LEARNING, PRIVATE INFORMATION AND THE ECONOMIC EVALUATION OF RANDOMIZED EXPERIMENTS

Tat Y. Chan and Barton H. Hamilton
John M. Olin School of Business
Washington University in St. Louis
Campus Box 1133
One Brookings Drive
St. Louis, MO 63130

Chan@wustl.edu; HamiltonB@wustl.edu

July 2006

ABSTRACT

Many randomized experiments are plagued by attrition, even among subjects receiving more effective treatments. We estimate the subject's utility associated with the receipt of treatment, as revealed by dropout behavior, to evaluate treatment effects. Utility is a function of both "publicly observed" outcomes and side effects privately observed by the subject. We analyze an influential AIDS clinical trial, ACTG 175, and show that for many subjects, AZT yields the highest level of utility despite having the smallest impact on the publicly observed outcome due to mild side effects. Moreover, although subjects enter the experiment uncertain of treatment effectiveness (and often the treatment received), the learning process implies that early dropout in ACTG 175 is primarily driven by side effects, while later attrition reflects declining treatment effectiveness.

Keywords: Learning, Side Effects, Attrition, Treatment Evaluation, Randomized Experiments, Clinical Trial, Discrete Choice Dynamic Programming, Simulation Estimation, AIDS, ACTG 175, CD4 Count.
JEL Classifications: I10, D83, C90

We are grateful to Martin Cripps, Pedro Mira, Jim Heckman, Bill Vogt, Bob Pollack, Glenn MacDonald, Don Kenkel, Dean Lillard, 3 anonymous referees and seminar participants at Rand, Cornell, Penn, Iowa State, Chicago, Washington University, the 13th Annual Health Economics Conference at Carnegie Mellon University, and the Duke/UNC Conference on Structural Models in Labor, Aging, and Health for many helpful comments. Early versions of this paper circulated under the title "Structural Econometric Analysis of Randomized Clinical Trials: The Case of ACTG 175."

LEARNING, PRIVATE INFORMATION AND THE ECONOMIC EVALUATION OF RANDOMIZED EXPERIMENTS

1. INTRODUCTION

Since Fisher (1935), randomized experiments have been the gold standard for treatment evaluation. In the standard evaluation of the randomized experiment, the outcomes of subjects randomly assigned to treatment and control groups are observed over a set period of time. Comparisons of outcomes across these groups are then used to recover treatment effects of interest. The primary advantage of this approach is that randomization eliminates the possibility that bias may be introduced into the estimation when individuals are able to choose their treatment status.[1] The randomized experiment has become the standard approach for the evaluation of medical and pharmaceutical treatments, and is increasingly viewed as a particularly attractive framework in economic and social contexts (Burtless (1995)).[2,3]

Some have questioned the inferences obtained from many randomized experiments, due in part to substantial attrition and non-compliance that plagues many experiments (e.g., Efron and Feldman (1991); Heckman, Smith, and Taber (1998)). These studies generally view attrition as primarily a statistical issue that should be addressed using statistical or econometric methods accounting for sample selection.[4] However, as discussed by Heckman and Smith (1998), attrition may be an important indication of how subjects themselves are evaluating the experiment, and may

---

[1] See Manski (1997) for a discussion of alternative treatment effects that may be generated from randomized experiments. See also Rubin (1974, 1978).

[2] Heckman and Smith (1995) discuss the advantages and disadvantages of social experiments.

[3] Hundreds of randomized clinical trials investigating medical efficacy are funded by the US National Institute of Health each year alone (Office of Technology Assessment (1983)), while Greenberg and Shroder (1997) report on 143 social experiments conducted through 1996.

[4] Recently, a number of studies in the epidemiology/statistics literature have addressed the issue of "broken" randomized experiments resulting from non-compliance or attrition. The focus of these studies has been on addressing the impact of sample selection on the assessment of outcomes, rather than the behavioral implications of non-compliance. See Frangakis and Rubin (1999) and Barnard et al (2003) for a discussion of these issues in a statistics/epidemiology context.

reveal information concerning subject preferences toward treatments and outcomes in the experiment.[5] Dropout decisions are likely to be a function of at least four factors: (1) the "publicly observed" outcomes, which are explicitly measured by the experiment's investigators and which are the focus of the experiment;[6] (2) "private" outcomes that the subject observes, but which are not explicitly measured by the investigators. We term these privately observed outcomes *side effects*;[7] (3) the subject's preferences for the treatment and the publicly and privately observed outcomes; (4) the treatment options available outside the experiment.[8] This paper investigates how the specification and estimation of explicit economic models of subject behavior in a randomized experiment, based on public and private information and preferences, may be used to generate new insights regarding the evaluation of the experiment or clinical trial.

To highlight the contrasts between the approach of this paper and the standard approach taken in the medical literature, consider a prototypical randomized clinical trial analyzing the impact of two drugs, A and B, for the treatment of AIDS (the example is hypothetical). As is common in many AIDS trials, health status is measured by the trial subjects' CD4 counts, with lower CD4 values implying greater compromise of the individual's immune system. Patients are initially randomized into groups receiving each drug treatment, and then measurements of their CD4 counts are taken over a number of periods. Figure 1 illustrates the outcome data associated with Drugs A and B from this hypothetical experiment. The figure shows that Drug A has greater positive impact on patient health status than Drug B, on average, although health status appears to decline for all

[5] Philipson and Desimone (1997) examine the impact of subject evaluation on inferences regarding treatment effects in randomized clinical trials in medicine. Philipson and Hedges (1998) note that investigators conducting the experiment may also be accumulating information concerning treatment effectiveness which potentially conflicts with subject evaluations.

[6] Examples of publicly observed outcomes include wage rates in a job training experiment or CD4 counts in an AIDS clinical trial.

[7] Examples of side effects include nausea or vomiting caused by the consumption of the drug in a medical trial, or the costs incurred when attending job training classes in a social experiment.

[8] Heckman (1992), Heckman, Hohmann, and Smith (2000), and Palaca (1989) discuss how subjects may substitute

participants over the course of the trial. One would likely conclude using the standard approach that Drug A is superior to Drug B.

Although evaluation of the relative effectiveness of the two drugs would seem to be straightforward from Figure 1, conclusions regarding treatment efficacy are less clear when trial attrition is taken into account. Figure 2 presents the cumulative fraction of patients remaining in each arm of the trial over the first 20 periods from baseline. Despite greater apparent effectiveness, subjects receiving Drug A are much more likely to drop out of the trial than are those receiving Drug B. How should the treatments be evaluated in this case? While the potential bias introduced by dropout behavior shown in Figure 2 may be statistically adjusted for in one fashion or another when comparing the outcomes in Figure 1,[9] the dropout behavior itself provides information that is useful in the evaluation of the treatment. The high dropout rate among subjects receiving Drug A suggests that despite its higher effectiveness, Drug A generates substantial side effects that generate disutility for many subjects. Consequently, attrition reveals the extent to which subjects are willing to trade off higher side effects for greater direct effects on health status, as measured by the publicly observed CD4 count, as well as treatment options that lay outside the randomized experiment.

We construct an economic model of decision-making in which individuals make utility maximizing choices concerning dropout/compliance that provides a rich framework for evaluating randomized experiments. We estimate the subject's utility associated with the receipt of alternative treatments, as revealed by dropout or compliance behavior, to evaluate treatment effectiveness. The idea of using discrete choices to infer preferences has a long history in economics (see McFadden (2001) for a summary). For example, Heckman (1974) used the labor force participation decisions of married women to infer preferences toward labor supply. Our empirical framework is similar to

---

outside treatments for those available in the experiment.

[9] See Scharfstein et al (1999) for an example of this approach applied to an AIDS clinical trial.

that used by Miller (1984) to study occupational choice, where workers have prior beliefs concerning the uncertain match quality of different occupations. After choosing a job, workers gradually learn about the quality of the match, update their beliefs concerning the future stream of returns associated with the job, and decide whether to continue in the occupation.[10] Similarly, subjects in randomized experiments are uncertain of the outcomes of the treatment they receive at the outset of the experiment. Furthermore, they may be uncertain as to which treatment they are receiving if the experiment is single- or double-blind. In our behavioral model, subjects acquire information over time through their participation in the experiment, learn about the effectiveness of the treatment (and update their beliefs concerning the type of treatment received) from their outcomes and decide whether to continue in the experiment or choose the outside option.[11] Consequently, we gain insight not only into the impact of the treatment on publicly observed outcomes, which has been the focus of the literature, but we also examine other features of the treatments, such as the importance of privately observed outcomes to subjects and the roles of learning and uncertainty in explaining behavior in the experiment.

We use our framework to analyze data from the AIDS randomized clinical trial ACTG 175 (see Hammer et. al. (1996)). ACTG 175 was a landmark randomized double-blind clinical trial designed to evaluate the effectiveness of four alternative therapies for HIV infected individuals with CD4 cell counts of between 200 and 500/mm^3: zidovudine (AZT); didanosine (ddI); zidovudine plus didanosine (AZT+ddI); and zidovudine plus zalcitabine (AZT+ddC). CD4 counts are a widely used

[10] This framework has been used to investigate a variety of economic phenomena, such as job turnover and the impact of tenure on wages (Jovanovic (1979)); learning about the value of patents (Pakes (1986)); the impact of advertising on the purchase of "experience" goods (Erdem and Keane (1996); Ackerberg (2003)); and the purchase of pharmaceuticals where individuals differ in their reactions to drugs (Crawford and Shum (2005)).

[11] Our model differs from Miller (1984) in that we utilize data on the publicly observed outcome, CD4 counts, while he does not use wage data associated with occupations. This data is important in identifying the separate impacts of treatment effectiveness and side effects on subject utility. Previous studies also do not address the issue of uncertainty regarding the type of treatment being received. Malani (2006) discusses the importance of the learning process in understanding placebo effects in medical clinical trials.

marker for the status of an individual's immune system, and hence the progression of AIDS in the patient. CD4 counts were measured at trial baseline and then at weeks 8, 20, 32…104 of the trial. The "publicly observed" outcomes are therefore the subjects' CD4 counts at each trial week. A notable feature of ACTG 175 was that roughly half of the subjects dropped out (i.e., did not return for follow-up) by the end of the second year of the trial. While the substantial attrition was noted in the initial evaluation of the trial (Hammer et. al. (1996)), Scharfstein et al (1999) was the first to attempt to account for dropout in the estimation of the treatment effect in the trial. However, Scharfstein et al do not consider the behavioral implications of the evolution of attrition over the course of the trial.

Our empirical specification is related to models found in the literature analyzing the demand for pharmaceuticals, in which individuals are initially uncertain of the effectiveness of a drug, but infer this over time via Bayesian updating. While Ching (2000) assumes that consumers are myopic, we follow Philipson and Desimone (1997) and Crawford and Shum (2005) and assume that subjects are forward-looking, since the effectiveness of HIV drugs may grow (or diminish) with consumption. In contrast to these studies, we are able to distinguish between the direct impact of drug effectiveness (as measured by CD4 counts) and side effects due to the availability of individual-level data on publicly observed outcomes. We can therefore examine how subjects weigh these factors when deciding to drop out or continue using the drug, which may be particularly important when deciding on which subgroups to target when the treatment is released to the general patient population.

Our findings suggest that the standard evaluation criteria (impact on CD4 counts) may indeed yield misleading conclusions regarding treatment effectiveness. Hammer et al. (1996) claim that AZT alone is an inferior to the other treatments in ACTG 175. While AZT does have the smallest impact on CD4 counts, our structural estimates imply that a non-negligible fraction of subjects prefer

AZT to all other therapies. In fact, the fraction of subjects who prefer AZT is about the same as the percentage that prefers AZT+ddC, which has a greater impact on CD4 counts but more negative side effects. The results also suggest that the subject learning process is a key component in understanding trial behavior. Random assignment to treatment arm and the double-blind design enable us to infer that attrition at the onset of ACTG 175 is driven primarily by side effects, since subjects have prior beliefs concerning effectiveness that are independent of treatment assignment. Subjects are balanced in terms of unobserved characteristics and preferences. Therefore, early on in the trial, subjects with low treatment effectiveness may remain while those with high treatment effectiveness may drop out due to side effects. As the trial proceeds, subjects update their beliefs concerning treatment effectiveness by observing their sequence of CD4 counts, and attrition is driven by the variation across treatment arms in terms of the impact on CD4 levels. This process explains the puzzling fact that dropout among AZT patients is initially similar to that of other subjects, despite the relative ineffectiveness of the drug, but by the end of the trial attrition among these individuals is substantially greater than for the other treatments.

2. THE ACTG 175 DATA

The data for this paper comes from AIDS Clinical Trial Group Study 175 (ACTG 175) that compares the impact of monotherapy to combination therapy for 2467 HIV-infected adults with screening CD4 cell counts from 200 to 500 per cubic millimeter. ACTG 175 followed a double-blind design, and subjects were recruited from AIDS Clinical Trials Units and National Hemophilia Foundation sites in the United States and Puerto Rico. Individuals were randomly assigned to one of four daily treatment regimens: 600 mg of zidovudine (AZT); 400 mg of didanosine (ddI); 600 mg of zidovudine plus 400 mg of didanosine (AZT+ddI); or 600 mg of zidovudine plus 2.25 mg of zalcitabine (AZT+ddC). AZT was first introduced into the market in 1987, and was a standard treatment for HIV during 1991 and 1992 when the trial was initiated. Appendix Table 1 shows that

56% of subjects had taken AZT prior to enrollment in ACTG 175. The remaining 3 treatments in the trial were considered experimental. For each subject, the "publicly observed" treatment response was measured by the longitudinal progression of CD4 counts recorded at different intervals. Subject CD4 counts were reported at baseline (week 0), week 8, and then every twelve weeks thereafter for a period of 104 weeks or more. Full details of the trial may be found in Hammer et al. (1996).

The analysis sample consists of 2200 subjects treated at 89 different trial sites.[12] We examine the effect of these treatments for the first two years after baseline, so that a subject potentially has CD4 counts recorded in periods $t = 0, 1, 2 \ldots 9$, corresponding to weeks $w =$ 0,8,20...104. The two-year window was chosen because the clinical endpoint of the trial for all subjects occurred at least two years after baseline. For the purpose of our analysis, attrition occurs when a subject chooses to end the treatment assigned at baseline prematurely (i.e., before reaching week 104).

Figure 3 plots the relationship between the median change in CD4 counts and weeks in the trial for subjects in each treatment arm. The CD4 profiles shown in the figure indicate that the (median) CD4 count of individuals receiving AZT declines in week 8 relative to baseline, while CD4 counts actually increased for subjects in the remaining three treatment groups. By week 104, median CD4 counts of AZT patients dropped by 37 units, but remained higher than the baseline level for the other subjects.[13]

Figure 4 plots the survivor function over the two-year period for each of the four treatment sub-samples. The figure exhibits a number of notable features. First, attrition appears to be a potentially important confounder in assessing the progression of CD4 counts observed in Figure 3,

[12] Most of the 267 subjects excluded from our analysis were dropped due to missing CD4 information in 1 or more weeks. The excluded observations appear to be randomly distributed across treatment arms (conditional upon non-attrition. Formally modeling the missing data process is beyond the scope of the paper.

[13] Median (and OLS) regression estimates indicate that the differences in the CD4 profiles across treatment arms are highly statistically significant.

since only about 35%-50% of subjects continue in the trial through week 104. The high amount of dropout has led authors such as Scharfstein et al (1999) to consider statistical selection corrections when analyzing the impact of treatment on CD4 counts. Second, not only does dropout vary by treatment arm, but there is some crossover and "fanning out" of the survivor functions. In particular, AZT patients are more likely to remain in the trial through week 20 than are AZT+ddC recipients, despite the observation from Figure 3 that AZT has a much lower impact on CD4 counts from the start of the trial. However, by week 104, the probability that an AZT+ddC patient remains in the trial is approximately 30% higher than the empirical survival probability of AZT patients. One explanation that is consistent with Figure 4 is that subjects are initially unsure of the impact of the drug they are taking on their CD4 count, and require a few periods to learn its effectiveness. This would explain why a higher fraction of AZT patients do not immediately drop out.

A subtle feature of the data shown in Figure 4 concerns the fact that ddI patients have a lower probability of leaving the trial by week 8 than do subjects in other treatment arms, while AZT+ddC subjects are the most likely to drop out. Table 1 presents estimates of a discrete hazard model applied to the ACTG 175 data, where the probability of leaving the trial in week *w* is estimated via a probit regression. Column (1) describes how the dropout hazard for AZT patients varies over the course of the trial, while the coefficients in columns (2) – (4) indicate whether the attrition hazard in the particular week differs significantly from that for AZT. The first row of the table confirms that ddI subjects have a significantly lower probability of leaving the trial in week 8 than AZT recipients. This difference should not reflect the direct effectiveness of the treatments, as measured by CD4 counts, since trial subjects are not supposed to receive information on CD4 counts before week 8. In addition, randomization implies that unobserved subject characteristics and preferences are balanced across treatment arms.

The week 8 difference in dropout hazards is consistent with the view that ddI subjects experience fewer side effects that would lead them to discontinue their medication, while AZT+ddC patients would appear to experience the most negative side effects. AZT, ddC, and ddI are known to cause a number of potential side effects, many of which are immediately apparent to the patient, but the side effects associated with ddI are less numerous.[14] However, the incidence and severity of these side effects may vary substantially across patients.[15]

The reason a subject drops out of the trial is reported in the ACTG 175 data. While the reported reasons are hardly definitive, and may mask multiple causes, some insight on the potential role of side effects in inducing attrition may be gained by examining the reason for dropout data in Table 2. The first row of each panel in the table shows that, conditional upon leaving the study for any reason, the death of a subject while participating in the trial is relatively uncommon. Moreover, few patients are explicitly removed from the trial by the ACTG 175 investigators. Consequently, dropout appears to be in large part the subject's decision.

Comparison of column (1) across Panels A-D shows that patients receiving AZT+ddC are more likely to request to leave the trial at some point due to toxic reactions to the treatment than are those receiving the other treatments, perhaps suggesting greater side effects of AZT and/or ddC. On the other hand, ddI subjects are the least likely to report toxic side effects as a reason for dropout. In addition, the fraction of patients who report dropping out due to toxic reactions declines markedly for those leaving the study in weeks 8 – 44 (column (2)) compared to dropouts in weeks 56 plus

[14] Side effects associated with AZT, ddI, and ddC include nausea, vomiting, rashes, and headaches. In some patients, AZT may also cause insomnia, fatigue, muscle pain, and anemia. Additional side effects of ddC may include mouth ulcers, chest pain, fever, and peripheral neuropathy, while ddI may also cause diarrhea, and, rarely, pancreatitis.

[15] Attrition was significantly lower among subjects who had used AZT prior to entering ACTG 175 (see Table 3 below). For example, the fraction dropping out in week 8 was 9.9% (14.8%) for those with (no) prior AZT use. These individuals may have learned to better cope with the effects of antiretroviral drugs, and their experience may lower the cost of complying with an exacting drug regimen. However, the relative pattern of attrition across trial arms in week 8 (and later), shown in Table 1, was not significantly affected by prior AZT use, suggesting that ddI

(column (3)). Of course, even patients dropping out in weeks 56 plus are likely to have learned fairly quickly about these toxic side effects and may have chosen to stay in the trial as long as they believed that the impact of the treatment on CD4 counts outweighed these negative side effects. Thus, it appears that patients learn quickly about the side effects of AZT and/or ddC. In contrast, subjects are more likely to request to leave the trial without a particular reason being reported later in the trial, perhaps reflecting the weaker effect of the drug received on CD4 counts. It may be the case that a subject with a strong reaction to the treatment may simply not return to the trial physician and be classified as lost to follow-up rather than leaving due to toxicity.

2.1 CD4 PROFILES BY ATTRITION GROUP

Another way to examine the relative importance of side effects across treatment arms is to compare the profiles of CD4 counts of non-dropouts with those of dropouts. Small cell sizes prevent us from constructing separate profiles for individuals dropping out in week 20, 32...104. Instead, subjects were classified into 3 groups: (a) those who did not drop out between week 0 and 104 (non-dropouts, 47% of subjects); (b) subjects dropping out between weeks 68 and 104 (Year 2 dropouts, 18% of subjects); and (c) subjects that dropped out between weeks 8 and 56 (Year 1 dropouts, 35% of subjects). Figure 5 plots the mean CD4 count for the three groups by treatment status. Note that for Year 1 dropouts, CD4 counts are available for all subjects in week 8, but the value shown in the figure for weeks 20, 32, and 44 only reflect the CD4 counts of subjects surviving that long (all Year 1 dropouts have left by week 56). Similarly, for the Year 2 dropouts, CD4 counts are available for all subjects prior to week 68.

Not surprisingly, each graph in Figure 5 shows that dropouts tended to have lower CD4 counts than stayers. This suggests that the decline in CD4 counts is an important factor explaining dropout behavior in ACT 175 for each treatment group. A notable difference across treatment

---

has fewer side effects regardless of prior AZT usage.

groups concerns the similarity in CD4 profiles for Year 1 and Year 2 dropouts who received AZT+ddC, shown in the lower left quadrant of Figure 5. If declining CD4 counts were the only factor in explaining attrition, then we would expect the CD4 counts of Year 2 dropouts to be greater than those of subjects dropping out in Year 1 in weeks 8-44, as is the case for the AZT recipients, for example. However, the similarity of the two profiles shown for AZT+ddC patients indicates that side effects are likely to play an important role in explaining dropout behavior among these subjects. Similar to the conclusions from Table 2, AZT+ddC subjects appear to drop out early on despite relatively low decline in CD4 counts compared to those receiving other treatments, suggesting greater negative side effects associated with this treatment.

3. MODELING SUBJECT BEHAVIOR IN THE RANDOMIZED EXPERIMENT

Subjects in the experiment decide each period whether to remain in the trial or drop out and seek alternative treatment. The decision to remain in the trial depends in part on the subject's evaluation of the direct impact on health status of the treatment received in ACTG 175, denoted by H_{it}, as well as the side effects experienced by the patient when taking the trial medication, S_{it}. While H_{it} is a measured outcome in the trial (e.g., the CD4 count) and is typically the focus of the evaluation of the efficacy of alternative treatments by trial investigators, side effects are assumed to be private information to the subject. Variation in the side effects associated with particular treatments potentially leads to the situation described in the Introduction where a treatment has a strong positive impact on health status, but subjects receiving the treatment are much more likely to drop out of the experiment in spite of the observed positive impact on H_{it}.

We consider a model of attrition that is structural in the sense that we specify subject utility functions and separately identify direct effectiveness and side effect distributions. Because drug therapies may take time to be effective, and patients may want to remain in the trial in order to keep the option of taking the experimental drug in future periods of the trial, subjects are allowed to be

forward-looking and will choose to remain in the trial for one more period if the discounted stream of expected current and future utilities is greater than the value of the outside option. We assume that subjects immediately observe side effects S_{it}, which is consistent with the results in Table 2 showing that patients quickly discover whether they have a toxic reaction to their treatment. However, the direct effectiveness of the treatment on H_{it} is not immediately known. The subject is uncertain about treatment effectiveness both because he is blinded to his assigned treatment arm, and because there is subject-level variation in the impact of each treatment on health status. Subject i learns about the effectiveness of treatment by observing the sequence $\{H_{it}\}$ over the course of the trial and updating beliefs via a Bayesian learning process. Dynamic behavior reflects the learning process as well as unexpected period-specific shocks and changes in the outside option over the course of the RE. The model is similar to that of Crawford and Shum (2005) in their analysis of the demand for anti-ulcer medication, but those authors did not have an observable measure of H_{it}; hence, they could not separately identify S_{it} from H_{it}. Conversely, in our application the data on CD4 counts and random assignment to treatment (as well as blinding) enables us to empirically distinguish the separate impacts of treatment effectiveness and side effects on subject utility.

3.1 THE SUBJECT'S DECISION PROBLEM

We now describe the subject's decision problem. To fix ideas, let $t = 0,1,\ldots,\tau$ index the time periods in the trial (corresponding to weeks 0, 8, 20,…,104) and $i = 1,\ldots,N$ denote subjects. The indicator variables $d_{it}^{\,k}$ indicate whether the subject remains in the experiment ($k = T$) or chooses the outside option and drops out ($k = o$) in period t. A subject may be in only one state in each period, so that $d_{it}^{\,T} + d_{it}^{\,o} = 1$. The objective of the subject is to choose a sequence of actions that maximizes the present value of lifetime utility:

$$(1) \qquad \max_{\{d_{it}^k\}} E\Big[\sum_{s=t}^{\tau} \beta^{s-t} \sum_{k=\{o,T\}} E[U_{is}^k \mid I_{is}] d_{it}^k \mid I_{it}\Big] \; + \; \beta^{s-\tau} V_i^o(I_{it}, \; \tau+1),$$

where U_{it}^{k} is the period-specific flow of to the subject from alternative *k*, I_{it} is the subject's information set at time *t*, and β is the discount rate. We assume that at the end of the trial period, all subjects leave the trial and receive the discounted lifetime flow of expected utility from the outside option, $V_i^o(I_{it}, \tau+1)$.

In ACTG 175 subjects are not allowed to re-enter the trial once they drop out, so we impose the constraint $d_{it+1}^{o} = 1$ if $d_{it}^{o} = 1$. Therefore, at any time *t* during the trial, the lifetime flow of expected utility available to the subject if he leaves the trial is given by

$$(2) \quad V_i^o(I_{it},\ t) \ = \ E[\sum_{s=t}^{\tau} \beta^{s-t} E[U_{is}^o \mid I_{it}] \ + \ \beta^{s-\tau} V_i^o(I_{it},\ \tau+1)].$$

Given the subject's information set at time *t*, I_{it}, the expected return of remaining in the trial in period $t < \tau$ may be written as

$$(3) \quad V_i^T(I_{it},t) = E[U_{it}^T \mid I_{it}] + \beta E[\max\{V_i^T(I_{it+1},t+1), V_i^o(I_{it+1},t+1)\} \mid I_{it}], \quad 0 \le t < \tau,$$

In the terminal period of the trial, $t = \tau$, the subject receives

$$(4) \quad V_i^T(I_{i\tau},\tau) = E[U_{i\tau}^T \mid I_{i\tau}] + \beta E[V_i^o(I_{i\tau+1},\tau+1) \mid I_{i\tau}]$$

by remaining in the trial, since once the trial is completed we assume that the patient receives the outside treatment option. The maximal lifetime utility of the subject at time *t* is then given by the maximum of the value of remaining in the trial for at least one more period and the value of leaving the trial and receiving the outside treatment forever, since subjects cannot re-enter the trial once they have dropped out:

$$(5) \quad V_i(I_{it},t) = \max\{V_i^T(I_{it},t), V_i^o(I_{it},t)\}.$$

For a positive discount rate, the decision-making framework outlined in equations (1) – (5) implies that forward-looking subjects may remain in the trial despite low current period utility, perhaps resulting from painful side effects, if they expect the future benefits of trial participation to

be particularly high. In addition, remaining in the trial for an additional period augments the subject's information set, which also may increase future utility as the patient learns more about the effectiveness of the treatment.

3.2 SPECIFICATION OF PREFERENCES

Subject *i*'s per-period utility obtained by participating in ACTG175 is a function of health status in period *t* as measured by CD4 count (i.e., the extent to which his immune system has been compromised), H_{it}, as well as the side effects (e.g., pain or nausea associated with consumption of treatment drugs) experienced by the individual when taking the trial drug, S_{it}. Because these side effects are not reported in ACTG 175 and are identified in part through random treatment assignment, we assume that utility is additively separable in H_{it} and S_{it}. While the separability assumption limits substitution possibilities, we do allow the impact of health status on utility to be non-linear, so that

$$(6) \qquad U_{it}^{T}(H_{it}, S_{it}) = -\exp(-\gamma H_{it}) - S_{it},$$

where $\gamma > 0$ is the coefficient of absolute risk aversion. The declining marginal utility of H_{it} implied by equation (6) captures the feature suggested by Figure 5 that side effects may be less important in driving attrition among individuals with low CD4 counts, but highly important in explaining dropout among healthy subjects in the trial. Our utility specification therefore allows the willingness of individuals to trade treatment effectiveness for improved side effects to vary depending on health status.

No information is available from ACTG 175 concerning the subject after he drops out of the experiment. Consequently, we specify the stream of lifetime utility of subjects dropping out of the experiment at time *t* defined in equation (2) to be a linear function of observed characteristics x_{oit} as well as a stochastic component:

(7) $V_i^o(I_{it},\ t) = x_{oit}\delta + \varepsilon_{iot}$.

3.3 SPECIFICATION OF HEALTH STATUS AND SIDE EFFECTS

Studies in the literature evaluating the impact of alternative treatments for HIV on patient health typically focus on the change in CD4 count relative to its baseline value to measure the degree to which the subject's immune system has been compromised (Boscardin et. al. (1998)). Therefore, subject *i*'s CD4 count at time *t*, H_{it}, is assumed to be a function of the subject's health status at the beginning of the trial, H_{i0}; the subject-level response to the treatment received; a subject (and treatment) specific time trend; and a period specific error term:

(8) $H_{irt} - H_{i0} = \theta_{ri} + \lambda_{ri}t + \upsilon_{it}$,

where the subscript r = {AZT, AZT+ddI, AZT+ddC, ddI} indicates the treatment arm the subject is randomly assigned to. The subject-level random intercept in equation (8) is normally distributed and allowed to depend on time-invariant subject characteristics, while the λ_{ri} do not:

(9) $\theta_{ri} \sim N(z_i\beta_r, \sigma^2_{\theta r}),\quad \lambda_{ri} \sim N(\lambda_r, \sigma^2_{\lambda r})$.

With the exception of treatment specific intercepts, the β_r are restricted to be equal across trial arm.

While the ACTG 175 data contain an observable measure of H_{it} (the subject's CD4 count), no direct measures of S_{it} are available. Consequently, side effects are specified to be a linear function of observed characteristics, including indicators for treatment arm, and a period-specific stochastic component:

(10) $S_{it} = x_{Si}\alpha_i + \varepsilon_{iSt}, \qquad \alpha_i \sim N(\alpha,\ \sigma^2_\alpha)$.

The random coefficient specification in equation (10) allows the impact of side effects to vary across individuals and enables subject-level estimation of the impact of the treatments on utility.

3.4 INCORPORATING SUBJECT LEARNING

Subjects decide whether to remain in the trial in period *t* before health status H_{it} is observed,

and must form expectations concerning its value. Subjects do not observe the idiosyncratic component υ_{it} in equation (8), but they do know its distribution. In addition, subjects are blinded as to which treatment arm they have been assigned. Subjects form expectations regarding their health status H_{it} following a two-step process. Subjects first update their beliefs regarding the impact of trial participation on health status, conditional upon assignment to treatment arm r. They then update their beliefs concerning the probabilities $\{p_{ir}\}$ of being assigned to each treatment arm.

Conditional upon assignment to treatment arm r, we assume that subjects believe their CD4 counts over the course of the trial follow the linear relationship

$$(11) \qquad H_{it} = H_{i0} + \mu_{irt} + \upsilon_{it} = H_{i0} + \mu_{0ir} + \mu_{1ir} t + \upsilon_{it}.$$

This specification of the prior mean captures the view that trial treatments may have cumulative effects that are not immediately observed. In addition, forward-looking subjects will take into account the expected future health status profile when choosing to remain in the trial. For example, if μ_{1ir} were positive, a subject may decide to remain in the trial due to the anticipation of high future benefits of treatment.

The prior beliefs for μ_{0ir} and μ_{1ir} are given by

$$(12) \quad \begin{pmatrix} \mu_{0ir} \\ \mu_{1ir} \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_{00r} \\ \mu_{10r} \end{pmatrix}, \Sigma_{\mu 0r} \right), \; \Sigma_{\mu 0r} = \begin{pmatrix} \sigma^2_{\mu 00r} & 0 \\ 0 & \sigma^2_{\mu 10r} \end{pmatrix}.$$

The prior variance $\Sigma_{\mu 0r}$ reflects the precision of the prior beliefs concerning treatment r of the subjects in the trial. While all trial participants have the same prior mean and variance at baseline, upon the commencement of the trial the subject observes a sequence of health status measures H_{it}, which are used to update the subject's prior beliefs according to the Bayesian rule (DeGroot (1970)). The posterior mean and variance of μ_{irt} in period t is given by

$$(13)\qquad E[\mu_{irt} \mid I_{it}] = \mu_{irt}^{t} = \begin{pmatrix} 1 & t \end{pmatrix}\left(\frac{A_{t-1}' A_{t-1}}{\sigma_{\upsilon}^2} + \Sigma_{\mu 0 r}^{-1} \right)^{-1} \left(\frac{A_{t-1}' H_{it-1}^{*}}{\sigma_{\upsilon}^2} + \Sigma_{\mu 0 r}^{-1} \begin{pmatrix} \mu_{00r} \\ \mu_{10r} \end{pmatrix} \right),$$

$$Var(\mu_{irt} \mid I_{it}) = \sigma_{\mu,r,t}^2 = \begin{pmatrix} 1 & t \end{pmatrix}\left(\frac{A_{t-1}' A_{t-1}}{\sigma_{\upsilon}^2} + \Sigma_{\mu 0 r}^{-1} \right)^{-1} \begin{pmatrix} 1 \\ t \end{pmatrix},$$

where

$$A_{t-1} = \begin{pmatrix} 1,1,...,1 \\ 1,2,...,t-1 \end{pmatrix}', \quad H_{it-1}^{*} = \left(H_{i1} - H_{i0}, ..., H_{it-1} - H_{i0} \right)'.$$

Expressions similar to (13) could be constructed to describe the learning process for side effects. However, given the evidence in Section 2, we assume that side effects are learned before week 8 by the patient (the idiosyncratic factors influencing side effects at time *t*, ε_{iSt} are observed by the subject, but not by the econometrician, when the dropout decision is made). This may be a reasonable assumption, given that many side effects reflect immediate discomfort (e.g., nausea, vomiting) associated with consumption of the treatment. Such features of the trial treatment seem likely to be observed within the first eight weeks. One issue in assuming that subjects learn their side effects immediately is that they may then infer treatment assignment from their side effects. However, this does not appear to be an important concern for ACTG 175 for a number of reasons. First, only AZT is an existing treatment in the market. As observed above, initial attrition is lower among AZT subjects despite its inferior effectiveness. If subjects inferred they were taking AZT from side effects, they should be more likely to drop out earlier. The estimates of side effects presented below also indicate substantial overlap in their distributions among the 4 drugs. This may in part reflect the fact that many side effects (e.g., nausea, headache, rash) are common to all drugs in the trial, implying that it may be very difficult to pinpoint the AZT treatment arm from side effects. Finally, we estimated the structural model presented below using the sub-sample of prior AZT users, who would be most likely to infer their treatment arm from side effects. We found very

similar results to those presented in Tables 3 and 4 below for the entire sample, suggesting that prior AZT receipt does not significantly influence the learning process.[16]

Let $R_{ir} = 1$ if subject *i* belongs to trial arm *r*, and 0 otherwise, so that $\sum_r R_{ir} = 1$. We assume that in period *t* R_{ir} follows a multinomial distribution with prior mean p_{ir}^t. After observing H_{it} and updating beliefs of the effectiveness of treatment *r*, μ_{irt}^t, subject *i* then updates beliefs concerning the probability of assignment to treatment arm *r* using the updating rule

$$(14) \qquad p_{ir}^{t+1} = E[R_{ir} = 1 \mid I_{it}] = \frac{p_{ir}^t \Pr(H_{it} - H_{i0} \mid I_{it}, R_{ir} = 1)}{\sum_{a \in \{AZT, AZT+ddI, AZT+ddC, ddI\}} p_{ia}^t \Pr(H_{it} - H_{i0} \mid I_{it}, R_{ia} = 1)},$$

where $\Pr(H_{it} - H_{i0} \mid I_{it}, R_{ir} = 1)$ is the posterior density after subject *i* observes H_{it}, conditional on belonging to treatment *r*. Because subjects are told at trial entry that they have an equal chance of being assigned to each of the four treatment groups, $p^0_{ir} = 0.25$.

Given the utility specification in equation (6) and the assumptions regarding subject learning described above, we show in the Appendix (section A) that the expected current period *t* utility of remaining in the trial given the subject's information set I_{it} may be written as

$$(15) \quad E[U_{it}^T(H_{it}, S_{it}) \mid I_{it}] = \sum_r p_{ir}^t \{(-\exp(-\gamma(H_{i0} + \mu_{irt}^t) + \frac{\gamma^2}{2}(\sigma_\upsilon^2 + \sigma_{\mu,r,t}^2)) - x_{Si}\alpha_i - \varepsilon_{iSt}\}.$$

Note that $\gamma > 0$ implies that increases in the expected value of health status increase utility, as does declining uncertainty regarding the subject's true state of health. Equation (15) may then be substituted into equations (3) and (4) to obtain the expected return to remaining in the experiment for at least one more period.

3.5 RANDOMIZATION, BLINDING, AND THE SPECIFICATION OF EXPECTATIONS

[16] Results available from the authors upon request. If prior AZT recipients used side effects to infer treatment effectiveness, we would have expected significantly different learning and side effect estimates for this subsample compared to the sample as a whole. Inferring the treatment arm from side effects may be more likely in placebo

Our discussion of the structural model of subject dropout behavior suggests that randomization at baseline is unlikely to identify the impact of treatments on CD4 counts, since non-random attrition over the course of the experiment introduces the possibility of sample selection bias. However, randomization does aid in the identification of preference parameters, including the side effect distributions associated with each treatment. As a result of random assignment, at the commencement of the trial subjects are balanced across treatment arms in terms of the levels of their unobserved characteristics and preferences. Consequently, differential attrition by treatment arm should arise from variation in treatment side effects and subject-specific treatment effectiveness rather than unobserved individual characteristics.

Subject blinding at baseline with respect to treatment receipt also helps to identify the side effect distributions. Blinding implies that a subject's prior beliefs concerning treatment effectiveness will be uncorrelated with treatment assignment. If subjects could observe their treatment assignments at baseline (or if they could immediately infer the assigned treatment from side effects), then they could construct different sets of prior beliefs depending on the treatment received. If so, differences in attrition across trial arms might reflect differences in prior beliefs, as well as side effects. A crucial difference between medical randomized trials and many social experiments is that blinding is usually not possible at the outset of the experiment in the latter context. For example, subjects are aware that they receive free job training or extended unemployment insurance benefits. Consequently, randomization and blinding are important aids in distinguishing between side effects and subject beliefs concerning the direct effectiveness of treatment receipt as explanations for subject choice behavior.

A final set of issues concerns the specification of the prior intercept and trend parameters μ_{00r} and μ_{10r} from equation (12) and the role of expectations. In many experiments, the control group

---

controlled trials.

may receive a placebo or there is substantial past experience with a treatment that is highly informative in specifying values for μ_{00r} and μ_{10r}.[17] For experimental treatments or relatively new therapies, such prior information is generally not available. These parameters appear to be poorly identified from data consisting of only trial participants, and so are difficult to estimate without making further assumptions. Data on individuals choosing not to enter the experiment may be necessary for more precise estimates of μ_{00r} and μ_{10r}.[18] In the absence of this data, we follow Crawford and Shum (2005) and assume that subjects have rational expectations concerning drug effectiveness. This assumption is not as strong as it might first appear, given the result below that learning in ACTG 175 was relatively rapid. Consequently, the rational expectations assumption will influence predicted attrition probabilities primarily at the beginning of the trial. Moreover, prior beliefs at baseline are independent of treatment assignment so that the rational expectations assumption does not affect relative dropout probabilities across arms at week 8. Finally, because AZT+ddI, AZT+ddC, and ddI were all experimental treatments in ACTG 175, we assume that subjects had common prior beliefs concerning the effectiveness of these three therapies (i.e., $\mu_{00AZT+ddI} = \mu_{00AZT+ddC} = \mu_{00ddI}$ and $\mu_{10AZT+ddI} = \mu_{10AZT+ddC} = \mu_{10ddI}$). Of course, learning is still likely to be important because subjects do not know the drug they are receiving or their subject-specific reaction to the drug.

4. ECONOMETRIC IMPLEMENTATION

The estimation procedure for our structural model is complicated by the fact that dropout decisions are correlated over time, as are the observed health status measures $\{H_{it}\}$. Therefore, the

[17] In the case of a placebo-controlled experiment, one might specify $\mu_{00r} = \mu_{10r} = 0$ in the control arm (assuming that subjects do not believe there is a placebo effect; see Malani (2006)).

[18] We estimated the model presented in Section 5 including μ_{00r} and μ_{10r} as parameters. Our estimates of these values indicated that subjects' prior beliefs were that their CD4 counts would increase if they participated in the trial. However, we found that different starting values led to different estimates of μ_{00r} and μ_{10r} (though still positive) with virtually no change in the likelihood value. We also estimated the model assuming that individuals believe the three experimental

econometric approach relies on simulation-based methods and proceeds in two steps. Consider a subject who drops out in period *t*. The likelihood contribution of this subject is the joint probability of observing the sequence of health outcomes $H_{i1}, H_{i2}\ldots H_{it-1}$ and dropout in period *t* and is given by

$$(16)\quad L_i = \Pr(H_{i1},...,H_{it-1}, d_{i1}^T = 1,...,d_{it-1}^T = 1, d_{it}^T = 0 \mid \Theta_1, \Theta_2, x_i, z_i),$$

where the parameters from the dropout decision and health status equations are denoted by $\Theta_1 = \{\beta, \gamma, \delta, \alpha, \sigma_\alpha, \Sigma_{\mu 0 r}\}$ and $\Theta_2 = \{\beta_r, \sigma_{\beta r}, \lambda_r, \sigma_\upsilon\}$, respectively, and $x_i = \{x_{Si}, x_{oit}\}$. Equation (16) can be rewritten as

$$(17)\quad L_i = f(H_{i1},...,H_{it-1} \mid d_{i1}^T = 1,...,d_{it-1}^T = 1, d_{it}^T = 0, \Theta_1, \Theta_2, x_i, z_i) \\ * \Pr(d_{i1}^T = 1,...,d_{it-1}^T = 1, d_{it}^T = 0 \mid \Theta_1, \Theta_2, x_i, z_i).$$

Using equations (3), (4), and (7), this is equivalent to

$$(18)\quad L_i = f(H_{i1},...,H_{it-1} \mid V_{i1}^T - V_{i1}^o > 0,...,V_{it-1}^T - V_{it-1}^o > 0, V_{it}^T - V_{it}^o < 0, \Theta_1, \Theta_2, x_i, z_i) \\ * \Pr(V_{i1}^T - V_{i1}^o > 0,...,V_{it-1}^T - V_{it-1}^o > 0, V_{it}^T - V_{it}^o < 0 \mid \Theta_1, \Theta_2, x_i, z_i).$$

We form the expected future value of the value function using the law of iterated expectations, which accounts for the fact that if the subject remains in the trial an additional period, equation (13) indicates that uncertainty is reduced, which yields increased utility as demonstrated by equation (15). The full derivation of the likelihood in the forward-looking case is described in the Appendix (Section B).

Both x_{Si} and x_{oit} and ε_{Sit} and ε_{oit} enter linearly into V_{it}^T - V_{it}^o in equation (18), so that only contrasts between the two are identified. We assume that Var(ε_{Sit} - ε_{oit}) = 1 because the scale of the difference in values is not observed. Directly maximizing (17) is difficult because the sequence of subject-period observations is not independent for individual *i*. Consequently, we use simulated

---

drugs would have the same effectiveness as AZT. The results from this specification were very similar to those reported in Table 3, mainly because of the rapid learning rate described below.

maximum likelihood to obtain the estimates.[19]

5. RESULTS

The literature on ACTG 175, such as Hammer et al (1996), suggests that the following variables might affect the subject-level treatment impact on CD4 counts, and hence are included in z_i: (1) demographics: age, gender, and race; (2) extent of disease at baseline: the screening CD4 count measured prior to baseline that determined entry into the trial; presence of symptomatic HIV infection (indicating greater spread of the disease); and prior antiretroviral therapy (also suggesting greater spread of the disease).

A key feature of the empirical analysis is the specification of the factors that influence the outside option available to trial participants at time *t*. Perhaps the most important factor is the set of AIDS treatments that become available over the course of the trial. As more effective drugs enter the market, subjects will be more likely to exit the trial to obtain them. While we do not observe such alternative treatments directly, we assume that their availability is correlated with calendar time. Intake into ACTG 175 occurred during the early 1990s, during which time new AIDS treatments were becoming available.[20] The vector x_{iot} thus contains a variable indicating the current calendar quarter of the trial since April 1991, as well as an indicator of whether the subject is an IV drug user, since these individuals are less likely to follow up in the trial (the other categories are haemophiliac and homosexual). Subjects with more financial resources may be more able to search and afford non-ACTG 175 treatments. While income is not reported, subject age, gender, and race may be correlated with the financial resources of the patient, and are included as proxies.

Side effects are specified to be a function of treatment assignment, and, as shown in equation

[19] See Hajivassiliou and Ruud (1994) for an overview of classical simulation methods for limited dependent variable problems.
[20] Duggan and Evans (2006) examine the impact of the introduction of new HIV drugs on mortality and pharmaceutical expenditure.

(10), the impact of the side effects is assumed to be subject-specific. In addition, we allow the side effects associated with each treatment to differ depending on whether the subject has received antiretroviral therapy prior to ACTG 175 and whether the individual has a symptomatic HIV infection, in addition to the CD4 count at screening.

5.1 RESULTS FOR THE STRUCTURAL ATTRITION MODELS

The parameter estimates from alternative structural models of the decision to remain in ACTG 175 or seek alternative treatment are presented in Table 3. Positive coefficients indicate a higher probability of remaining in the trial. The value of β is initially set to equal an annual discount rate of 0.95. We present results from different specifications of the structural model to assess the sensitivity of the estimates to alternative assumptions regarding discount rates, prior beliefs, and the outside option.

Turning to our base forward-looking specification in column (1) of Table 3, the positive estimate of γ implies that subjects with higher expected CD4 counts are more likely to remain in the trial. Figure 6 plots the marginal utility of a change in CD4 count implied by equation (6) (with no uncertainty) for values of H_{it} ranging from 1 to 650. The figure implies that for individuals with CD4 counts less than 200 (indicating serious immune system damage and high risk of opportunistic infection) an incremental rise in CD4 counts substantially increases the utility associated with remaining in the trial, reducing dropout. Conversely, for individuals with CD4 counts greater than 500, which are generally toward the normal range, changes in CD4 count have little impact on utility. For these subjects, dropout is thus likely to be driven primarily by the side effects of treatment. These findings are consistent with Figure 5, which showed that dropouts tend to have lower and declining CD4 counts, particularly in the case of AZT subjects. Consequently, in the absence of side effects, and with perfect information concerning the impact of treatment on CD4

counts, subjects receiving combination therapy or ddI should be more likely to remain in ACTG 175 than those receiving AZT alone.

Subjects who had prior antiretroviral therapy are more likely to remain in the trial, as are symptomatic individuals who are experiencing health problems similar to those induced by treatment side effects (symptoms include diarrhoea, minor skin or oral conditions, lack of energy, swollen glands, etc.). These factors may thus have less influence on dropout behavior. Intravenous drug users are much more likely to drop out of the trial in each treatment arm. If these individuals are to be studied further, incentives must be provided in order to induce them to remain in the trial. AIDS treatments available outside ACTG 175 appear to become increasingly attractive over time, since subjects who enrolled later in the trial are more likely to drop out. Finally, older individuals tend to be less likely to drop out of the trial, perhaps because they are less willing to move in order to investigate alternative treatments outside the trial.

The specification in column (1) of Table 3 assumes that the outside option does not depend on health status. However, healthier individuals may have a wider range of options. For example, alternative non-medical therapies are more likely to be taken if the risk of developing AIDS-related health complications is relatively low. Sicker individuals require closer monitoring and more intensive treatment, which may be subsidized or offered for free while the individual is enrolled in the trial. In column (2), we allow the outside option at time t to depend on the subject's CD4 count. While the pattern of side effects remains unchanged, the estimates imply that the outside option is indeed more attractive when the individual's (expected) CD4 count is higher, all else equal. Consequently, even if health status deteriorates more than expected during ACTG 175, a subject may still choose to remain in the trial because the outside option is poor. Conversely, some individuals for whom the trial treatment is effective may drop out because the marginal utility associated with an increase in the CD4 count is relatively low and the outside option more attractive.

We consider the overall fit of the model by comparing the predicted probabilities of dropout generated by the coefficient estimates in column (2) with the observed frequencies in ACTG 175. For each subject, we predict the probability of dropout in weeks 8, 20, 32…104 and then construct sample averages. We also use the estimates in column (3) specifying myopic subject behavior (i.e., $\beta = 0$) to construct predicted probabilities. Figure 7 shows that both specifications fit the data well, on average. For the most part, the mean predicted probabilities of dropout are quite close to the observed attrition frequencies. The main difference occurs in week 8; both models over-predict the survival probability, with the forward-looking model producing a slightly worse fit. This likely reflects the assumptions concerning prior beliefs regarding the trend in CD4 profiles and the specification of time-invariant side effects. However, the discrepancy between the observed and predicted week 8 dropout is not more than five percentage points, suggesting that the additional complication of a richer side-effect specification may not yield substantial extra benefits. The similar fits of the myopic and forward-looking models to the data has been observed in other studies applying dynamic discrete choice models (Keane and Wolpin (1997)).

5.2 HOW FAST DO SUBJECTS LEARN?

The estimated standard deviations of the prior intercept and time slope parameters are shown in the bottom rows of Table 3. We restricted the standard deviation of the prior intercept to be equal for the AZT and non-AZT treatment arms since our initial estimate of $\sigma_{\mu_{00\,AZT}}$ was imprecise.[21] The estimates indicate substantial variation in subjects' prior beliefs concerning the effectiveness of trial treatment. The standard deviation of the prior time slope is lower in the AZT arm, as might be

[21] When we allowed for treatment arm-specific standard deviations of the prior intercepts, the estimated values were 7.03 for AZT and 65.23 for the non-AZT arms. However, due to the large standard error associated with the estimated value of $\sigma_{\mu_{00\,AZT}}$, we restricted the standard deviations to be equal in the specifications reported in Table 3. This restriction did not impact the estimates of the other parameters in the model. Use of prior information concerning variation in subjects' beliefs is likely to be useful in separately identifying these parameters. For example, one might assume that there is little variation in beliefs concerning the effectiveness of a placebo. Gelman

expected given the experience of many individuals with AZT prior to enrolling in ACTG 175. In order to assess implications of these estimates for the speed with which expected beliefs converge to the actual CD4 count for a particular subject, we constructed the expected CD4 count in each period (using the parameter estimates from column (2) of Table 3) conditional upon assignment to treatment arm *r*, equal to $H_{i0} + \mu_{itr}^{t}$, using equation (13) and the rational expectations assumption for a representative subject. Assuming the subject has a baseline CD4 count of 350, we calculate the subject's unconditional expected CD4 count in each period as $\Sigma p_{ir}^{t}(H_{i0} + \mu_{itr}^{t})$ using equation (14).

Figure 8 presents the subject's expected CD4 count conditional upon assignment to AZT ("Belief |AZT") or one of the other three treatment arms ("Belief |not AZT"), as well as the unconditional expected ("Belief") and actual CD4 counts for a representative subject actually taking AZT.[22] Under the assumption of rational expectations, the subject's prior beliefs concerning his CD4 count at week 8 differ markedly depending on treatment arm. The subject has optimistic initial beliefs concerning the impact of the experimental (not AZT) treatments, so that the expected CD4 count at week 8 substantially overstates the actual level. If the subject knew he was receiving AZT, his beliefs would be close to the actual change in the CD4 count at week 8. Taking the weighted average of these values to form unconditional beliefs, the subject is therefore more likely to remain in the trial at week 8 than he would be if he knew he was receiving AZT. Because subjects are blinded and have not received information on their CD4 counts in the trial before week 8, initial differences in attrition across trial arms result from the side effects of the different treatments. After observing the week 8 CD4 count, the model estimates imply that the subject revises downward his expected week 20 CD4 count for the non-AZT arms, which increases the likelihood of dropout compared to the full information case. However, by week 32 posterior beliefs under each trial arm

---

et al (2004) discuss the use of restrictions to identify the parameters of mixture distributions.

[22] Similar features are observed for subjects receiving AZT+ddC, AZT+ddI and ddI.

have generally converged to the actual CD4 count, indicating that learning occurs fairly rapidly in ACTG 175. The quick convergence is generated by the large estimated standard deviations in the prior beliefs of intercept and time slope for both types of treatments. Attrition is less likely to result from mistaken beliefs concerning treatment effectiveness in later periods of the trial.[23]

5.3 SIDE EFFECT DISTRIBUTIONS

To provide an indication of how patients value side effects, the results from Table 3 suggest that the average AZT patient with a CD4 count of 350 would be willing to trade a 5.2% reduction in his CD4 count for a 1% improvement in the side effects associated with AZT. Due to the declining marginal utility of H_{it} shown in Figure 6, the same patient with a CD4 count of 200 would only be willing to accept a 2.6% reduction in his CD4 count in exchange for a 1% improvement. Comparing treatments, the coefficient estimates of the side effects in Table 3 indicate that patients receiving ddI are less likely to drop out of the trial, after accounting for its expected impact on CD4 counts, than are other subjects, suggesting that ddI recipients experience the fewest side effects. Conversely, AZT+ddC generates the worst side effects, on average. However, the subject level variation in side effects suggests that many individuals in this group may still experience fewer side effects than those receiving AZT, AZT+ddI, or ddI. Plots of the subject-level side effect distributions for the four treatment groups shown in Figure 9 confirm that most ddI patients are less likely to drop out of ACTG 175, all else equal, than the average AZT+ddC subject, for example. Nevertheless, there is still substantial overlap of the side effect distributions. One interpretation of these findings that is consistent with the difference in the treatment specific survivor functions plotted in Figure 4 is that AZT+ddC subjects are initially more likely to drop out due to the immediately perceived higher side effects associated with combination therapy. Over time, the fact that AZT subjects experience

[23] The posterior variance also declines over the course of the trial. From equation (14) and the positive estimate of γ, reductions in the posterior variance increase the expected utility associated with remaining in the experiment.

greater declines their CD4 counts offsets the lower side effects associated with this treatment relative to AZT+ddC, leading to increased attrition among patients receiving AZT.

5.4 THE IMPACT OF ACTG 175 TREATMENTS ON CD4 COUNTS

Table 4 presents the estimates for equation (8) measuring the change in CD4 counts that incorporates treatment specific random intercepts and time trends. The results indicate that AZT and AZT+ddC patients experience the smallest initial increase in CD4 counts, on average, compared to those receiving ddI alone or in combination with AZT. On the other hand, AZT+ddC subjects experience the least deterioration in their immune systems over the course of the trial, while AZT patients have the greatest mean decline in CD4 count. By week 104, the estimates in Table 4 suggest that AZT+ddC subjects have the highest CD4 count, on average, and AZT patients the lowest. There is substantial subject-level heterogeneity in CD4 profiles, particularly for patients receiving AZT+ddI; AZT generates the least amount of subject-level variation among the four treatments. The magnitudes of the standard deviations in the intercept and time trend parameters suggests that there may be some subjects for whom AZT has the greatest impact on CD4 counts, despite the fact that on average this is the most inferior treatment. With regard to the other covariates, the major finding appears to be that individuals with prior antiretroviral treatment experience greater declines in CD4 counts, perhaps indicating that their immune system was more compromised at baseline. Higher screening CD4 counts are positively associated with changes in CD4 levels, although this finding is not statistically significant.

6. SIMULATING THE IMPACT OF TREATMENTS ON SUBJECT UTILITY

Investigation of the treatment effect distributions associated with side effects and CD4 counts indicates that the ranking of the ACTG 175 therapies on each of these measures did not necessarily coincide. For example, AZT+ddC has the most negative side effects, while it is clearly superior to AZT when examining CD4 counts. We now turn to the major goal of our paper: assessing the

impact of the ACTG 175 treatments on subject utility, which incorporates impacts on both side effects and CD4 counts. This analysis addresses the question of which treatment the subject would have chosen had he or she been able to do so.

Using the structural estimates in column (2) of Table 3 and the CD4 profile estimates in Table 4, we construct the expected discounted values at baseline of the stream of utilities associated with the four ACTG 175 treatments for each subject in the trial. We assume that each subject has full information regarding the subject-specific side effect and direct effectiveness distributions of each treatment, and that the subject knows his or her "type" (i.e., the subject specific side effect and CD4 effect for each treatment). Subjects choose the treatment that maximizes expected discounted utility and are assumed to remain in their chosen treatment for 9 periods corresponding to weeks 8 – 104. One thousand draws are taken for each subject in the simulation, and the fraction of individuals choosing each of the therapeutic alternatives is recorded. This approach is similar to the "voting criterion" for treatment evaluation discussed by Heckman and Smith (1998) that more closely mimics purchasing decision-making behavior by consumers.

One difficulty in conducting simulations of this type is that information on the correlation of the subject-level side effects (or the direct effect on CD4 counts) across the four treatments is not known; we do not observe the same subject receiving two or more of the treatments in the data. Consequently, simulations are conducted under alternative assumptions regarding the correlations of the four side effect and four drug effectiveness distributions derived from Tables 3 and 4. Table 5 presents the fraction of patients preferring each treatment under alternative scenarios concerning the correlation of the subject-level side effect and drug effectiveness distributions across the four treatments. We conduct simulations assuming: no correlation (column (1)); moderate correlation (positive and/or negative – columns (2) – (5)); and perfect correlation (positive and negative - columns (6) and (7)). Finally, the simulation results in the table constrain subjects to choose one of

the four trial treatments.

The preference data presented in Table 5 indicates three notable findings. First, AZT+ddI or ddI alone is the preferred therapy by a plurality of subjects in each of the seven scenarios. This is not surprising; these treatments have the fewest negative side effects, on average, and have the largest initial impact on CD4 counts. However, the substantial subject-level heterogeneity associated with the effects of these treatments suggests that other therapies may be preferred by many patients. In fact, no treatment is preferred by than 35% of subjects in any of the scenarios, with the exception of the simulation assuming perfect positive correlations in column (6), where 45% of individuals prefer ddI. The greater subject-level heterogeneity associated with AZT+ddI and ddI also explains why their shares increase as the correlation across the subject-specific distributions increases in magnitude.

An unexpected result from Table 5 is that a non-negligible fraction of subjects prefer AZT alone to the other treatments, despite the fact that AZT had the smallest mean impact on CD4 counts. Moreover, approximately the same number of individuals appears to prefer AZT+ddC, despite the fact that AZT+ddC has the largest positive impact on CD4 counts at the end of the trial. For many individuals, the negative side effects of AZT+ddC outweigh the positive impact of the therapy on CD4 counts. The preference of many subjects for AZT does not simply reflect the constraint of not choosing the outside option. When allowed, under the scenario in column (2) 25% of subjects would choose the outside option (which may include AZT) over the trial treatments. However, under this scenario 13% continue to prefer AZT compared to 23% choosing ddI (25% and 14% prefer AZT+ddI and AZT+ddC, respectively). In fact, the decline in preference share for ddI from 33% to 23% in this scenario suggests that many individuals choosing the outside option appear to be

those who took ddI in column (2).[24] Overall, the simulations suggest that ddI, either alone or in combination with AZT, is the treatment preferred by a majority of trial subjects under a variety of assumptions regarding the correlation of subject-level side or direct effects, although the fraction of subjects preferring these drugs is generally not more than two-thirds. Moreover, even though AZT might be judged to be the least desirable treatment on the basis of CD4 counts, many subjects still prefer this therapy, perhaps due to its relatively mild side effects when compared to drug combinations such as AZT+ddC. One implication of these simulations is that offering a menu of therapies may enhance patient welfare, since no single treatment yields the highest utility for all patients.

6.1 LEARNING AND SUBJECT CHOICE

The simulations in Table 5 assume that subjects know their side effects and impact on CD4 counts with certainty. We now evaluate the impact of learning and uncertainty on the optimal choice of drugs by trial subjects. We ask the following question: How do the choices of subjects differ from those shown in Table 5 if subjects have t periods of information on the side effects and effectiveness of each of the four drugs? To answer this question, we simulate subject choices under different amounts of information using the parameter estimates from Tables 3 and 4 and the Bayesian updating formulas in equations (13) and (14), and compare those choices to those made under perfect certainty as shown in Table 5. We assume that the correlation of side effects and direct effects across treatments is 0.5 (corresponding to column (2) of Table 5 and again assume that subjects cannot choose the outside option.

Figure 10 presents the results of these simulations for the four treatments under different

[24] Under the scenario of perfect positive correlation in column (7) of Table 5, 35% of subjects choose the outside option, while the shares of AZT, AZT+ddI, AZT+ddC, and ddI are 7%, 27%, 7%, and 24%, respectively. With either zero or negative correlations in side effects or effectiveness, virtually no subjects prefer the outside option. In these scenarios the subject is likely to receive a good draw from the side effect or effectiveness distributions of at least one trial treatment, leading them to forego the outside option.

levels of information. We denote a subject as "mismatched" in period *t* if the individual's drug choice with *t* periods of information is different from the full information choice. The figure shows that approximately one-third of subjects whose full-information choice from Table 5 is AZT would choose a different drug if they only had 1 period of information (i.e., 1 CD4 count observation) for each treatment. With only one period of information, prior beliefs concerning the superiority of the 3 experimental drugs dominate expectations of treatment effectiveness, leading many subjects whose full information choice is AZT to choose the "wrong" treatment. Similarly, the negative side effects of AZT+ddC lead many to initially choose an alternative drug. Many of these mismatched subjects choose AZT+ddI or ddI, which have the least negative side effects and the largest initial effect on CD4 counts. However, the figure indicates that learning occurs quickly (within five periods in our simulation) as accumulating experience outweighs prior beliefs. The share of individuals choosing AZT+ddI or ddI declines as subjects acquire information concerning treatment effectiveness. The speed of the learning process suggests a significant economic value of experience, since even one period of use substantially reduces the number of mismatches and hence significantly enhances patient welfare.

7. CONCLUSION

This paper presents a framework for evaluating longitudinal randomized experiments, such as clinical trials conducted to evaluate medical treatments, based on subject utility rather than solely on the "publicly observed" outcomes that have typically been the focus of the literature. The standard approach does not capture a variety of features often observed in randomized experiments, such as attrition among subjects receiving the more effective (in terms of publicly observed outcomes) treatment and changing dropout patterns across treatment arms over the course of the experiment. Our evaluation framework incorporates these factors by viewing subjects as utility maximizing agents who compare the stream of utility associated with remaining in the trial for at

least one more period with the value of the outside option. Subject utility is specified to be a function of both publicly observed outcomes and the side effects of the treatment that are private information to the subject and are not measured by the experiment's investigators. Randomization and subject blinding aid in the identification of these side effect distributions. In addition, subjects form prior beliefs concerning treatment effectiveness (and treatment assignment) and then update these beliefs over the course of the trial as they accumulate information. Our framework therefore allows us to distinguish between learning, side effects, and the direct effect of treatments when explaining subject behavior in the experiment. We are then able to evaluate treatment effectiveness by jointly considering both the direct and side effects of each treatment on subject utility.

We apply our framework to evaluate the impact of four alternative drug therapies for AIDS using data from a well known longitudinal randomized clinical trial, ACTG 175. Previous evaluations of ACTG 175 tout the benefits of combination therapies for the treatment of AIDS, rather than the use of AZT alone, due to the superior impact of combination therapy on patient CD4 counts in the experiment (Hammer et al (1996)). Our framework generates different conclusions. Using a structural model of attrition, we find that for a significant fraction of subjects (generally about 18%-20%), AZT alone yields higher utility than the other treatments. Moreover, AZT+ddC has the most negative side effects, on average, among the trial treatments even though it has the highest impact on CD4 counts by the end of the trial. Due to the declining marginal utility of CD4, subjects with relatively high CD4 counts are willing to trade off the direct effect of the treatment for improvements in side effects. As a result, AZT+ddC is only preferred by about as many subjects as those preferring AZT alone. AZT+ddI and ddI are predicted to yield the highest stream of utility for the majority of patients, due to mild side effects and positive impact on CD4 counts of these treatments. Finally, substantial learning is observed over the course of the experiment, so that early dropout is primarily driven by side effects, while later attrition reflects declining CD4 counts for

many subjects. Overall, an important implication of our findings, not recognized using the standard evaluation approach, is that patient welfare may be enhanced by offering a menu of therapies, since no single treatment is preferred by a majority of patients.

A limitation of the evaluation of ACTG 175 reported in this paper is the lack of data on subjects choosing not to participate in the experiment. As noted by other authors (e.g., Heckman and Smith (1995)), the lack of information on non-participants limits the generalizability of the results, since the outside options of non-participants (or their beliefs about the effectiveness of trial treatment) is likely to differ from those of participants. If data on non-participants could be obtained, the decision to enrol in the trial could be incorporated into the framework through the specification of the distribution of prior beliefs concerning treatment effectiveness, and perhaps the findings could be generalized to the HIV population (with CD4 counts between 200 and 500) as a whole.

Our evaluation framework can be easily extended to capture other features often observed in randomized experiments, such as partial or non-compliance with treatments or crossover behaviour, if subjects are allowed to choose an alternative experimental treatment. Such features may provide greater information concerning correlations in side effects or direct effects across treatments. Multiple publicly observed outcomes are also easily incorporated into the utility specifications. In fact, a particular advantage of the structural approach with multiple outcomes is that one can estimate the relative value subjects place on each of the outcomes of the experiment. This is particularly important if treatment effectiveness varies across outcomes. Finally, while we applied the framework to the evaluation of a medical clinical trial, it can also be used to investigate social experiments, although the likely lack of subject blinding would have to be incorporated into the specification of prior beliefs.

TABLE 1
DISCRETE HAZARD ESTIMATES OF THE PROBABILITY OF ATTRITION, BY WEEK AND TRIAL ARM

Variable	Trial Arm			
	AZT	AZT+ddI	AZT+ddC	ddI
	(1)	(2)	(3)	(4)
Week 8	-1.148	0.017	0.096	-0.247
	(0.068)	(0.096)	(0.095)	(0.103)
Week 20	-1.170	-0.232	-0.060	-0.203
	(0.073)	(0.111)	(0.106)	(0.108)
Week 32	-1.307	-0.126	-0.119	-0.227
	(0.084)	(0.122)	(0.123)	(0.124)
Week 44	-1.264	-0.178	-0.102	-0.255
	(0.086)	(0.127)	(0.125)	(0.127)
Week 56	-1.246	-0.154	-0.087	-0.168
	(0.090)	(0.130)	(0.130)	(0.128)
Week 68	-1.321	-0.153	-0.187	-0.084
	(0.099)	(0.143)	(0.147)	(0.137)
Week 80	-1.246	-0.338	-0.178	-0.096
	(0.100)	(0.152)	(0.146)	(0.138)
Week 92	-1.286	-0.269	-0.316	-0.189
	(0.108)	(0.158)	(0.164)	(0.152)
Week 104	-1.137	-0.234	-0.123	-0.300
	(0.106)	(0.150)	(0.149)	(0.152)
p-Value of Test of Joint Significance of Trial Arm Coefficients Relative to:				
AZT Trial Arm		0.016	0.285	0.003
AZT+ddI			0.807	0.281
AZT+ddC				0.032

Note: Standard errors in parentheses. Coefficient estimates in columns (2) – (4) are relative to column (1).

TABLE 2
REPORTED REASON FOR DROPOUT, CONDITIONAL UPON ATTRITION

Reason	Panel A: AZT			Panel B: AZT+ddI		
	Overall	Dropout in Weeks 8-44	Dropout in Weeks 56+	Overall	Dropout in Weeks 8-44	Dropout in Weeks 56+
	(1)	(2)	(3)	(1)	(2)	(3)
Death	0.04	0.04	0.04	0.01	0.01	0.00
Toxicity of Treatment, Patient Request	0.21	0.23	0.13	0.25	0.26	0.22
Request of Patient	0.35	0.31	0.43	0.34	0.31	0.42
Request of Investigator	0.02	0.03	0.01	0.02	0.01	0.04
Lost to Follow-Up	0.16	0.19	0.11	0.20	0.22	0.15
Other	0.22	0.20	0.28	0.18	0.19	0.17
	Panel C: AZT+ddC			**Panel D: ddI**		
Death	0.02	0.01	0.05	0.05	0.05	0.05
Toxicity of Treatment, Patient Request	0.33	0.39	0.17	0.16	0.19	0.12
Request of Patient	0.26	0.21	0.41	0.38	0.30	0.51
Request of Investigator	.03	0.03	0.00	0.05	0.07	0.01
Lost to Follow-Up	0.22	0.23	0.21	0.17	0.21	0.09
Other	0.14	0.13	0.16	0.19	0.18	0.22

TABLE 3
PARAMETER ESTIMATES FOR PROBABILITY OF REMAINING IN TRIAL

Variable		Model		
		(1)	(2)	(3)
γ		0.009 (0.002)	0.009 (0.002)	0.005 (0.002)
Side Effects				
AZT	Mean	0.779 (0.252)	0.640 (0.268)	0.936 (0.287)
	S.D.	1.333 (0.144)	1.465 (0.176)	1.237 (0.161)
AZT+ddI	Mean[1]	0.138 (0.107)	0.134 (0.116)	0.030 (0.111)
	S.D.	1.449 (0.165)	1.580 (0.195)	1.285 (0.167)
AZT+ddC	Mean[1]	-0.108 (0.099)	-0.138 (0.108)	-0.159 (0.107)
	S.D.	1.375 (0.154)	1.505 (0.187)	1.294 (0.163)
ddI	Mean[1]	0.224 (0.104)	0.227 (0.112)	0.178 (0.110)
	S.D.	1.322 (0.154)	1.435 (0.186)	1.162 (0.169)
Screening CD4/100		0.0009 (0.042)	-0.001 (0.045)	-0.081 (0.048)
Prior AZT Use		0.309 (0.071)	0.340 (0.078)	0.354 (0.076)
Symptomatic		-0.158 (0.090)	-0.172 (0.097)	-0.120 (0.092)
Outside Option				
Age		0.028 (0.005)	0.031 (0.005)	0.028 (0.005)
Male		0.155 (0.111)	0.157 (0.120)	0.194 (0.116)
White		0.005 (0.077)	0.004 (0.082)	0.0003 (0.081)
IV Drug User		-0.565 (0.107)	-0.587 (0.118)	-0.569 (0.110)
Homosexual		0.067 (0.090)	0.064 (0.097)	0.065 (0.095)
E[H_{it}]			-0.0004 (0.002)	-0.001 (0.0003)
Calendar Quarter Since 4/91		-0.149 (0.022)	-0.163 (0.025)	-0.090 (0.023)
Prior Beliefs				
Prior Intercept[2]	S.D	72.052 (26.231)	51.684 (22.064)	54.171 (25.512)
Prior Slope (AZT)	S.D.	23.716 (22.500)	26.704 (15.192)	29.786 (23.176)
Prior Slope (not AZT)	S.D.	61.139 (9.571)	59.194 (8.571)	90.450 (27.324)
Discount Factor (annual)		0.95	0.95	0
Log-Likelihood		-3573.4	-3571.7	-3562.7

Note: Standard errors in parentheses. Estimates based on the 2200 subjects in ACTG 175.
[1]Parameter estimate of mean drug effect relative to AZT.
[2] Standard deviations of prior intercepts constrained to be equal for all treatment arms.

TABLE 4
PARAMETER ESTIMATES FOR CD4 PROFILES
Dependent Variable is $CD4_{it} - CD4_{i0}$

Variable		Estimate (Standard Error)
AZT	Mean	4.214 (13.059)
	S.D.	52.010 (2.755)
AZT+ddI	Mean	73.009 (13.434)
	S.D.	90.286 (2.737)
AZT+ddC	Mean	26.661 (13.534)
	S.D.	60.984 (2.925)
ddI	Mean	54.011 (13.226)
	S.D.	67.192 (2.833)
AZT*t	Mean	-8.810 (0.848)
	S.D.	9.699 (0.735)
AZT+ddI*t	Mean	-8.158 (0.972)
	S.D.	14.397 (0.733)
AZT+ddC*t	Mean	-1.875 (0.898)
	S.D.	11.871 (0.673)
ddI*t	Mean	-7.619 (0.826)
	S.D.	11.000 (0.652)
Age		0.076 (0.219)
Male		-5.612 (5.344)
White		7.182 (4.426)
Screening CD4 Count		0.029 (0.023)
Prior Antiretroviral Treatment		-31.825 (3.799)
Symptomatic		-5.944 (5.312)
σ_v		76.461 (0.253)
Log-Likelihood		-74357.934

Note: Standard errors in parentheses. Estimates based on the 2200 subjects in ACTG 175.

TABLE 5
SIMULATIONS OF PATIENT PREFERENCES FOR ACTG 175 THERAPIES

Correlations in Side Effects	0	0.5	-0.5	0.5	-0.5	1	-1
Correlations in Drug Effectiveness	0	0.5	-0.5	-0.5	0.5	1	-1
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
AZT	0.22	0.18	0.22	0.21	0.22	0.11	0.22
	[0.16, 0.29]	[0.06, 0.34]	[0.18, 0.28]	[0.17, 0.28]	[0.19, 0.27]	[0, 0.44]	[0.18, 0.28]
AZT+ddI	0.28	0.31	0.30	0.32	0.30	0.33	0.31
	[0.21, 0.36]	[0.14, 0.49]	[0.24, 0.36]	[0.26, 0.40]	[0.25, 0.36]	[0.05, 0.70]	[0.24, 0.38]
AZT+ddC	0.22	0.18	0.23	0.23	0.23	0.11	0.22
	[0.15, 0.29]	[0.06, 0.35]	[0.18, 0.28]	[0.19, 0.28]	[0.19, 0.28]	[0.01, 0.44]	[0.18, 0.28]
ddI	0.28	0.33	0.25	0.24	0.25	0.45	0.25
	[0.20, 0.36]	[0.14, 0.54]	[0.19, 0.32]	[0.17, 0.34]	[0.19, 0.32]	[0.06, 0.82]	[0.19, 0.32]

Note: Simulations based on 2200 subjects from ACTG 175. Table entries are the fraction of patients preferring the indicated treatment. Entries in brackets are 5^{th} and 95^{th} percentiles of patients preferring the treatment, respectively. 1000 draws are taken per subject.

APPENDIX TABLE 1
SUMMARY STATISTICS BY TREATMENT STATUS

Variable	AZT	AZT+ddI	AZT+ddC	ddI
Age at Baseline	34.97 (8.85)	34.87 (8.50)	35.34 (8.71)	35.15 (8.58)
Male	0.83 (0.38)	0.82 (0.38)	0.84 (0.37)	0.83 (0.38)
White	0.72 (0.45)	0.73 (0.45)	0.71 (0.45)	0.70 (0.46)
Symptomatic HIV Infection	0.18 (0.38)	0.19 (0.39)	0.19 (0.39)	0.18 (0.38)
Screening CD4 Count	346.0 (83.4)	343.2 (84.4)	350.4 (82.7)	344.0 (84.1)
Prior AZT Therapy	0.56 (0.50)	0.56 (0.50)	0.55 (0.50)	0.57 (0.50)
IV Drug User at Baseline	0.13 (0.34)	0.15 (0.35)	0.15 (0.36)	0.14 (0.34)
Homosexual at Baseline	0.65 (0.48)	0.66 (0.47)	0.68 (0.47)	0.68 (0.47)
Number of Observations	560	543	546	551
Number of Sites	89			

Note: Standard deviations in parentheses.

REFERENCES

Ackerberg, D. (2003). "Advertising, Learning, and Consumer Choice in Experience Good Markets: A Structural Empirical Examination," *International Economic Review*, forthcoming.

Barnard, J., Frangakis, C., Hill, J., and Rubin, D. (2003). "A Principal Stratification Approach to Broken Randomized Experiments: A Case Study of Vouchers in New York City," *Journal of the American Statistical Association*, 98, 299-323.

Burtless, G. (1995). "The Case for Randomized Field Trials in Economic and Policy Research," *Journal of Economic Perspectives*, 9, 63-84.

Boscardin, W., Taylor, J., and Law, N. (1998). "Longitudinal Models for AIDS Marker Data," *Statistical Methods in Medical Research*, 7, 13-27.

Ching, A. (2000). "Dynamic Equilibrium in the U.S. Prescription Drug Market after Patent Expiration," mimeo, Department of Economics, University of Minnesota.

Crawford, G. and Shum, M. (2005). "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 73, 1137-1174.

Duggan, M. and Evans, W. (2006). "Estimating the Impact of Medical Innovation: A Case Study of HIV Antiretroviral Treatments," mimeo, Department of Economics, University of Maryland.

Efron, B. and Feldman, D. (1991). "Compliance as an Explanatory Variable in Clinical Trials," *Journal of the American Statistical Association*, 86, 9-17.

Erdem, T. and Keane, M. (1996). "Decision-Making Under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 15, 1-20.

Fisher, R. (1935). The *Design of Experiments*. London: Oliver and Boyd.

Frangakis, C. and Rubin, D. (1999). "Addressing the Complications of Intention-to-Treat Analysis in the Combined Presence of All-or-None Treatment-Noncompliance and Subsequent Missing Outcomes," *Biometrika*, 86, 365-379.

Gelman, A., Carlin, J., Stern, H. and Rubin, D. (2004). Bayesian *Data Analysis, 2nd ed.* Chapman & Hall: New York.

Greenberg, D. and Shroder, M. (1997). *The Digest of Social Experiments, 2nd ed.* The Urban Institute Press: Washington, D.C.

Hajivassiliou, V. and Ruud, P. (1994). "Classical Estimation Methods for LDV Models Using Simulation," in *Handbook of Econometrics, Vol. 4*, eds. R Engle and D. McFadden, Elsevier Science B.V., New York, 2383-2441.

Hammer, S., et. al. (1996). "A Trial Comparing Nucleoside Monotherapy with Combination Therapy in HIV-Infected Adults with CD4 Cell Counts from 200 to 500 per Cubic Millimeter," *New England Journal of Medicine*, 335, 1081-1090.

Heckman, J. (1974). "Shadow Prices, Market Wages, and Labor Supply," *Econometrica*, 42, 679-694.

Heckman, J. (1992). "Randomization and Social Policy Evaluation," in *Evaluating Welfare and Training Programs*, eds. C. Manski and I. Garfinkel, Harvard University Press, Cambridge, 201-230.

Heckman, J., Hohmann, N., and Smith, J. (2000). "Substitution and Dropout Bias in Social Experiments: A Study of an Influential Social Experiment," *Quarterly Journal of Economics*, CXV, 651-694.

Heckman, J. and Robb, R. (1985). "Alternative Models for Evaluating the Impact of Interventions," in *Longitudinal Analysis of Labor Market Data*, eds. J. Heckman and B. Singer, Cambridge University Press, New York, 156-245.

Heckman, J. and Smith, J. (1995). "Assessing the Case for Social Experiments," *Journal of Economic Perspectives*, 9, 85-110.

Heckman, J. and Smith, J. (1998). "Evaluating the Welfare State," in *Econometrics and Economic Theory in the 20th Century: The Ragner Frisch Centennial Symposium*, e. S. Strom, Cambridge University Press: New York, 241-318.

Heckman, J. and Smith, J., and Taber, C. (1998). "Accounting for Dropouts in the Evaluation of Social Programs," *Review of Economics and Statistics*, LXXX, 1-14.

Jovanovic, B. (1979). "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, 972-990.

Keane, M. and Wolpin, K. (1997). "The Career Decisions of Young Men," *Journal of Political Economy*, 105, 473-522.

Malani, A. (2006). "Identifying Placebo Effects with Data from Clinical Trials," *Journal of Political Economy*, 114, 236-256.

Manski, C. (1996). "Learning About Treatment Effects from Experiments with Random Assignment of Treatments," *Journal of Human Resources*, 31, 707-733.

Miller, R. (1984). "Job Matching and Occupational Choice," *Journal of Political Economy*, 92, 1086-1120.

Office of Technology Assessment. (1983). "The Impact of Randomized Clinical Trials on Health Policy and Medical Practice: Background Paper." OTA-BP-H-22. Washington, D.C.

Pakes, A. (1986). "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54, 755-784.

Palaca, J. (1989). "AIDS Drug Trials Enter New Age," *Science Magazine*, October 6, 19-21.

Philipson, T. and Desimone, J. (1997). "Experiments and Subject Sampling," *Biometrika*, 84, 619-631.

Philipson, T. and Hedges, L. (1998). "Subject Evaluation in Social Experiments," *Econometrica*, 66, 381-408.

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

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

Scharfstein, D., Rotnitzky, A., and Robins, J. (1999). "Adjusting for Non-Ignorable Drop-Out Using Semiparametric Nonresponse Models (with discussion)," *Journal of the American Statistical Association*, 94, 1096-1146.

APPENDIX

A. Bayesian Updating of Beliefs

We assume that subjects' beliefs of their CD4 counts over the course of the trial, analogous to equation (11), follow the linear relationship

$$H_{it} = H_{i0} + \mu_{it} + \upsilon_{it}$$

where $\mu_{it} = \sum_r R_{ir}\mu_{irt}$, and μ_{irt} is the treatment effect in period *t* conditional on *i* belonging to treatment arm *r*, and $R_{ir} = 1$ if subject *i* belongs to trial arm *r*, and 0 otherwise. Since a subject belongs to only one treatment arm, $\sum_r R_{ir} = 1$. The prior belief of effectiveness, conditional on *i* belonging to *r* and with information set I_{it}, is $\mu_{irt} \sim N(\mu_{irt}^t, \sigma_{i,\mu,r,t}^2)$.

The conditional expected utility function of remaining in trial in period *t* is[25]

$$E[U_{it}^T(H_{it}) \mid I_{it}] = \int U_{it}^T(H_{it}) dF(\mu_{irt}, R_{ir}, \forall r; \upsilon_{it} \mid I_{it})$$
$$= \int U_{it}^T(H_{i0} + \mu_{it} + \upsilon_{it})\, dF(\upsilon_{it} \mid \mu_{irt}, R_{ir}, \forall r; I_{it})\, dF(\mu_{irt}, R_{ir}, \forall r \mid I_{it})$$
$$= \int E_\upsilon[U_{it}^T(H_{i0} + \mu_{it} + \upsilon_{it}) \mid \mu_{it}] dF(\mu_{irt} \mid R_{ir}, I_{it}) \cdot dF(R_{ir}, \forall r \mid I_{it})$$

The last equality comes from the assumption that the distribution of υ_{it} is independent of treatment arm *r*. The conditional distribution $F(\mu_{irt} \mid R_{ir}, \forall r; I_{it})$ can be written as $\sum_r R_{ir} \cdot N(\mu_{irt}^t, \sigma_{i,\mu,r,t}^2)$. Hence we have the following expression

$$E[U_{it}^T(H_{it}) \mid I_{it}] = \sum_r \int R_{ir} \cdot E_\mu[E_\upsilon[U_{it}^T(H_{i0} + \mu_{it} + \upsilon_{it}) \mid \mu_{it}] \mid R_{ir} = 1, I_{it}] dF(R_{ir}, \forall r \mid I_{it})$$

Based on the multinomial distribution assumption for R_{ir}, we can further write

$$(A.1) \quad E[U_{it}^T(H_{it}) \mid I_{it}] = \sum_r p_{ir}^t \cdot E_\mu[E_\upsilon[U_{it}^T(H_{i0} + \mu_{it} + \upsilon_{it}) \mid \mu_{it}] \mid R_{ir} = 1, I_{it}]$$

where

$$(A.2) \quad E_\mu[E_\upsilon[U_{it}^T(H_{i0} + \mu_{it} + \upsilon_{it}) \mid \mu_{it}] \mid R_{ir} = 1, I_{it}]$$
$$= -\exp(-\gamma(H_{i0} + \mu_{irt}^t) + \frac{\gamma^2}{2}(\sigma_\upsilon^2 + \sigma_{i,\mu,r,t}^2).$$

This is analogous to equation (15) in the paper.

[25] For simplicity we omit the effect of "side effects" on utility here.

In period t subject i forms his expected utility in period $s > t$. As above, we can derive the conditional expected utility function of remaining in trial in period $s > t$ as

$$E[U_{it}^{T}(H_{is}) \mid I_{it}]$$
$$= \sum_{r} p_{ir}^{t} \cdot E_{\mu}[E_{\upsilon}[U_{it}^{T}(H_{i0} + \mu_{is} + \upsilon_{is}) \mid \mu_{is}] \mid R_{ir} = 1, I_{it}]$$

Conditional on i belonging to trial r, $\mu_{irt} = \mu_{0ir} + \mu_{1ir}t$. The prior beliefs in the first period for μ_{0ir} and μ_{1ir} are given by

$$\begin{pmatrix} \mu_{0ir} \\ \mu_{1ir} \end{pmatrix} \sim N\left(\begin{pmatrix} \mu_{00r} \\ \mu_{10r} \end{pmatrix}, \Sigma_{\mu 0r} \right), \ \Sigma_{\mu 0r} = \begin{pmatrix} \sigma_{\mu 00r}^{2} & 0 \\ 0 & \sigma_{\mu 10r}^{2} \end{pmatrix}.$$

Let

$$A_{t-1} = \begin{pmatrix} 1,1,...,1 \\ 1,2,...,t-1 \end{pmatrix}',$$

and

$$H_{it-1} = (H_{i1} - H_{i0}, ..., H_{it-1} - H_{i0})'.$$

Using the Bayesian updating formula, the posterior means and variances of (μ_{0ir}, μ_{1ir}) in period t are

$$(A.3) \quad E\left[\begin{pmatrix} \mu_{0ir} \\ \mu_{1ir} \end{pmatrix} \mid I_{it} \right] = \mu_{irt}^{t} = D_{rt} d_{rt},$$

$$(A.4) \quad \text{var}\left[\begin{pmatrix} \mu_{0ir} \\ \mu_{1ir} \end{pmatrix} \mid I_{it} \right] = \sigma_{i,\mu,r,t}^{2} = D_{rt},$$

where

$$D_{r,t} = \left(\frac{A_{t-1}' A_{t-1}}{\sigma_{\upsilon}^{2}} + \Sigma_{\mu 0r}^{-1} \right)^{-1},$$

and

$$d_{r,t} = \frac{A_{t-1}' H_{i,t-1}}{\sigma_{\upsilon}^{2}} + \Sigma_{\mu 0r}^{-1} \cdot \begin{pmatrix} \mu_{10r} \\ \mu_{20r} \end{pmatrix}.$$

Hence we have equation (13) for μ_{irt}^{t} and $\sigma_{i,\mu,r,t}^{2}$.

B. Derivation of $V_i(I_{it}, t)$

To derive L_i in (17) and (18) we need to compute $V_i(I_{it}, t) = \max\{V_i^{T}(I_{it}, t), V_i^{o}(I_{it}, t)\}$ in (5).

First define a discrete decision process $D_{itu} = \{d_{it} = 1, ..., d_{it+u} = 1, d_{it+u+1} = 0, ..., d_{iT} = 0\}$, where d_{is} is a discrete decision in period *s*, and $d_{is}=1$ means that subject *i* stays in the experiment and $d_{is}=0$ means dropout. D_{itu} represents the decision that subject *i* stays in the experiment until $t+u \leq T$, and drops out in period $t+u+1$.

Define

$$(A.5) \quad \hat{V}_i^T(D_{itu}; I_{it}, t) = E[U_{it}^T - U_{it}^0 \mid I_{it}] + \beta E[U_{it+1}^T - U_{it+1}^0 \mid I_{it}] + ... + \beta^u E[U_{it+u}^T - U_{it+u}^0 \mid I_{it}]$$

Then

$$(A.6) \quad V_i(I_{it}, t) = \max\{\hat{V}_i^T(D_{it0}; I_{it}, t), \hat{V}_i^T(D_{it1}; I_{it}, t), ..., \hat{V}_i^T(D_{itT-t}; I_{it}, t), 0\}$$

Analogous to equation (15), $E[U_{it}^T - U_{it}^0 \mid I_{it}] = \sum_r p_{ir}^t \cdot E[U_{irt}^T - U_{irt}^0 \mid R_{ir} = 1, I_{it}]$, where p_{ir}^t is the prior mean for $R_{i,r}$ in period *t*, and

$$(A.7) \quad E[U_{irt}^T - U_{irt}^0 \mid R_{ir} = 1, I_{it}] \equiv \hat{u}_{irt}$$
$$= -\exp(-\gamma * (H_{i0} + \mu_{irt}^t) + \frac{\gamma^2}{2}(\sigma_\upsilon^2 + \sigma_{i,\mu,r,t}^2)) - x_{Si}\alpha_i - \varepsilon_{iSt},$$

Based on (A.7) and using iterative expectations we obtain, for period $t+1$, $E[U_{it+1}^T - U_{it+1}^0 \mid I_{it}] = \sum_r p_{ir}^t \cdot E[U_{irt+1}^T - U_{irt+1}^0 \mid R_{ir} = 1, I_{it}]$, and

$$E[U_{irt+1}^T - U_{irt+1}^0 \mid R_{ir} = 1, I_{it}] \equiv \hat{u}_{irt+1}$$
$$= E[E[U_{irt+1}^T - U_{irt+1}^0 \mid R_{ir} = 1, I_{it+1}] \mid R_{ir} = 1, I_{it}]$$
$$= -E[\exp(-\gamma * (H_{i0} + \mu_{irt+1}^{t+1}) + \frac{\gamma^2}{2}(\sigma_\upsilon^2 + \sigma_{i,\mu,r,t+1}^2)) \mid I_{it}] - x_{Si}\alpha_i,$$

since $E[\varepsilon_{iSt+1} \mid I_{it}] = 0$ by assumption. After some manipulation we obtain

$$(A.8)\quad \hat{u}_{irt+1} = -\exp\{-\gamma * [H_{i0} + \begin{pmatrix}1\\ t+1\end{pmatrix}' D_{rt+1} \cdot (\frac{A_{t-1}' H_{it-1}}{\sigma_\upsilon^2} + \Sigma_{\mu 0r}^{-1} \cdot \mu_{0r})$$

$$+ \frac{\begin{pmatrix}1\\ t+1\end{pmatrix}' D_{rt+1} \cdot \begin{pmatrix}1\\ t+1\end{pmatrix}}{\sigma_\upsilon^2} \cdot \begin{pmatrix}1\\ t\end{pmatrix}' D_{rt} \cdot (\frac{A_{t-1}' H_{it-1}}{\sigma_\upsilon^2} + \Sigma_{\mu 0r}^{-1} \cdot \mu_{0r})]$$

$$+ \frac{\gamma^2}{2}\left[\left(\sigma_\upsilon^2 + \begin{pmatrix}1\\ t+1\end{pmatrix}' D_{rt+1} \cdot \begin{pmatrix}1\\ t+1\end{pmatrix}\right) + \left(\frac{\begin{pmatrix}1\\ t+1\end{pmatrix}' D_{rt+1} \cdot \begin{pmatrix}1\\ t+1\end{pmatrix}}{\sigma_\upsilon^2}\right)^2 \cdot \left(\begin{pmatrix}1\\ t\end{pmatrix}' D_{rt} \cdot \begin{pmatrix}1\\ t\end{pmatrix} + \sigma_\upsilon^2\right)\right]\}$$

$$- x_{Si}\alpha_i,$$

where $\Sigma_{\mu 0r}^{-1}$, A_{t-1}, and H_{it-1} are defined as above, $\mu_{0r} = \begin{pmatrix}\mu_{10r}\\ \mu_{20r}\end{pmatrix}$, and $D_{rt} = \left(\frac{A_{t-1}' A_{t-1}}{\sigma_\upsilon^2} + \Sigma_{\mu 0r}^{-1}\right)^{-1}$, $D_{rt+1} = \left(\frac{A_t' A_t}{\sigma_\upsilon^2} + \Sigma_{\mu 0r}^{-1}\right)^{-1}$.

We can use the same method and iterate for k periods. Let $E[U_{it+k}^T - U_{it+k}^0 \mid R_{ir} = 1, I_{it}] \equiv \hat{u}_{irt+k}$. For simplicity of notation we omit the subscript r below. By method of induction

$$(A.9)\quad \hat{u}_{it+k} = -\exp\left\{-\gamma \cdot \left[H_{i0} + \psi_{tk}' \tilde{D}_{tk} \cdot \left(\frac{A_{t-1}' H_{it-1}}{\sigma_\upsilon^2} + \Sigma_{\mu 0}^{-1} \cdot \mu_0\right)\right] + \frac{\gamma^2}{2}\left[(\psi_{tk})^{2'} (\phi_{tk} + \sigma_\upsilon^2)\right]\right\} - x_{Si}\alpha_i,$$

where

$$\tilde{D}_{tk}=\begin{pmatrix}\tilde{D}_{it+k}\\ \tilde{D}_{it+k-1}\\ \dots\\ \tilde{D}_{it}\end{pmatrix};\quad \tilde{D}_{it+l}=\begin{pmatrix}1\\ t+l\end{pmatrix}' D_{t+l}=\begin{pmatrix}1\\ t+l\end{pmatrix}'\left(\frac{A_{t+l-1}' A_{t+l-1}}{\sigma_\upsilon^2}+\Sigma_{\mu 0}^{-1}\right)^{-1};$$

$$\phi_{tk}=\begin{pmatrix}\sigma^2_{\mu,t+k}\\ \sigma^2_{\mu,t+k-1}\\ \dots\\ \sigma^2_{\mu,t}\end{pmatrix};\quad \sigma^2_{\mu,t+l}=\tilde{D}_{it+l}{}^2\cdot\begin{pmatrix}1\\ t+l\end{pmatrix};$$

$$\psi_{tk}=\begin{pmatrix}1\\ \hat{\psi}_{tk}[1]\\ \hat{\psi}_{tk}[1]+\hat{\psi}_{tk}[2]\\ \dots\\ \sum_{l=1}^{k}\hat{\psi}_{tk}[l]\end{pmatrix};\quad \psi_{tk}=\begin{pmatrix}\tilde{\psi}_{tk}[1]\\ \tilde{\psi}_{tk}[1]\times\tilde{\psi}_{tk}[2]\\ \dots\\ \tilde{\psi}_{tk}[1]\times\tilde{\psi}_{tk}[2]\times\dots\times\tilde{\psi}_{tk}[k]\end{pmatrix};\quad \tilde{\psi}_{tk}=\begin{pmatrix}\sigma^2_{\mu,t+k}/\sigma^2_\upsilon\\ \sigma^2_{\mu,t+k-1}/\sigma^2_\upsilon\\ \dots\\ \sigma^2_{\mu,t+1}/\sigma^2_\upsilon\end{pmatrix},$$

and *[l]* denotes the *l-th* element in a *k×1* vector.

From above

$$\begin{aligned}&\hat{V}_i^T(D_{itk};I_{it},t)\\ &=\sum_r p_{ir}^t\cdot(\hat{u}_{irt}+\beta\hat{u}_{irt+1}+\dots+\beta^k\hat{u}_{irt+k})-\frac{1-\beta^{k+1}}{1-\beta}\cdot x_{Si}\alpha_i-\varepsilon_{iSt}\\ &=W_i^T(D_{itk};I_{it},t)-\frac{1-\beta^{k+1}}{1-\beta}\cdot x_{Si}\alpha_i-\varepsilon_{iSt}\end{aligned}$$

and therefore

$$(A.10)\quad V_i(I_{it},t)=\max\left\{\max\begin{Bmatrix}W_i^T(D_{it0};I_{it},t)-x_{Si}\alpha_i, W_i^T(D_{it1};I_{it},t)-(1+\beta)x_{Si}\alpha_i,\\ \dots, W_i^T(D_{itT-t};I_{it},t)-\frac{1-\beta^{T-t+1}}{1-\beta}x_{Si}\alpha_i\end{Bmatrix}+\varepsilon_{iSt},0\right\}$$

We use simulated maximum likelihood to evaluate the likelihood function in equation (18), using a frequency simulator drawing ε_{iSt} in (A.10) from the assumed normal distribution. It is well-known that, for the consistency of the estimator, the number of simulations must increase with the number of observations in data; in practice the number is fixed at 1,000 in our model estimation.

FIGURE 1
AVERAGE CD4 COUNT PROFILES FOR HYPOTHETICAL RCT

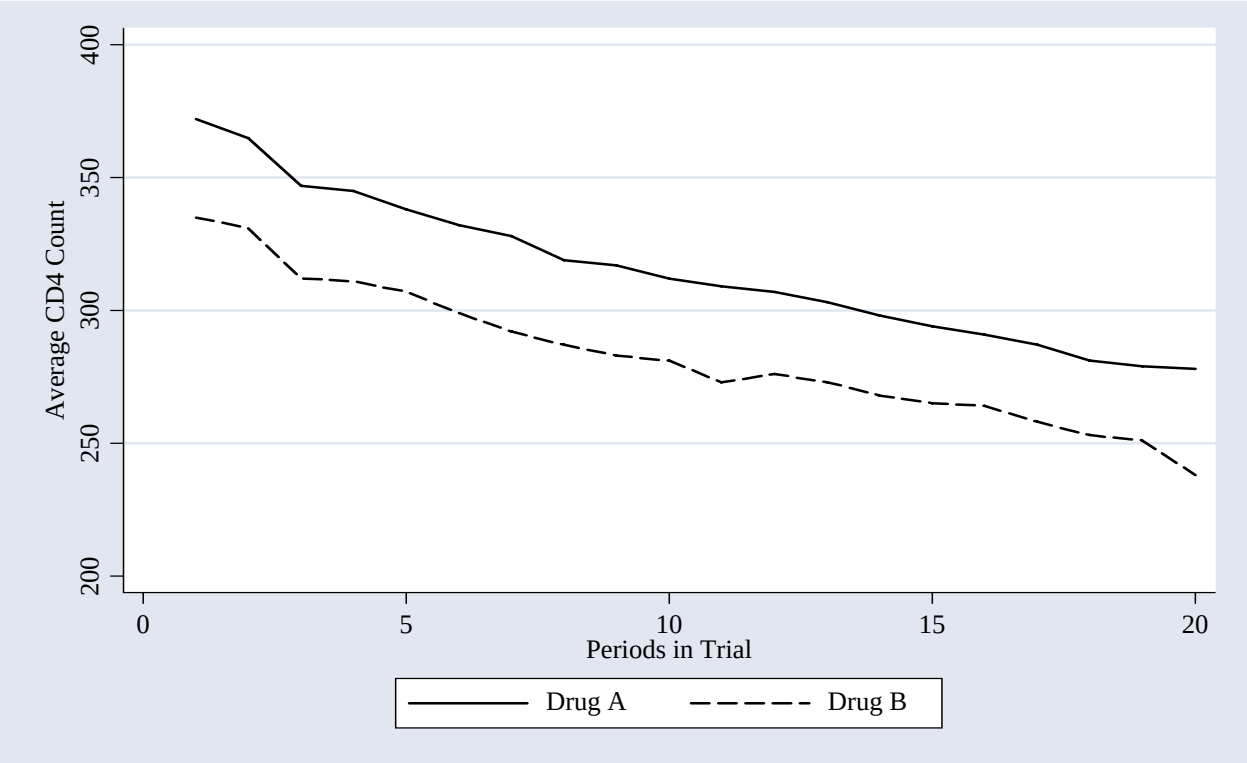

FIGURE 2
PROBABILITY OF REMAINING IN HYPOTHETICAL TRIAL

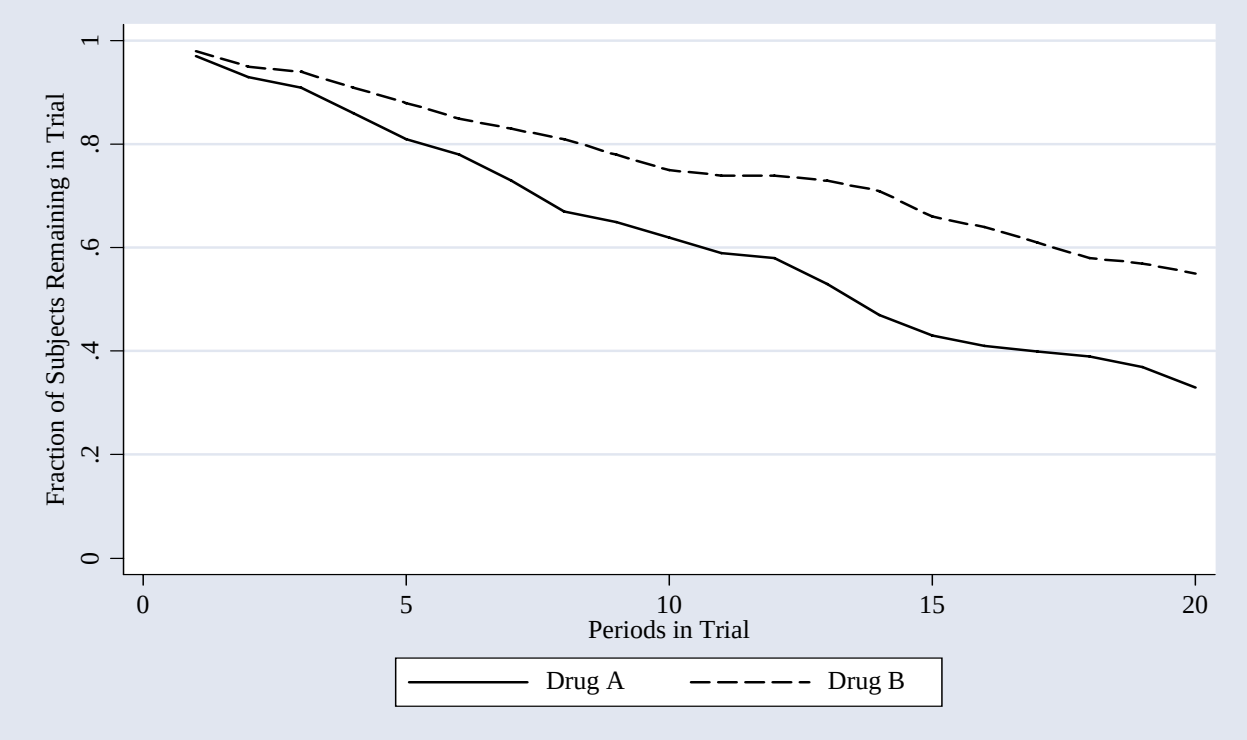

FIGURE 3
CD4 COUNT PROFILES, BY TREATMENT ARM

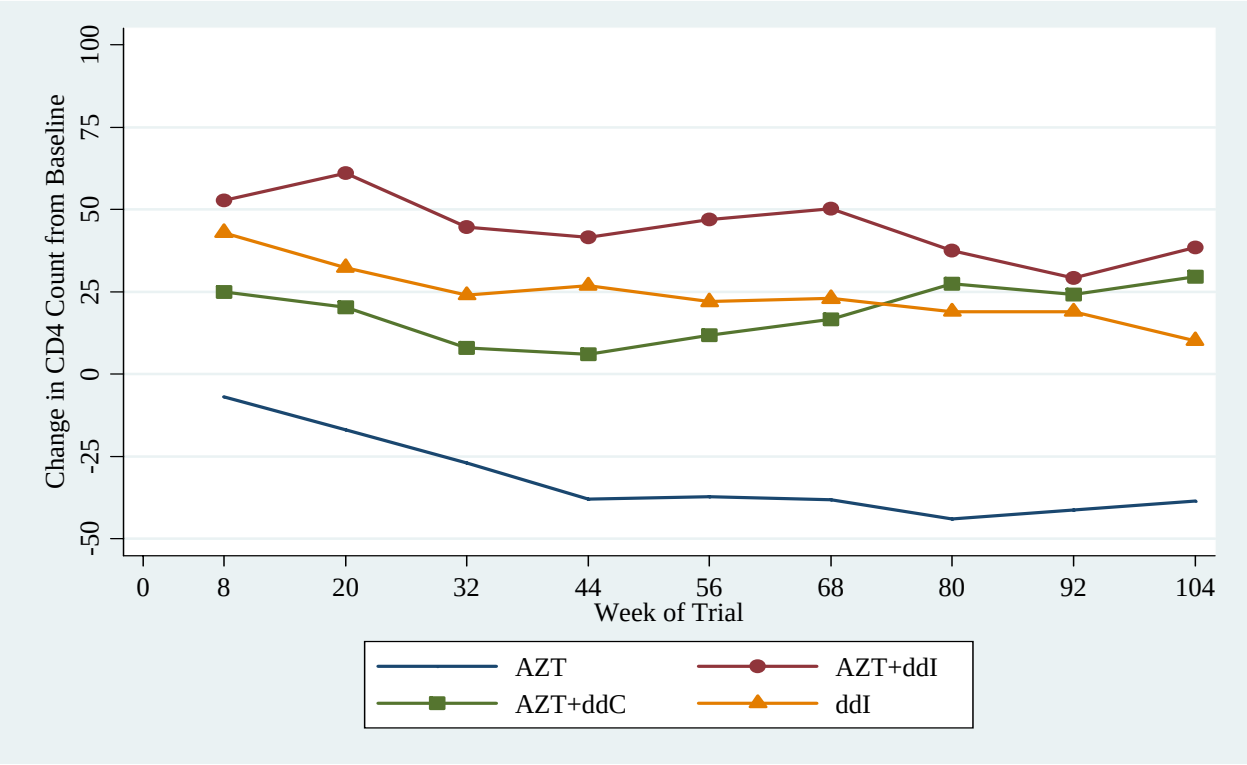

FIGURE 4
PROBABILITY OF REMAINING IN ACTG 175, BY TREATMENT ARM

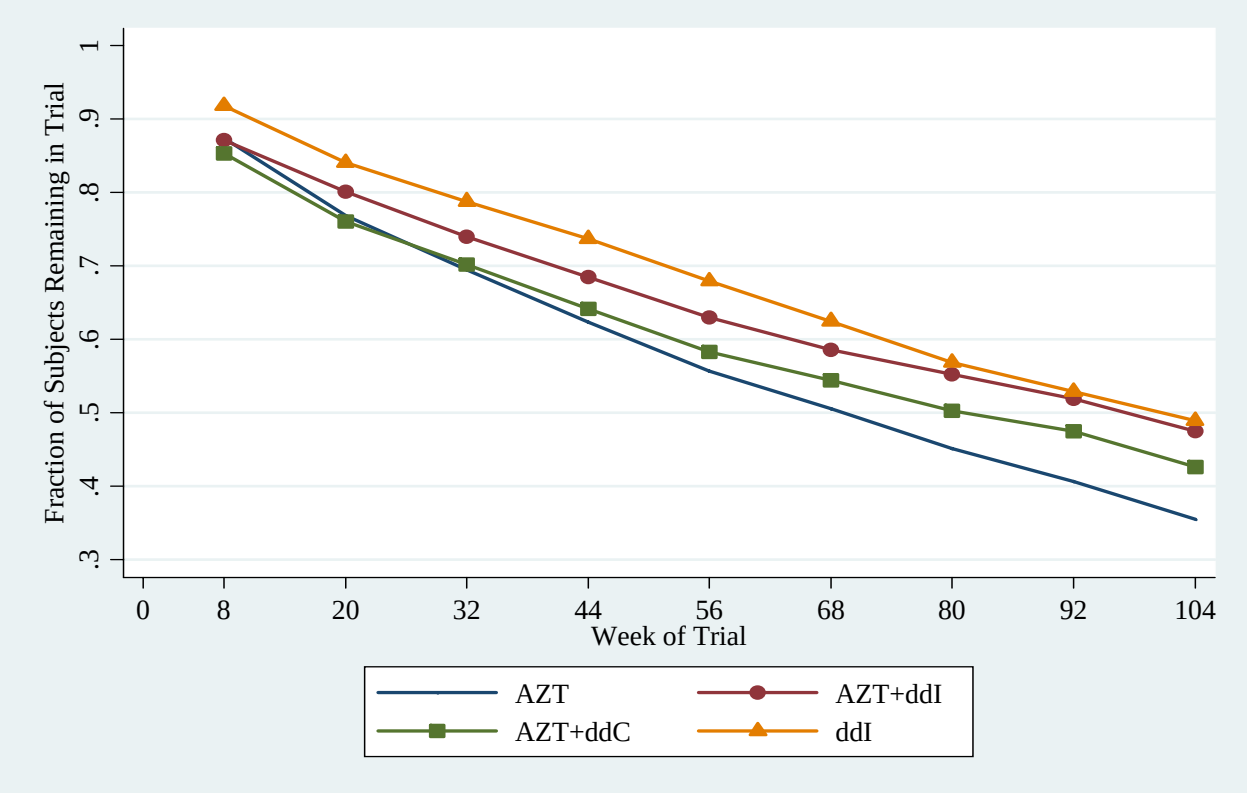

FIGURE 5
CD4 COUNT PROFILES, BY TREATMENT AND ATTRITION GROUP

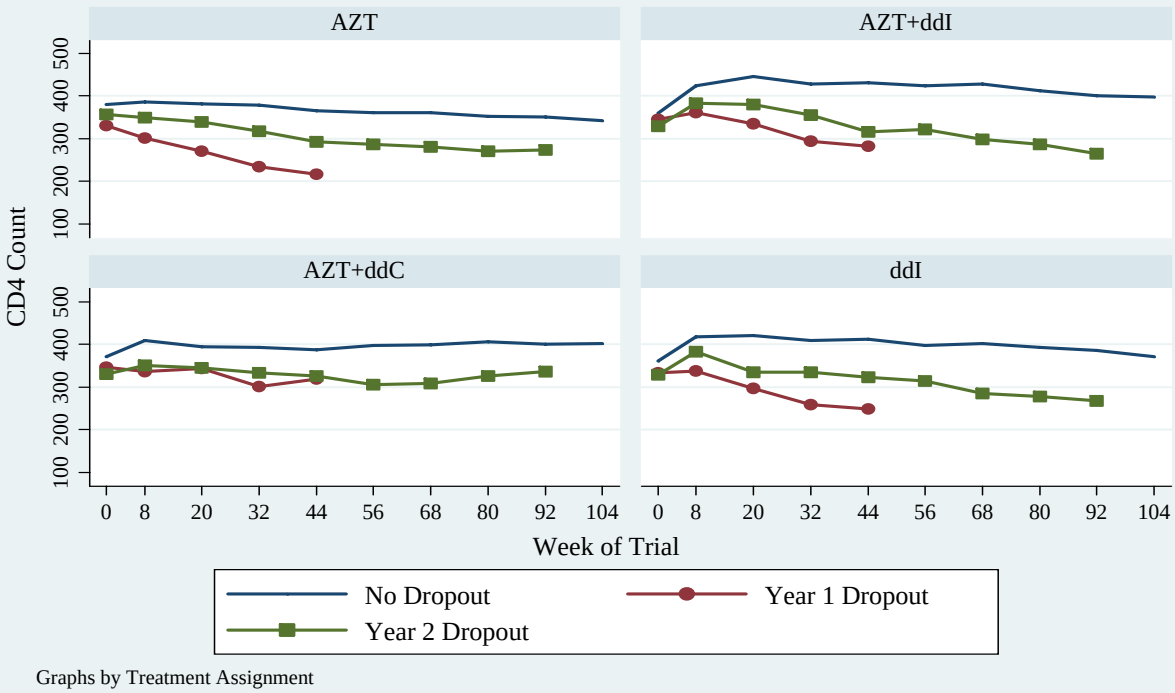

FIGURE 6
MARGINAL UTILITY OF H_{it} BY CD4 COUNT

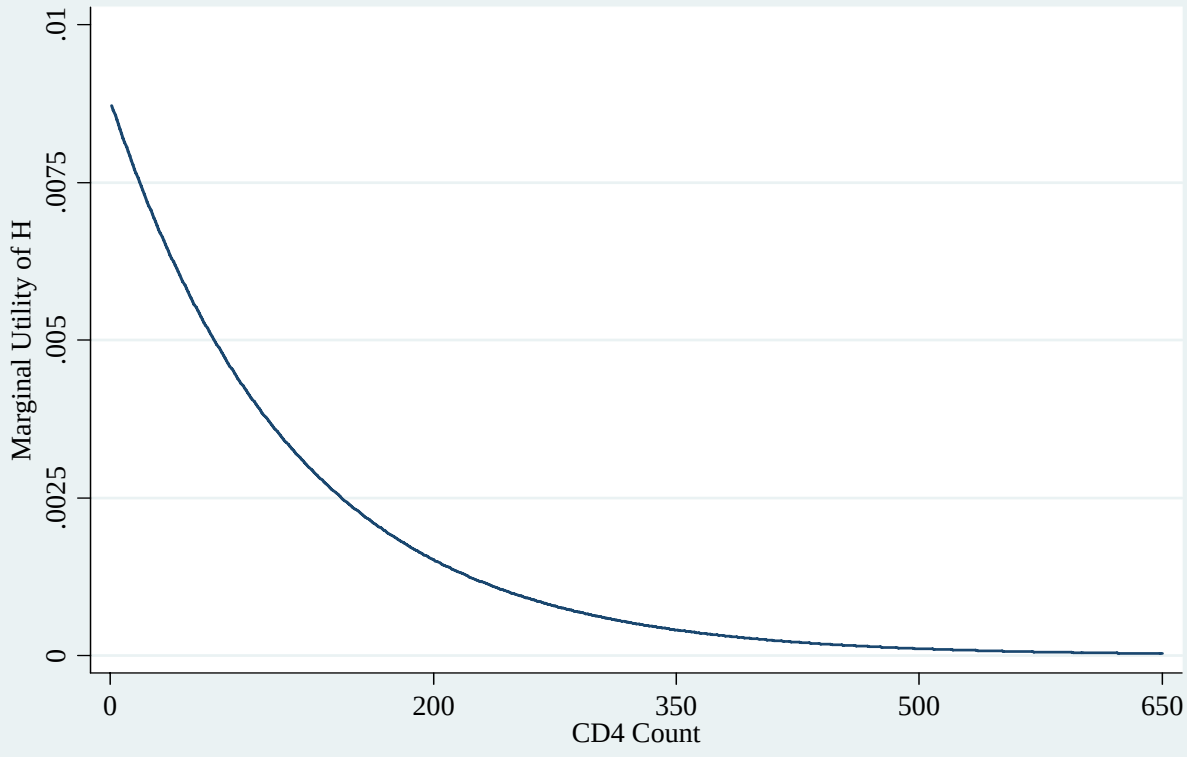

FIGURE 7
OBSERVED AND PREDICTED DROPOUT PROBABILITIES

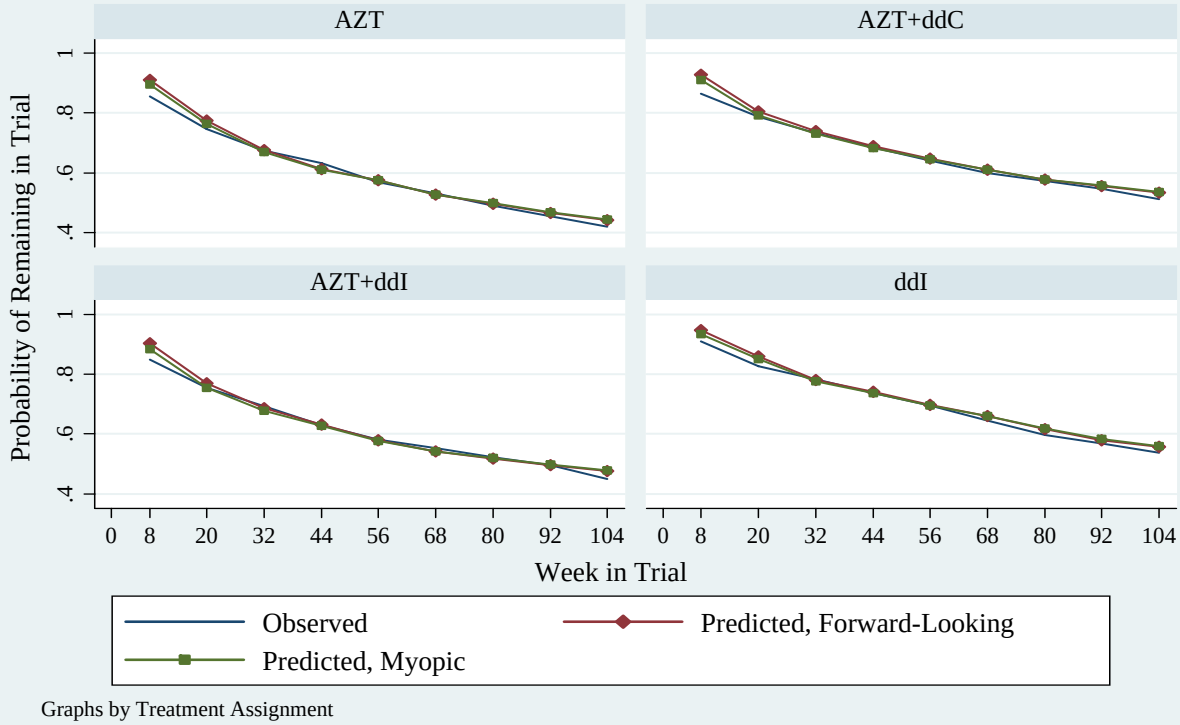

FIGURE 8
CONVERGENCE BETWEEN ACTUAL AND EXPECTED CD4 COUNTS

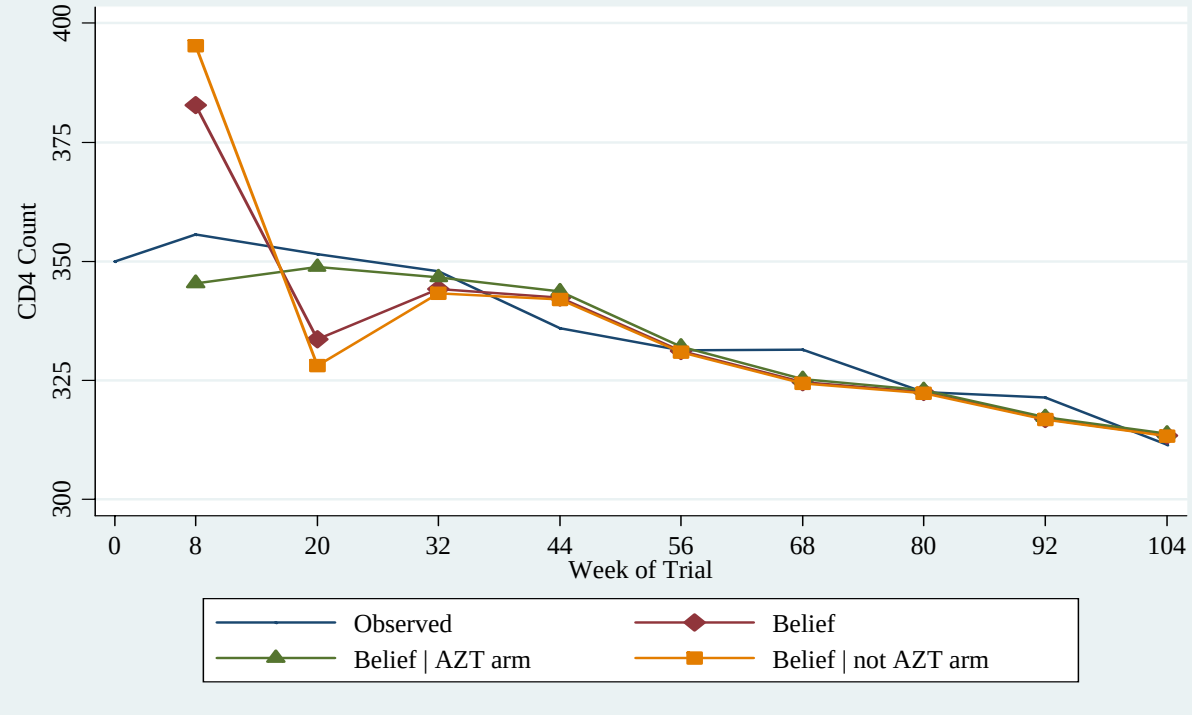

FIGURE 9
DISTRIBUTION OF SUBJECT SIDE EFFECTS, BY TREATMENT GROUP

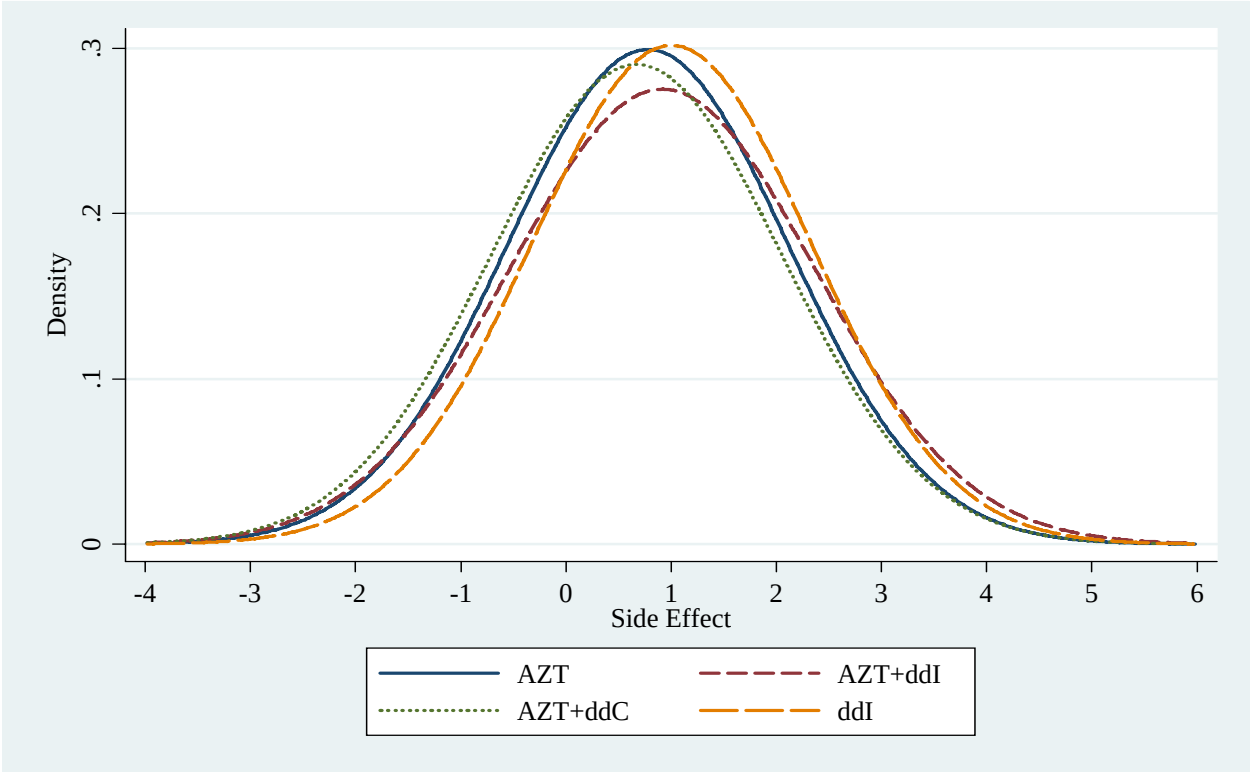

FIGURE 10
SIMULATED IMPACT OF INFORMATION ON TREATMENT CHOICE