effects of time-varying treatments in experimental and observational longitudinal studies. However, those methods are not directly applicable to the hybrid designs discussed in this paper.

The first objective of this paper is to propose a general model for assessing preference effects on the outcomes of interest. Our second objective is to propose a framework based on recent statistical ideas of causal inference for assessing the hybrid designs described above, and extensions. The basic idea is to classify individuals in the population into strata, which may or may not be observed for participants in the trial, and then assess assumptions required to identify the causal effects of

treatments within these strata. Causal effects are defined as the difference in average outcome if all individuals within a stratum were assigned to treatment A and if all individuals within a stratum were assigned to treatment B (e.g. Rubin, 1974).

The paper was motivated by the “Women Take Pride” study (Janevic et al., 2003), which utilized a DRPT to assess behavioral interventions for women with heart disease. In Section 2 we discuss the design of that study, and show how to estimate causal effects of subpopulations defined by preference treatments and preference effects using the causal inference framework and the method of moments, an analysis similar to that proposed by Rücker (1989) but within the causal inference framework. In Section 3, we propose a more general likelihood-based method that accommodates covariates and estimates the model causal parameters using the EM algorithm. In Section 4, we analyze the WTP data using the method developed in Section 3. In Section 5, we present a simulation study to evaluate the performance of the proposed method. In Section 6 we expand these methods to handle a broader class of designs, and discuss alternative designs from the perspective of the strength of assumptions required to make causal inference. Section 7 presents some concluding remarks.

2 The “Women Take Pride” Study

2.1 The Study

The “Woman Take Pride” (WTP) intervention study (Janevic et al., 2001) concerns women aged 60 years and older with diagnosed cardiac disease being treated by daily heart medication. The interventions are behavioral programs aimed at enhancing the women's ability to manage their disease, based on principles of self-regulation from social cognitive theory. The comparison is between two versions of an intervention consisting of six weekly units and each participant i was assigned to one of these two treatments: a Group format ($A_i = 1$), where 6-8 women meet for 2-2 1/2 hours, and in a Self-directed format ($A_i = 0$) where the participant studies at home following orientation. Motivation and support are provided through the social environment in the Group version, and through weekly telephone calls from the health educator and peer leader in the Self-directed format. These two versions of the intervention present the same material and only differ in format.

A DRPT design where some participants choose their treatment was seen as preferable to a completely randomized design in this setting, since the choice more accurately reflects a clinical situation. In a “real world” setting, patients may well have a preference for one intervention or treatment over another and preference may well impact the motivation and adherence of the participants and hence the success of the treatment. The design is summarized in Figure 1. At the first stage, a total of 3079 women with heart disease were randomized to a Random arm ($W_i = 1$) when treatments were

randomly assigned, and a Choice arm ($W_i = 0$) when participants received their treatment of choice: within the Random arm ($n = 1613$), 575 (35.6%) women agreed to participate; within the Choice arm ($n = 1466$), 496 (33.8%) agreed to participate. At the second stage, the women in the Random arm were randomized to three groups: control ($n = 184$) (for simplicity ignored in our analyses), the Group format ($n = 190$) and the Self-Directed format ($n = 201$); women in the Choice arm were asked to choose between the Group format ($n = 321$) and the Self-Directed format ($n = 175$).

Primary outcomes of the study are measures of improved physical and psychosocial functioning, frequency and severity of symptoms, and health-care utilization measured at baseline, 4, 12 and 18 months. Since data on these variables were available only at baseline at time of writing, we focus on a secondary outcome of interest in understanding treatment adherence, namely Y = treatment compliance, which is defined as a binary variable of whether a woman completed at least one unit of materials; other outcomes will be considered in the future when data become available. Table 1 summarizes the compliance data.

Table 1 shows similar compliance rates (76%) for both treatments in the Random arm, while a higher compliance rate for the Group format ($A_i = 1$) than for the Self-Directed format ($A_i = 0$) in the Choice arm (93% vs 77%), where A_i indicates the treatment assignment. The comparison of compliance rates between treatments in the Random arm is valid because of the randomized allocation. The compliance rates in the Random arm might be expected to be lowered than those in the Choice arm due to the fact that some participants are not randomized to their treatment of choice. The comparison of compliance rates in the Choice arm is potentially biased by the effects of self-selection. In particular, the compliance rate for the Group format ($A_i = 1$) measures compliance to the Group format in the subpopulation that would choose the Group format (denote this subpopulation by $C_i = 1$), and the compliance rate for the Self-Directed format ($A_i = 0$) measures compliance to the Self-Directed format in the subpopulation that would choose the Self-Directed format (denote this subpopulation by $C_i = 0$). Since $C_i = 1$ and $C_i = 0$ are different subpopulations, the comparison of the two compliance rates in the choice arm does not estimate a causal effect in a particular population, which is the key requirement for a causal effect in Rubin's (1974) sense. A direct comparison of these two compliance rates requires the very debatable assumption that $C_i = 1$ and $C_i = 0$ are equivalent with respect to treatment outcomes. This assumption might be improved by regression adjustment for known characteristics of participants in the two groups, but as in any observational setting, such adjustments do not necessarily remove the bias.

2.2 Causal Inference Using the Method of Moments

A causal analysis that addresses the issue of choice is to compare the compliance rates of being assigned to two different treatments, i.e., the Group format ($A_i = 1$) and the Self-Directed format ($A_i = 0$), within each of the two preference subpopulations, i.e., the subpopulation of women who prefer the Group format ($C_i = 1$) and the subpopulation of women who prefer the Self-Directed format ($C_i = 0$). This question is not addressable from data in the Choice arm alone, because it requires compliance information for participants who do not receive their treatment of choice. However, it can be addressed from the DRPT, since some participants in the Randomized arm do not receive their treatment of choice, and treatment assignment remains random within the two preference $C_i = 1$ and $C_i = 0$ subpopulations.

More specifically, we partition the the population into two subpopulations, people who prefer the Group Format ($C_i = 1$) and people who prefer the Self-Directed format ($C_i = 0$). In the choice arm, the subpopulations are observed. In the random arm, each group represents the whole population and is a mixture of the two subpopulations. Define

- $\mu(0)$ =overall compliance rate if assigned to the Self-Directed format ($A_i = 0$),
- $\mu_0(0)$ =compliance rate if assigned to the Self-Directed format ($A_i = 0$) in the subpopulation that prefer the Self-Directed format ($C_i = 0$)
- $\mu_1(0)$ =compliance rate if assigned to the Self-Directed format ($A_i = 0$) in the subpopulation that prefer the Group format ($C_i = 1$),

and define $\mu(1)$, $\mu_0(1)$ and $\mu_1(1)$ as the corresponding compliance rates if assigned to the Group format ($A_i = 1$), in the overall and the two subpopulations respectively. Let π_1 be the proportion of the population that prefers the Group Format ($C_i = 1$). Then

$$\begin{aligned}\mu(0) &= \pi_1\mu_1(0) + (1 - \pi_1)\mu_0(0) \\ \mu(1) &= \pi_1\mu_1(1) + (1 - \pi_1)\mu_0(1).\end{aligned}$$

From the choice arm, we estimate $\hat{\mu}_1(1) = 299/321 = 0.93$, $\hat{\mu}_0(0) = 134/175 = 0.77$, and $\hat{\pi}_1 = 321/496 = 0.65$. From the random arm, we estimate $\hat{\mu}(0) = 0.76$ and $\hat{\mu}(1) = 0.76$. Thus, $\hat{\mu}_1(0) = (0.76 - 0.35 \times 0.77)/0.65 = 0.75$, and $\hat{\mu}_0(1) = (0.76 - 0.65 \times 0.93)/0.35 = 0.44$.

Denote by θ_1 the causal effect of the treatments of being assigned to the Group format ($A_i = 1$) compared to being assigned to the Self-Directed format ($A_i = 0$) on compliance in the subpopulation

that prefers the Group format ($C_i = 1$). The forgoing calculations suggest that θ_1 can be estimated as

$$\hat{\theta}_1 = \hat{\mu}_1(1) - \hat{\mu}_1(0) = 0.93 - 0.75 = 0.18,$$

This means for women who prefer the Group format, the compliance rate when assigned to the Group format is estimated to be 18% higher than the compliance rate when assigned to the Self-Directed format. Denote by θ_0 the causal effect of the treatments ($A_i = 1$ vs $A_i = 0$) on compliance in the subpopulation that prefer the Self-Directed format ($C_i = 0$). Then θ_0 can be estimated by

$$\hat{\theta}_0 = \hat{\mu}_0(1) - \hat{\mu}_0(0) = 0.44 - 0.77 = -0.33.$$

This means for women who prefer the Self-Directed format, the compliance rate when assigned to the Self-Directed format is estimated to be 33% higher than the compliance rate when assigned to the Group format. These results indicate that women are more likely to comply to the treatment they prefer. The preference effect, defined as the treatment effect difference between the two subpopulations, is thus estimated as $\hat{\theta}_1 - \hat{\theta}_0 = 0.51$.

The standard errors of $\hat{\theta}_1$, $\hat{\theta}_0$ and $\hat{\theta}_1 - \hat{\theta}_0$ are estimated as 0.05, 0.10 and 0.12, respectively, computed using the observed information matrix as discussed in the next section. It follows that the treatment effects on compliance are highly significant in the two subpopulations and they are significantly different. The results suggest that the very similar compliance rates for the two intervention groups in the Random arm mask strong preference effects, with much higher compliance rates for the preferred treatments. This analysis method is a more formal description of the method of Rücker (1989) within the causal inference framework, applied to our data set. It does not require a distributional assumption and is applicable to estimating causal effects of an arbitrary outcome, e.g., the causal effect of a health outcome at a post-intervention time point such as 12 month. In the next section we describe the underlying assumptions in this analysis and generalize the analysis to include covariates.

3 A General Model for a DRPT

Janevic, et al. (2003) found that the probability of choosing each intervention format was affected by demographical variables and disease severity. The compliance probabilities within the two subpopulations ($C_i = 1$ and $C_i = 0$) are also likely to be affected by covariates besides the treatments. We hence extend the analysis of the previous section to accommodate covariates. We describe formally in this section the implicit assumptions in the analysis assumed in the previous assumption. As before, we consider a DRPT with two treatments, generically denoted as 1 and 0.

3.1 The Assumptions

The method of moments in the absence of covariates discussed in Section 2 and the maximum likelihood method in the presence of covariates discussed in Sections 3.2 and 3.3 are based on the following assumptions:

1. Stable Unit Treatment Value Assumption (SUTVA) (Rubin, 1978): The potential outcomes for each individual do not depend on the treatment status of other individuals in the sample.
2. No Selection Bias from Randomization (NSBR): The random arm participants and the choice arm participants are random samples of the same population.
3. “Exclusion Restriction” (ER) Assumption: The outcome for an individual randomly assigned to treatment m ($m = 0, 1$) is the same as if that individual had chosen treatment m ($m = 0, 1$). We follow other authors by calling this the “exclusion restriction” (ER), since it is an example of an exclusion restriction in the sense that the term is used in econometrics (Angrist and Rubin, 1996).

We will discuss in Section 4 how these assumptions might be relaxed.

3.2 The Model

Suppose the data are comprised of n subjects. For subject i , let Y_i denote the observed outcome of interest, A_i denote the treatment assignment (1/0), C_i denote the treatment preference (1/0), $W_i = 1$ if subject i is randomized to the random arm and $W_i = 0$ if subject i is randomized to the choice arm, X_{1i} be a set of covariates associated with Y_i , and X_{2i} be a set of covariates associated with C_i , where X_{1i} and X_{2i} may overlap.

Let $Y_i(1)$ be the potential outcome of interest for subject i , when assigned treatment 1, and $Y_i(0)$ be the potential outcome of interest for subject i , when assigned treatment 0. The average causal effect of treatment assignment for the whole population is

$$\theta = E\{Y_i(1)\} - E\{Y_i(0)\},$$

The average causal effect of treatment assignment for the subpopulation preferring treatment m , ($m = 0, 1$) and having covariates X_1 is

$$\theta_m(X_1) = E\{Y_i(1)|C_i = m, X_1\} - E\{Y_i(0)|C_i = m, X_1\}.$$

Averaging over the distribution of X_1 , the average causal effect of treatment assignment for the subpopulation preferring treatment m equals to

$$\theta_m = E_{X_1|C_i=m} [E\{Y_i(1)|C_i = m, X_1\} - E\{Y_i(0)|C_i = m, X_1\}].$$

Note that these causal parameters $\theta_m(X_1)$ and θ_m depend on potential outcomes. We hence need to relate these conditional means that involve potential outcomes $Y_i(1)$ and $Y_i(0)$ to their counterparts that can be estimated using the observed data Y_i . This can be done under the assumptions stated in Section 3.1. Specifically, since subjects are randomized to the Choice arm and the Random arm, we have

$$E\{Y_i(j)|C_i = m, X_1\} = E\{Y_i(j)|C_i = m, X_1, W_i\},$$

where $m = 0, 1$ and $j = 0, 1$. In the Random arm, since subjects are randomized to treatment groups, we have

$$\begin{aligned} E\{Y_i(j)|C_i = m, X_1, W_i = 1\} &= E\{Y_i(j)|A_i = j, C_i = m, X_1, W_i = 1\} \\ &= E(Y_i|A_i = j, C_i = m, X_1, W_i = 1). \end{aligned}$$

In the Choice arm, subjects are assigned to treatments they prefer, that is, $A_i = C_i$. Thus for $j = 0, 1$,

$$\begin{aligned} E\{Y_i(j)|C_i = j, X_1, W_i = 0\} &= E\{Y_i(j)|A_i = C_i, C_i = j, X_1, W_i = 0\} \\ &= E(Y_i|A_i = j, C_i = j, X_1, W_i = 0). \end{aligned}$$

However, we cannot estimate $E\{Y_i(0)|C_i = 1, X_1, W_i = 0\}$ and $E\{Y_i(1)|C_i = 0, X_1, W_i = 0\}$ only using data from the Choice arm. Under the NSBR and ER Assumptions,

$$\begin{aligned} E\{Y_i(0)|C_i = 1, X_1, W_i = 0\} &= E\{Y_i(0)|A_i = 0, C_i = 1, X_1, W_i = 1\} = E\{Y_i|A_i = 0, C_i = 1, X_1, W_i = 1\} \\ E\{Y_i(1)|C_i = 0, X_1, W_i = 0\} &= E\{Y_i(1)|A_i = 1, C_i = 0, X_1, W_i = 1\} = E\{Y_i|A_i = 1, C_i = 0, X_1, W_i = 1\} \end{aligned}$$

We hence can use the data in the random arm in conjunction with the data in the choice arm to estimate these quantities by viewing the each group in the random arm is a mixture of two preference subpopulations.

Specifically, we assume that the distribution of Y_i given A_i , C_i , X_{1i} and W_i follows the exponential family

$$f(Y_i|A_i, C_i, X_{1i}, W_i) = \exp \left\{ \frac{Y_i \gamma_i - b(\gamma_i)}{\phi a_i^{-1}} + c(Y_i, \phi) \right\},$$

where a_i is a known constant, ϕ is a scale parameter, γ_i is the canonical parameter, $b(\cdot)$ and $c(\cdot)$ are known functions. The mean of Y_i is $\mu_i = E(Y_i|A_i, C_i, X_{1i}, W_i) = b'(\gamma_i)$ and follows a generalized linear model

$$g(\mu_i) = \beta_0 + X_{1i}^T \beta_{X_1} + A_i \beta_A + C_i \beta_P + A_i C_i \beta_{AP}, \quad (1)$$

where $g(\cdot)$ is a monotone differential link function (McCullagh and Nelder, 1989). The model is completed by assuming the treatment preference C_i given X_{2i} follows a logistic model with mean $\pi_i = E(C_i|X_{2i})$ satisfying

$$\text{logit}(\pi_i) = \alpha_0 + X_{2i}^T \alpha_{X_2}. \quad (2)$$

The causal effect of treatment assignment for the subpopulation preferring treatment m ($m = 0, 1$) given covariates X_1 is

$$\begin{aligned} \theta_m(X_1) &= E\{Y(1)|C = m, X_1\} - E\{Y(0)|C = m, X_1\} \\ &= g^{-1}(\beta_0 + X_1 \beta_{X_1} + \beta_A + m \beta_P + m \beta_{AP}) - g^{-1}(\beta_0 + X_1 \beta_{X_1} + m \beta_P) \end{aligned}$$

The marginal causal effects, θ_m , can be estimated using

$$\theta_m = \int \theta_m(X_1 = x_1) f(x_1|C = m) dx_1, \quad (3)$$

where $f(x_1|C = m)$ can be empirically estimated from the Choice arm, that is, $\hat{f}(x_1|C = m) = \sum_i I(C_i = m, X_{1i} = x_1) / \sum_i I(C_i = m)$, where $I(\cdot)$ is an indicator function. Alternative marginal causal effects θ_m^* calculated by standardizing using the distribution of X_1 in the overall population can be estimated by using $f(x_1)$ in place of $f(x_1|C = m)$ in (3), where $f(x_1)$ can be estimated from the data in the random arm.

In both the random arm and the choice arm, (Y_i, W_i) are observed. In the random arm, A_i is observed due to randomization but C_i is not observed. In the choice arm, both A_i and C_i are observed and $A_i = C_i$, since subjects are assigned to their preferred treatment. Denote by $\delta = (\alpha, \beta, \phi)$ the parameter vector. Define $Y = (Y_1, Y_2, \dots, Y_n)^T$, and C, X_1, X_2, W similarly. Noticing that each group in the random arm is a mixture of the two preference subpopulations, the observed data loglikelihood is given by

$$\begin{aligned} \ell(Y, C_{obs}|A, X_1, X_2, W; \theta) &= \sum_{W_i=0} [\log \{f(Y_i|A_i, X_{1i}, C_i)\} + C_i \log(\pi_i) + (1 - C_i) \log(1 - \pi_i)] \\ &\quad + \sum_{W_i=1} [\log \{\pi_i f(Y_i|A_i, X_{1i}, C_i = 1) + (1 - \pi_i) f(Y_i|A_i, X_{1i}, C_i = 0)\}], \end{aligned} \quad (4)$$

where C_{obs} denote the observed C values for the Choice arm, and $f(Y_i|A_i, X_{1i}, C_i)$ follows the generalized linear model (1) and π_i follows the logistic model (2).

3.3 Estimation Using the EM Algorithm

In this section, we develop an EM to calculate the maximum likelihood estimate of δ in the model proposed in Section 3.2. The complete data are $(Y_i, C_i, A_i, X_{1i}, X_{2i}, W_i)$ and C_i is missing in the random arm. If the C_i were observed for all subjects, the complete-data loglikelihood is

$$\begin{aligned}\ell(Y, C|A, X_1, X_2, W; \delta) &= \sum_i \ell(Y_i, C_i|A_i, X_{1i}, X_{2i}, W_i; \delta) \\ &= \sum_i [\log \{f(Y_i|A_i, X_{1i}, X_{2i}, W_i)\} + C_i \log(\pi_i) + (1 - C_i) \log(1 - \pi_i)].\end{aligned}\quad (5)$$

The EM algorithm iterates between an E step, which can be viewed here as imputing the preference probabilities in the random arm, and an M step, which maximizes the expected complete-data loglikelihood.

1. *E Step at the k th iteration.* We calculate the expected complete-data loglikelihood given Y, C_{obs}, A, X, Z, W and the current parameter estimates $\delta^{(k)}$, namely

$$\begin{aligned}& E\{\ell(Y, C|A, X_1, X_2, W)|Y, C_{obs}, A, X_1, X_2, W; \delta^{(k)}\} \\ &= \sum_{W_i=1} \sum_{m=0}^1 w_{i,m}^{(k)} \ell(Y_i, C_i = m|A_i, X_{1i}, X_{2i}; \delta^{(k)}) + \sum_{W_i=0} \ell(Y_i, C_i|A_i, X_{1i}, X_{2i}; \delta^{(k)}),\end{aligned}\quad (6)$$

where $m = 0, 1$, and for the random arm ($W_i = 1$),

$$\begin{aligned}w_{i,m}^{(k)} &= p(C_i = m|Y_i, A_i, X_i; \delta^{(k)}) \\ &= \frac{f(Y_i|A_i, X_{1i}, X_{2i}, C_i = m; \delta^{(k)})p(C_i = m|A_i, X_{1i}, X_{2i}; \delta^{(k)})}{\pi_i f(Y_i|A_i, X_{1i}, X_{2i}, C_i = 1; \delta^{(k)}) + (1 - \pi_i) f(Y_i|A_i, X_{1i}, X_{2i}, C_i = 0; \delta^{(k)})}.\end{aligned}$$

The E-step estimates the weights $w_{i,m}^{(k)}$ in the weighted complete-data log-likelihood (5) for the random arm.

2. *M Step at the k th iteration.* This step updates the parameter estimates $\delta^{(k+1)}$ by maximizing the expected complete-data loglikelihood (5). We first construct an augmented data set as follows. Each observation in the random arm, whose C_i is missing, is replaced by two “filled-in” observations, in which the missing treatment preference indicator is replaced by 0 and 1 respectively and the corresponding weights $w_{i,0}$ and $w_{i,1}$ are computed using the current estimates of the parameters. The observations in the choice arm are kept the same and the weights are assigned to be one. Using this augmented data set, we fit a weighted generalized linear model for $\beta^{(k+1)}$ and $\phi^{(k+1)}$, and a weighted logistic model for $\alpha^{(k+1)}$.

The resulting estimates from the EM algorithm at convergence give the MLEs of δ . For the special case of no covariates X_{1i}, X_{2i} , the ML estimates of β transformed to the mean scale reduce to the method of moments estimates in Section 2.2. A similar EM algorithm was proposed by Ibrahim (1990) for missing categorical covariates. Our EM algorithm extends his algorithm to the hybrid randomized-preference design by estimating the weights in the random arm and fixing the weights to be one in the choice arm. We consider estimating standard errors of the maximum likelihood estimates using bootstrap, the observed information and the approximation proposed Ibrahim (1990). The observed information is obtained by directly computing the second derivative of the observed likelihood using a symbolic differentiation algorithm.

4 Analysis of the WTP data

We reanalyze the WTP data described in Section 2.2 by incorporating the covariates using the method discussed in Section 3. In particular, we are interested in estimating the causal effect of treatment assignment on compliance rate within each treatment preference subpopulation after adjusting for covariates. In view of the results of Janevic, et al. (2003), we consider the following covariates in models (1) and (2): employment status, age, total symptom impact, which is a measure of symptom severity scored between 0 and 70 (Clark et al., 1997), and Sickness impact profile (SIP) physical dimension score, which is a measure of physical functional health status scored between 0 and 100 (Bergner et al., 1981).

Table 2 presents the estimates of the regression parameters in model (1) and model (2) and their estimated standard errors. We first discuss the covariate effects on treatment preference and compliance. Our results from model (2) show that employed women were more likely to choose the Self-Directed format (OR=1.99, $p=0.010$); women with a higher total symptom impact score at baseline were less likely to choose the Self-Directed format (OR=0.98, $p=0.028$); and women having greater physical limitations due to illness at baseline were more likely to choose the Self-directed format (OR=1.03, $p=0.007$). The magnitudes of these effects on treatment preference are comparable to those in Janevic et al. (2003). In terms of the covariate effects on compliance, our results show that employed women were less likely to comply than non-employed women (OR=0.36, $P<0.001$). The other covariates were found to have no significant effects on compliance.

We now discuss the causal effect of treatment assignment within each treatment preference subpopulation. The results in Table 2 show that accounting for covariate effects, women who preferred the Group format were more likely to comply if assigned to the Group format compared to if assigned to the Self-Directed format (OR=5.28, $p < 0.001$), while women who preferred the Self-Directed for-

mat were less likely to comply if assigned to the Group format compared to if assigned to the Self-Directed format (OR=0.22, $p < 0.001$). In other words, subjects were more likely to comply when assigned to their preferred treatment. The highly significant interaction between treatment preference and treatment assignment shows the treatment preference strongly modifies compliance to treatment assignment ($p < 0.001$) and indicates a strong preference effect.

Table 3 shows the estimates of the marginal compliance probabilities and their estimated standard errors by integrating over the distributions of the covariates using equation (3). The marginal causal treatment effect on compliance on the probability scale for the subpopulation preferring the Group format is $\hat{\theta}_1 = 0.929 - 0.725 = 0.204$, and the marginal causal treatment effect on compliance in the probability scale for the subpopulation preferring the Self-Directed format is $\hat{\theta}_0 = 0.424 - 0.771 = -0.347$. These results show that women who preferred the Group format were 20.4% more likely to comply if assigned to the Group format compared if assigned to the Self-Directed form, whereas women who preferred the Self-Directed format were 34.7% more likely to comply if assigned to the Self-Directed format compared to if assigned to the Group format. The SEs computed using the bootstrap method were similar to those computed using the observed information. The results using two different standardizations were also similar. Compared to the the no-covariate adjusted analysis discussed in Section 2.2, the covariate-adjusted treatment effects were slightly stronger.

5 A Simulation Study

We conducted a simulation study to evaluate the finite sample performance of the proposed method. The design of the simulation study was similar to that of the WTP study. Each data set consisted of 1000 observations. Independent binary observations C_i of the treatment preference indicator were generated using the logistic model

$$\text{logit}\{E(C_i)\} = \alpha_0 + X_{2i}\alpha_1,$$

where $\alpha_1 = -1$, $\alpha_1 = 2$, and the X_{2i} were generated from a uniform distribution on the interval $[0, 1]$. The outcome variable Y_i was assumed to be binary and generated independently using the logistic model

$$\text{logit}\{E(y_i)\} = \beta_0 + X_{1i}\beta_1 + A_i\beta_A + C_i\beta_P + A_iC_i\beta_{AP},$$

where $\beta_0 = -2$, $\beta_1 = \beta_A = \beta_P = 2$, $\beta_{AP} = -2$, the X_{1i} were generated from a uniform distribution on the interval $[0, 1]$. The arm indicators W_i for the random arm versus the choice arm were generated from a Bernoulli distribution with $P(W_i = 1) = 0.5$. The treatment assignment indicators A_i were set to the same as C_i for the choice arm ($W_i = 0$ arm), and were generated from a Bernoulli distribution

with $P(A_i = 1) = 0.5$ for the random arm ($W_i = 1$). The preference indicators C_i were set to be missing in the random arm ($W_i = 1$).

A total of 125 simulated data sets were generated and analyzed. Table 4 presents the simulation results. The point estimates were very close to the true values and there was little bias. We compared three methods for estimating the standard errors: the bootstrap method, the observed information, and the approximation given by Ibrahim's (1990). These estimated standard errors were compared with the empirical standard errors. Our results show that the observed information based standard errors were closest to the empirical standard errors. The bootstrap standard errors performed similarly to the observed information based standard errors except that they slightly overestimated the standard errors for the coefficients of treatment preference and interaction between treatment preference and treatment assignment. The approximation given by Ibrahim seemed to slightly underestimate the standard errors. This might be due to the fact that the estimated regression parameters β and α in model (1) and (2) were assumed to be independent in the approximation.

6 A General Framework for Hybrid Trials: Taxonomy and Principal Strata

The NSBR and the ER assumption discussed in Section 3.1, which underly the analysis of Sections 3 and 4, are open to question. It seems likely that an individual willing to participate in the random arm would also be willing to participate in the choice arm, and we make this assumption in what follows; but the reverse may not be the case, since some individuals willing to participate if given their choice of treatment might be unwilling to agree to being randomized, particularly if their treatment preference is strong. The fact that participation rates in the two arms are quite similar in the WTP study (Figure 1) suggests that the NSBR assumption is plausible in that case. ER is an assumption since it is possible that a treatment might be more effective if it is a conscious choice of the participant rather than randomly assigned, so the outcome differs under these two scenarios. This assumption is not capable of empirical verification in the DRPT design.

We now consider an expanded framework that makes these two assumptions more explicit, and consider alternatives to the DRPT design in terms of assumptions needed to identify the parameters. Some of the alternative designs do not require the NSBR and ER assumptions to estimate the causal effects of interest. As a consequence, one could check the NSBR and ER assumptions using these designs. We consider for simplicity trials with two treatments A and B, although the framework extends to more than two treatments in an obvious way.

We first stratify the target population into five groups (Figure 2(a)):

1. The set of individuals unwilling to participate even if given their choice of treatments ($\overline{P}$). Clearly we cannot learn anything about treatment effects for this group without making assumptions that relate it to a group we can study. We do not consider this group further here.
2. The set of individuals willing to participate if given the choice of treatment (P). We stratify this group into four subpopulations:
 - (a) The subpopulation that prefers A and will not participate unless allowed to choose A ($P\overline{R}A$).
 - (b) The subpopulation that prefers A but is nevertheless willing to participate in a randomized trial (PRA).
 - (c) The subpopulation that prefers B but are nevertheless willing to participate in a randomized trial (PRB).
 - (d) The subpopulation that prefers B and will not participate unless allowed to choose B ($P\overline{R}B$).

These groups are ordered in an approximate sense by degree of preference for A over B, in that the individuals unwilling to be randomized tend to be those with strong preferences for a particular treatment. However there may be some who oppose randomization for reasons other than a treatment preference, so the ordering is not precise.

For each of four subpopulations ($P\overline{R}A$), (PRA), (PRB), ($P\overline{R}B$), we consider two versions of treatment A, a version where an individual gets to choose A ($\mathcal{A}_C$) and a version where an individual is randomized to A ($\mathcal{A}_R$). Similarly for treatment B we define two versions, $\mathcal{B}_C$ and $\mathcal{B}_R$. Thus under ER, $\mathcal{A}_C = \mathcal{A}_R$ and $\mathcal{B}_C = \mathcal{B}_R$, but in principle we allow the two versions of each treatment to differ.

The combinations of stratum and treatment are displayed in Figure 2(b). Here and in later tables, the cells marked by a bold $\overline{F}$ are a-priori counterfactuals (Angrist, Imbens and Rubin, 1996) (AIR), in the sense that we cannot observe an outcome in these cells under any design. For example, we do not get to see the effect of randomizing to A ($\mathcal{A}_R$) in the subpopulation of individuals who are prefer A but will not participate in a randomized trial ($P\overline{R}A$). Opinions differ on the extent to which it is meaningful to consider outcomes in such cells, and this needs to be considered in the specific context of each trial. In our discussion we follow AIR and focus attention on outcomes that are measurable under some design, that is, the cells without $\overline{F}$'s in Figure 2(b). Causal comparisons (in Rubin's

sense) are comparisons of column effects within rows of Figure 2(b). Comparisons between different rows are not causal since they concern different subpopulations.

We use the framework of Figure 2 to discuss five alternative designs; PRPT and DRPT designs, which have been previously proposed in the literature: versions of PRPT and DRPT with initial questions about treatment preference (PRPTQ, DRPTQ), and a composite PRPT/DRPT design with a question about treatment preference (CRPTQ).

In a partially randomized preference trial (PRPT) (Figure 3(a)), individuals willing to be randomized are assigned to the random arm and receive $\mathcal{A}_R$ or $\mathcal{B}_R$, and individuals not willing to be randomized are allowed to choose the treatment, $\mathcal{A}_C$ or $\mathcal{B}_C$. Thus the PR individuals are randomized to $\mathcal{A}_R$ and $\mathcal{B}_R$, the $P\bar{R}A$ individuals receive $\mathcal{A}_C$ and the $P\bar{R}B$ individuals receive $\mathcal{B}_C$. Figure 3(b) indicates which of the cells in Figure 2(b) can be directly estimated under this design (E), which of the cells are a-priori counterfactual ($\bar{F}$), and which of the cells are not a-priori counterfactual but cannot be estimated without modeling assumptions, since outcomes for individuals in these cells are not directly observed ($\bar{E}$). The randomized part of the design allows comparisons of $\mathcal{A}_R$ and $\mathcal{B}_R$ for the population PR willing to be randomized, but does not distinguish preference effects for this population. The choice arm provides estimates for the outcome under $\mathcal{A}_C$ for $P\bar{R}A$ and the outcome under $\mathcal{B}_C$ for $P\bar{R}B$, but additional modeling assumptions are needed for this information to yield estimates of causal effects within the $P\bar{R}A$ and $P\bar{R}B$ rows of the table. Under the exclusion restriction, $\mathcal{A}_C = \mathcal{A}_R$ and $\mathcal{B}_C = \mathcal{B}_R$, and we can compare the outcomes of A and B in $PRA \cup PRB$, and estimate the outcomes of A in $P\bar{R}A$ and B in $P\bar{R}B$. However to make causal comparisons of the effect of A and B in $P\bar{R}A$ and $P\bar{R}B$, we need to assume that these subpopulations are equivalent. This assumption can be weakened by regression adjustment on observed covariates, but remains strong and not directly testable.

Suppose, however, that after adjustment for covariates, the average outcomes for treatment A were similar in $P\bar{R}A$ and $PRA \cup PRB$ (e.g. a test of the null hypothesis that the means are equal is not significant). Also, the average outcomes for treatment B were similar in $P\bar{R}B$ and $PRA \cup PRB$. Then one could argue that after covariate adjustment, $P\bar{R}A$ and $P\bar{R}B$ are homogeneous with respect to the outcomes of interest, and therefore the comparison of A in $P\bar{R}A$ with B in $P\bar{R}B$ is plausibly causal. The strata might be collapsed to provide a more precise estimate of the treatment effect. This kind of analysis would not be possible in a strictly randomized design that does not include the choice arm.

In a Doubly Randomized Preference Trial (DRPT) (Figure 4(a)), people willing to participate (P) are first randomized to a choice arm and a random arm. Within the choice arm, people receive

their treatment of choice (PA or PB). Within the random arm, individuals willing to be randomized (R) are randomized to $\mathcal{A}_R$ or $\mathcal{B}_R$, and those not willing to be randomized do not participate. The four cells of Figure 2 that are estimated directly under this design are shown with an E in Figure 4(b); note that these cells are all combinations of pairs of strata (rows). Under the NSBR and ER assumptions, this table collapses to that in Figure 4(c), and the effects of A and B are computed separately in the PRA and PRB subpopulations as discussed in the previous section. Thus unlike the PRPT design, the DRPT design does allow the estimation of preference effects under NSBR and ER assumptions.

A simple way of estimating preference effects is to ask a direct question about treatment preference, and use the answers to that question to classify people directly into preference subpopulations. An important issue with this approach is whether it is feasible to combine such a question with randomization of treatments, and in particular whether asking the participant to answer a question about preference undermines their willingness to agree to randomization. However, this approach has been found feasible in some randomized trials (Torgerson et al., 1996).

Figure 5 describes a PRPT design with a preliminary question about treatment preference, which we label PRPTQ, and shows the set of directly estimable effects for this design. The randomized arm does allow us to estimate causal effects of treatments within each treatment preference group, but the choice components of the design do not allow us to estimate causal effects without assumptions of homogeneity across strata. Figure 6 describes a DRPT design with a question about treatment preference, which we label DRPTQ, and shows the corresponding set of directly estimable effects. As with PRPTQ, the randomized arm allows us to estimate causal effects of treatments within each treatment preference group; the choice arm now provides additional information under one of the NSBR or ER assumptions. Specifically if NSBR is assumed then the ER can be checked, since the treatment effects can be estimated for both the randomized and chosen versions of the treatments. On the other hand if the ER is assumed, then the NSBR assumption can be assessed, since treatment effects can be estimated in the strata both willing and unwilling to be randomized. So our analysis favors DRPTQ over PRPTQ. Even with no NSBR or ER assumption, useful bounding information might be derived by assuming that the outcome under $\mathcal{A}_C(\mathcal{B}_C)$ is at least as good as the outcome under $\mathcal{A}_R(\mathcal{B}_R)$.

Figure 7 shows a complex design that includes both a preliminary question about preferences and combines the features of the PRPT and DRPT designs. We call this a combined randomized preference trial with a question about preference (CRPTQ). The corresponding set of estimable effects, shown in Figure 7(b), includes all the effects that are not a-priori counterfactual. In this case

all four versions of the treatment ($\mathcal{A}_R$, $\mathcal{A}_C$, $\mathcal{B}_R$, $\mathcal{B}_C$) can be compared for the subpopulations who prefer A and prefer B, and are willing to be randomized; the outcomes of the treatment can also be compared in the three subpopulations where they are not a-priori counterfactual. The NSBR and ER assumptions can be checked using this design. A further elaboration of the initial question is to ask the question using a Likert scale: SA = strongly prefer A; A = prefer A; N = neutral; B = prefer B, SB = strongly prefer B. This more detailed classification may provide more information for modeling treatment effects in the different principal strata.

7 Conclusions and Discussion

Although classical randomized trials enjoy many advantages, it is well recognized that they are subject to several limitations. Efforts to improve classical randomized trials are encouraged. Doubly Randomized Preference Trials provide an alternative to classical randomized trials to examine how treatment preference can influence treatment effects. However, the naive approach of comparing outcomes under two treatments is flawed due to selection bias in the choice arm in such designs. In this article, we have proposed a two stage model for estimating the causal treatment effects and the preference effect in a DRPT for discrete and continuous outcomes, with the preference indicator missing in the random arm. An EM algorithm has been used for parameter estimation. The analysis can be readily generalized to model other types of outcomes.

The analysis of alternative designs in Section 6 suggests an advantage of DRPTs over PRPTs, since the overlap of subpopulations in the DRPT allows causal effects to be estimated under certain assumptions. Our analysis also demonstrates the utility of designs that ask directly about treatment preference, particularly if the NSBR or ER assumptions are questionable. The utility of designs with questions about treatment preference needs to be weighed against the potential drawbacks of asking about treatment preference in the context of randomization. Pilot studies might be helpful in casting light on this issue.

The framework in Figure 2 is useful since it distinguishes effects that cannot be estimated by any design without modeling assumptions (comparisons involving a-priori counterfactuals) and effects that are estimable under some designs but not under others. Although modeling assumptions are essential for causal inferences involving a-priori counterfactuals, sometimes the more complex designs allow us to combine information across rows of Figure 2 with less stringent assumptions than would be required by simpler designs, as discussed in Section 3. More complex designs involve an administrative overhead, but follow W.G. Cochran's principle of providing the opportunity to learn more by making hypotheses more complex (Rubin, 1984). Designs that combine randomization and choice seem to us

a fruitful area for future research and application, and the ideas of causal inference and population stratification are helpful for guiding the design and resulting analysis.

References

- Altman, D. G., Whitehead, J., Parmar, M. K. B., Stenning, S. P., Fayers, P. M., and Machin, D. (1995), "Randomised Consent Designs in Cancer Clinical Trials," *European Journal of Cancer*, 31A (12), 1934-1944.
- Angrist, J. D., Imbens, G. W., and Rubin, D. B. (1996), "Identification of Causal Effects Using Instrumental Variables (Disc: p456-472)", *Journal of the American Statistical Association*, 91, 444-455.
- Bergner, M., Bobbitt, R. A., Carter, W. B., and Gilson, B. S. (1981), "The Sickness Impact Profile: Development and Final Revision of a Health Status Measure," *Medical Care*, 19,787-805.
- Brewin, C. R., and Bradley, C. (1989), "Patient Preferences and Randomised Clinical Trials," *British Medical Journal*, 299, 313-315.
- Clark, N. M., Janz, N. K., Dodge, J. A., Schork, M. A., Wheeler, J. R. C., Liang, J., Keteyian, S. J., and Santinga, J. T. (1997), "Self-management of Heart Disease by Older Adults: Assessment of an Intervention Based on Social Cognitive Theory," *Research on Aging*, 19, 362-382.
- Ellenberg, S. S., Finkelstein, D. M., and Schoenfeld, D. A. (1992), "Statistical Issues Arising in AIDS Clinical Trials," *Journal of the American Statistical Association*, 87, 562-569.
- Holland, P. W. (1986), "Statistics and Causal Inference(with discussion)," *Journal of the American Statistical Association*, 81, 945-970.
- Ibrahim, J. G. (1990), "Incomplete Data in Generalized Linear Models," *Journal of the American Statistical Association*, 85, 765-769.
- Janevic, M. R., Janz, N. K., Lin, X., Pan, W., Sinco, B. R., and Clark, N. M. (2003), "The Role of Choice in Health Education Interventional Trials: A Review and Case Study," *Social Science and Medicine*, 56(7), 1581-1594.
- Lambert, M. F., and Wood, J. (2000), "Incorporating Patient Preferences into Randomized Trials," *Journal of Clinical Epidemiology*, 53, 163-166.
- Little, R. J., and Rubin, D. B. (2001), "Causal Effects in Clinical and Epidemiological Studies via Potential Outcomes: Concepts and Analytical Approaches," *Annual Review of Public Health*,

21, 121-145.

McCullagh, P. and Nelder, J. (1989) *Generalized Linear Models*. Chapman and Hall, London.

Neyman, J. (1923), "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9," translated in *Statistical Science, (with discussion)*, 5, 465-480, 1990.

Robins, J.M. (1986), "A New Approach to Causal Inference in Mortality Studies with Sustained Exposure Periods - Application to Control of the Healthy Worker Survivor Effect," *Mathematical Modeling*, 7, 1393-1512.

Robins, J.M. (1987), Addendum to "A new Approach to Causal Inference in Mortality Studies with Sustained Exposure Periods - Application to Control of the Healthy Worker Survivor Effect," it *Computers and Mathematics with Applications*, 14, 923-945.

Rubin, D.B. (1974), "Estimating Causal Effects of Treatments in Randomized and Non-randomized Studies," *Journal of Educational Psychology*, 66, 688-701.

Rubin, D.B. (1977), "Assignment to a Treatment Group On the Basis of a Covariate," it *Journal of Educational Statistics*, 2, 1-26.

Rubin, D.B. (1978), "Bayesian Inference for Causal Effects: The Role of Randomization," *The Annals of Statistics*, 6, 34-58.

Rubin, D.B. (1984), "William G. Cochran's Contributions to the Design, Analysis and Evaluation of Observational Studies," In *W. G. Cochran's Impact on Statistics*, ed. PSRS Rao, J Sedransk, 37-69. New York: Wiley.

Rücker, G. (1989), "A Two-stage Trial Design for Testing Treatment, Self-selection, and Treatment Preference Effects," *Statistics in Medicine*, 8, 477-485.

Torgerson, D. J., Klaber-Moffett, J., and Russell, I. T. (1996), "Patient Preferences in Randomized Trials: Threat or Opportunity?" *Journal of Health Services Research and Policy*, 1(4), 194-197.

Wennberg, J. E., Barry, M. J., Fowler, F. J., and Mulley, A. (1993), "Outcomes Research, PORTs, and Health Care Reform," *Annals of the New York Academy of Sciences*, 703, 52-62.

Zelen, M. (1990), "Randomized Consent Designs for Clinical Trials: An Update," *Statistics in Medicine*, 9, 645-656.

Table 1. The Compliance Data from the WTP Study

	Random Arm		Choice Arm	
	Group	Self-Directed	Group	Self-Directed
Complier	144	152	299	134
Noncomplier	46	49	22	41
Total	190	201	321	175
Compliance Rate	0.76	0.76	0.93	0.77

Table 2. Parameter Estimates in the two-stage models (1) and (2) and For the WTP Data

Parameter	Estimate	SE ¹	SE ²
Outcome Model (1)			
Intercept	1.481	1.196	1.178
Age	0.002	0.0161	0.159
Employment	-1.015	0.299	0.277
Total Symptom Impact	0.0004	0.009	0.009
SIP Physical Dimension Score	-0.021	0.011	0.011
Treatment Assignment	-1.535	0.505	0.447
Treatment Preference	-0.362	0.403	0.371
Preference*Assignment	3.199	0.692	0.635
Model (2)			
Intercept	0.453	1.141	1.096
Age	0.003	0.016	0.015
Employment	-0.688	0.270	0.266
Total Symptom Impact	0.020	0.010	0.009
SIP Physical Dimension Score	-0.029	0.011	0.011

¹ Computed using the bootstrap method based on 1000 bootstrapped data sets.

² Computed by computing the observed information directly.

³ 1=Employed, 0=Unemployed.

Table 3. Estimates of the Compliance Probabilities by Incorporating Covariates for the WTP Data

Preference	Assignment	Estimate ¹	SE ²	Standardized Estimate ³	SE ²
Group	Group	0.929	0.014	0.928	0.014
	Self-Directed	0.725	0.054	0.723	0.054
Self-Directed	Group	0.424	0.091	0.441	0.093
	Self-Directed	0.771	0.033	0.786	0.032

¹ Calculated using equation (3).

² Calculated using the bootstrap method.

³ Calculated using equation (3) with $f(x_1|C = m)$ replaced by $f(1)$.

Table 4. Simulation Results Based on 125 Replications

Parameters	True Value	Estimate	Est. SE ¹	Est SE ²	Est. SE ³	Emp. SE ⁴
Outcome Model (1)						
Intercept	-2.00	-2.01	0.22	0.22	0.22	0.21
X_1	2.00	2.00	0.28	0.28	0.31	0.27
A	2.00	2.02	0.40	0.35	0.32	0.39
C	2.00	2.04	0.50	0.44	0.35	0.40
$A * C$	-2.00	-2.05	0.70	0.60	0.48	0.60
Preference Model (2)						
Intercept	-1.00	-0.98	0.18	0.18	0.16	0.19
X_2	2.00	1.97	0.32	0.32	0.27	0.31

¹ Estimated SE using the bootstrap method based on 300 bootstrapped data sets.

² Estimated SE using the observed information.

³ Estimated SE using the approximation in Ibrahim (1990).

⁴ Empirical SE.

Figure 1. Design of the WTP Study

Figure 2. Population Stratification and Principal Treatments

(a) Population Stratification

(b) Combinations of Population Stratum and Treatment

	Treatment			
Stratum	$\mathcal{A}_C$	$\mathcal{A}_R$	$\mathcal{B}_R$	$\mathcal{B}_C$
$P\bar{R}A$		$\bar{F}$	$\bar{F}$	$\bar{F}$
PRA				$\bar{F}$
PRB	$\bar{F}$			
$P\bar{R}B$	$\bar{F}$	$\bar{F}$	$\bar{F}$	

$\bar{F}$ = A-priori counterfactual; can not be estimated directly from data.

Figure 3. The Partial Randomized Preference Trial (PRPT)

(a) Design of the PRPT

(b) Estimable Effects from the PRPT

	Treatment			
Stratum	$\mathcal{A}_C$	$\mathcal{A}_R$	$\mathcal{B}_R$	$\mathcal{B}_C$
$P\bar{R}A$	E	$\bar{F}$	$\bar{F}$	$\bar{F}$
$PRA \cup PRB$	$\bar{E}$	E	E	$\bar{E}$
$P\bar{R}B$	$\bar{F}$	$\bar{F}$	$\bar{F}$	E

$\bar{F}$ = A-priori counterfactual; can not be estimated directly from data.

$\bar{E}$ = not an A-priori counterfactual; not estimable from this design.

E = estimable directly from this design.

Figure 4: The Doubly Randomized Preference Trial (DRPT)

(b) Estimable Effects from the DRPT

	Treatment			
Stratum	$\mathcal{A}_C$	$\mathcal{A}_R$	$\mathcal{B}_R$	$\mathcal{B}_C$
$P\bar{R}A$	E	$\bar{F}$	$\bar{F}$	$\bar{F}$
PRA		E	E	$\bar{F}$
PRB	$\bar{F}$			E
$P\bar{R}B$	$\bar{F}$	$\bar{F}$	$\bar{F}$	

(c) Estimable Effects under NSBR and ER assumptions from the DRPT

	Treatment	
Stratum	A	B
PA	E	E
PB	E	E

Figure 5. The Partial Randomized Preference Trial with the Preference Question (PRPTQ)

(b) Estimable Effects from the PRPTQ

	Treatment			
Stratum	$\mathcal{A}_C$	$\mathcal{A}_R$	$\mathcal{B}_R$	$\mathcal{B}_C$
$P\bar{R}A$	E	$\bar{F}$	$\bar{F}$	$\bar{F}$
PRA	$\bar{E}$	E	E	$\bar{F}$
PRB	$\bar{F}$	E	E	$\bar{E}$
$P\bar{R}B$	$\bar{F}$	$\bar{F}$	$\bar{F}$	E

Figure 6. The Doubly Randomized Preference Trial with the Preference Question (DPRPTQ)

(a) Design of the DPRPTQ

(b) Estimable Effects from the DRPTQ

	Treatment			
Stratum	$\mathcal{A}_c$	$\mathcal{A}_R$	$\mathcal{B}_R$	$\mathcal{B}_c$
$P\bar{R}A$	E	$\bar{F}$	$\bar{F}$	$\bar{F}$
PRA		E	E	$\bar{F}$
PRB	$\bar{F}$	E	E	
$P\bar{R}B$	$\bar{F}$	$\bar{F}$	$\bar{F}$	E

Figure 7. The Composite PRPT/DRPT (CRPTQ)

(a) Design of the CRPTQ

(b) Estimable Effects from the CRPTQ

	Treatment			
Stratum	$\mathcal{A}_C$	$\mathcal{A}_R$	$\mathcal{B}_R$	$\mathcal{B}_C$
$P\bar{R}A$	E	$\bar{F}$	$\bar{F}$	$\bar{F}$
PRA	E	E	E	$\bar{F}$
PRB	$\bar{F}$	E	E	E
$P\bar{R}B$	$\bar{F}$	$\bar{F}$	$\bar{F}$	E