

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@olin.wustl.edu; hamiltonb@olin.wustl.edu

March 2006

ABSTRACT

Randomized experiments are viewed as the gold standard for the treatment evaluation, but many experiments are plagued by attrition or non-compliance, even among subjects receiving the more effective treatment. This paper constructs an economic model of decision-making in which individuals make utility maximizing choices that provides a 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. Utility is a function of both the "publicly observed" outcomes that are typically the focus of evaluation studies, and treatment side effects that are the private information of the subject. Participants enter the experiment uncertain of treatment effectiveness (and often the treatment received), and update their prior beliefs over the course of the experiment when deciding whether to drop out. We use the framework to analyze an influential AIDS clinical trial, ACTG 175, which has been used to tout the benefits of combination therapies for AIDS over the use of AZT alone. However, our analysis indicates that for many subjects, AZT yields the highest level of utility, despite having the smallest impact on the publicly observed outcome of the study, the patient's CD4 count. Significant and rapid 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. 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.

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, 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 held as the gold standard for treatment evaluation. In the standard evaluation of the randomized experiment, the outcomes of subjects randomly assigned to a treatment group are observed over a set period of time. Comparisons of outcomes across groups are then used to recover treatment effects of interest. It is argued that the primary advantage of this approach is that randomization eliminates the possibility that bias may be introduced into the estimation if 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 dropout 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 RE;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 RE, 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 RE analyzing the impact of two drugs, A and B, for the treatment of AIDS (the example is hypothetical). As is common in many AIDS REs, 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 participants over the

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 RE, 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

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 RE.

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 REs. 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 that used by

outside treatments for those available in the RE.

⁹ See Scharfstein et al (1999) for an example of this approach in the context of a randomized clinical trial.

Miller (1984) to study occupational choice, where workers have prior beliefs concerning the uncertain match quality of different occupations. After choosing a job, workers 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, by construction subjects in REs 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) 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 RE.

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/mm³: zidovudine (AZT); didanosine (ddI); zidovudine plus didanosine (AZT+ddI); and zidovudine plus zalcitabine (AZT+ddC). At the time, AZT was the standard treatment for HIV in many patient subpopulations. CD4 counts are a widely used marker for the status of an individual's immune system, and hence the progression of AIDS in the patient.

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); Akerberg (2003)); and the purchase of pharmaceuticals where individuals differ in their reactions to drugs (Crawford and Shum (2003)). However, these studies do not address the issue of uncertainty regarding the type of treatment being received. Malani (2003) discusses the importance of the learning process in understanding placebo effects in medical clinical trials.

¹¹ Our model differs from Miller (1984) in that we utilize data on the publicly observed outcome, CD4 counts, while

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 by 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) and Currie and Park (2000) assume 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.¹² 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

he does not use wage data associated with occupations.

¹² Crawford and Shum (2005) analyze the demand for anti-ulcer medications in a non-RE setting. Ferrall (2002)

AZT to all other therapies. In fact, the fraction of subjects who prefer AZT is about the same (20%) 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. Because of random assignment to treatment arm and the double-blind design, we are able to infer that attrition at the onset of ACTG 175 is driven primarily by side effects, since subjects have prior beliefs (based on rational expectations) concerning effectiveness that are common across treatments and are balanced in terms of unobserved characteristics and preferences. Therefore, early on in the trial, subjects with low CD4 counts may remain while those with high CD4 counts may drop out. 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 any other treatment.

The remainder of the paper first presents initial summary evidence on the importance of attrition and side effects, and their impact on the observed evolution of CD4 counts over the course of ACTG 175. Section 3 then describes the structural model of dropout behavior, while Section 4 describes the econometric implementation of the economic models. The parameter estimates are presented in Section 5, along with a discussion of the side effects and direct effects of each treatment. Section 6 simulates patient treatment choices using the parameter estimates, constructs the “market shares” reflecting preferences for each drug, and investigates the importance of information and learning on drug choice. A short conclusion summarizes the findings and suggests avenues for future research.

considers the implications of forward-looking behavior for the interpretation of social and natural experiments.

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. Subjects were recruited from 43 AIDS Clinical Trials Units and 9 National Hemophilia Foundation sites in the United States and Puerto Rico. A total of 89 sites were used in the trial. 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 the 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 recorded 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).

After excluding observations with missing covariate or outcome data, the analysis sample consists of 2200 subjects treated at 89 different trial sites, 560 receiving AZT, 543 receiving AZT+ddI, 546 receiving AZT+ddC, and 551 receiving ddI. 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 $t = 0, 1, 2, \dots, 9$, corresponding to weeks $w = 0, 8, 20, \dots, 104$. The two-year window was chosen because the clinical endpoint of the trial for all subjects occurred at least two years after baseline.

2.1 CD4 PROFILES

The publicly observed outcome measure used in this study is the change in the CD4 count between week w and baseline week 0. This specification allows us to determine whether the rise in CD4 counts often observed 8 weeks after baseline continues to persist into subsequent weeks. 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.

To assess the significance of the differences in CD4 profiles, Table 1 presents estimates from a regression of the change in CD4 counts relative to baseline on indicator variables for treatment and treatment-specific linear time trends. Specification (1) shows that subjects receiving combination therapy or ddI alone have significantly higher CD4 counts at all time periods in the trial. Moreover, the intercept of the average CD4 profiles of those receiving ddI alone or in combination with AZT is significantly higher than that for subject taking AZT+ddC. However, unlike the other 3 groups, the time trend for AZT+ddC patients is positive, suggesting that the health status of these individuals may actually improve over the course of the trial. Specification (2) presents estimates for a specification that also includes controls for age, gender, whether the subject is white, the subject's screening CD4 count, whether the subject received ZDV prior to the start of the clinical trial, and whether the subject is HIV symptomatic. As might be expected from the random assignment of trial participants, the estimated treatment effects do not change when these controls are included in the model.

2.2 DROPOUT BEHAVIOR IN ACTG 175

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 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 in ACTG 175, since only about 35%-50% of subjects continue in the trial through week 104. Consequently, the OLS estimates of the impact of alternative treatments on CD4 profiles presented in columns (1) and (2) of Table 1 may be biased. 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 2 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 columns (2) – (4) indicate whether the attrition hazard in the particular week differs significantly across treatment arm. The first row of the table confirms that ddI subjects

have a significantly lower probability of leaving the trial in week 8. This difference cannot 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 and ddC are known to cause a number of potential side effects, many of which are immediately apparent to the patient, while the side effects associated with ddI are less numerous.¹³ However, the incidence and severity of these side effects may vary substantially across patients.

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 3. 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 a 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

¹³ Side effects associated with AZT include nausea, vomiting, headache, insomnia, muscle pain, muscle wasting, and anemia. Side effects of ddC include mouth ulcers and peripheral neuropathy, while ddI may cause diarrhea, and, in rare cases, pancreatitis.

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 (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.

2.3 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 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 3, 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. MODELLING SUBJECT BEHAVIOR IN THE RANDOMIZED EXPERIMENT

Subjects in the RE 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 two alternative models of attrition that differ in the restrictions imposed on subject behavior as well as the complexity of econometric implementation. Each of the models is structural in the sense that we specify subject utility functions and separately identify direct

effectiveness and side effect distributions. In the myopic learning model, we assume that patients maximize current period utility but are not forward-looking. Subjects also immediately observe side effects S_{it} , although they do not know which arm of the trial they are in. This is consistent with the results shown in Table 3, which indicate 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 the treatment on health status. Subject i learns about the effectiveness of treatment by observing the sequence $\{H_{it}\}$ over the course of the trial. This specification of utility maximization and learning is similar to the pharmaceutical demand model of Currie and Park (2000) in which beliefs concerning drug effectiveness are updated 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.

In the learning model described above, subjects are myopic and only take into account current period expected utility when deciding to remain in the RE. This model therefore seems restrictive, since 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. Consequently, we generalize the learning model to incorporate forward-looking behavior so that subjects 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. The model is similar to that of Crawford and Shum (2003), although those authors did not have an observable measure of H_{it} . In this case, dynamic behavior reflects expectations regarding future utility as well as learning.

3.1 THE SUBJECT'S DECISION PROBLEM

We now describe the forward-looking learning model and note the restrictions that are

imposed by the myopic learning model. To fix ideas, let $t = 0, 1, \dots, \tau$ index the time periods in the trial (corresponding to weeks 0, 8, 20, ..., 104) and $i = 1, \dots, N$ 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 ($k = o$). 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:

$$(1) \quad \max_{\{d_{it}^k\}} E \left[\sum_{s=t}^{\tau} \beta^{s-t} \sum_{k \in \{o, T\}} E[U_{is}^k | I_{is}] d_{it}^k | I_{it} \right] + \beta^{s-\tau} V_i^o(I_{it}, \tau + 1),$$

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

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

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

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

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

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

if he remains 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. Note that the learning model specification assumes that $\beta = 0$ so that subjects maximize current period utility only.

3.2 SPECIFICATION OF PREFERENCES

Subject i 's per-period utility obtained by participating in ACTG175 is a function of his 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} . To simplify the analysis, we assume that utility is additively separable in H_{it} and S_{it} . In addition, the impact of health status on utility is allowed to be non-linear. We also assume that utility is a linear function of side effects, implying that

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

where $\gamma > 0$ is the coefficient of absolute risk aversion.

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 RE at time t defined in equation (2) to be a linear function of observed characteristics x_{oit} as well as a stochastic component:

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

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 in the RE; a subject (and treatment) specific time trend; and a period specific error term:

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

where the subscript $r = \{AZT, AZT+ddI, AZT+ddC, ddI\}$ indicates the treatment arm the subject is randomized into. 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:

$$(9) \quad \theta_{ri} \sim N(z_i\beta_r, \sigma_{\theta_r}^2), \quad \lambda_{ri} \sim N(\lambda_r, \sigma_{\lambda_r}^2).$$

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:

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

Equation (10) incorporates a random coefficient specification to allow the impact of side effects to vary across subjects.

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 v_{it} in equation (8), but they do know its distribution. In addition, subjects are blinded as to which treatment arm they have been assigned. Consequently, from the subject's point of view at time t

$$(11) \quad H_{it} \sim N(H_{i0} + \mu_{it}, \sigma_v^2),$$

where μ_{it} is the unknown individual specific impact on health status of remaining in the trial. We assume that subjects' expectation of their CD4 counts over the course of the trial follows the linear relationship $\mu_{it} = \mu_{0i} + \mu_{1i}t$. 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_{1i} 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_{0i} and μ_{1i} are given by

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

The prior variance $\Sigma_{\mu 0}$ reflects the precision of the prior beliefs 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_{it} in period t is given by

$$(13) \quad E[\mu_{it} | I_{it}] = \mu_{it}^t = (1 - t) \left(\frac{A_{t-1}' A_{t-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \right)^{-1} \left(\frac{A_{t-1}' H_{it-1}^*}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \begin{pmatrix} \mu_{00} \\ \mu_{10} \end{pmatrix} \right),$$

$$\text{Var}(\mu_{it} | I_{it}) = \sigma_{\mu,t}^2 = (1 - t) \left(\frac{A_{t-1}' A_{t-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \right)^{-1} \begin{pmatrix} 1 \\ t \end{pmatrix},$$

where

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

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 the end of week 8 of the RE by the patient. As a result, 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 short-term discomfort associated with consumption of the treatment. Such features of the trial treatment seem likely to be observed within the first eight weeks.

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

$$(14) \quad E[U_{it}^T(H_{it}, S_{it}) | I_{it}] = -\exp(-\gamma^*(H_{i0} + \mu_{it}^t) + \frac{\gamma^2}{2}(\sigma_v^2 + \sigma_{\mu,t}^2)) - x_{iSt} \alpha_i - \varepsilon_{iSt},$$

since H_{it} is a normal random variable. 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 (14) may then be substituted into equations (3) and (4) to obtain the expected return to remaining in the RE for at least one more period.

3.5 RANDOMIZATION, BLINDING, AND THE SPECIFICATION OF EXPECTATIONS

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 as is typically believed, since non-random attrition over the course of the RE introduces the possibility of sample selection bias. However, randomization does aid in the identification of 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 reflect variation in treatment side effects rather than subject heterogeneity.

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 the subjects could observe their treatment assignments at baseline, 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 REs and many social experiments is that blinding is usually not possible 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 issue concerns the specification of expectations. We assume that subjects in the RE have rational expectations, so that the prior intercept and trend parameters μ_{00} and μ_{10} from equation (12) correspond to observed sample averages. These parameters appear to be poorly identified from data consisting of only trial participants, and so are difficult to estimate when rational expectations

are not imposed. Data on individuals choosing not to enter the RE may be necessary for more precise estimates of μ_{00} and μ_{10} .¹⁴ 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

This section presents the econometric method used to estimate the economic model described in Section 3. The estimation procedure is complicated by the fact that dropout decisions are correlated over time, as are the observed health status measures $\{H_{it}\}$. Therefore, our econometric approach relies on simulation-based methods and proceeds in two steps. To begin, consider a subject who drops out in period t . The likelihood contribution of this subject is then the joint probability of observing the sequence of health outcomes $H_{i1}, H_{i2}, \dots, H_{it-1}$ and dropout in period t and is given by

$$(15) \quad L_i = \Pr(H_{i1}, \dots, H_{it-1}, d_{i1}^T = 1, \dots, d_{it-1}^T = 1, d_{it}^T = 0 \mid \Theta_1, \Theta_2, x_i, z_i),$$

where the parameters from the dropout decision and health status equations are denoted by $\Theta_1 = \{\beta, \gamma, \delta, \alpha, \sigma_\alpha, \Sigma_{\mu 0}\}$ and $\Theta_2 = \{\beta_r, \sigma_{\beta_r}, \lambda_r, \sigma_v\}$, respectively, and $x_i = \{x_{Si}, x_{oit}\}$. Equation (15) can be rewritten as

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

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

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

The form of V_{it}^T will depend on the model that is being estimated, and increases in

¹⁴We estimated the model presented in Section 5 including μ_{00} and μ_{10} 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_{00} and μ_{10} (though still positive) with virtually no change in the likelihood value.

complexity as we incorporate forward-looking behavior. 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 (14). The full derivation of the likelihood in the forward-looking case is described in Appendix 1.

Both x_{Si} and x_{oit} and ε_{Sit} and ε_{oit} enter linearly into $V_{it}^T - V_{it}^o$ in equation (17), so that only contrasts between the two are identified. We assume that $\text{Var}(\varepsilon_{Sit} - \varepsilon_{oit}) = 1$ 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 maximum likelihood to obtain the estimates.

5. RESULTS

We fit both the myopic and forward-looking learning models to the ACTG 175 data using the econometric methods described above. The literature on ACTG 175, such as Hammer et al (1996), suggests that the following variables might be expected to affect the subject-level treatment impact on CD4 counts, and hence are included in z_i : (1) demographic variables such as age, gender, and race; (2) variables measuring the extent of disease at baseline, such as the screening CD4 count measured prior to baseline that determined entry into the trial; whether the subject has a symptomatic HIV infection (which suggests a greater spread of the disease); and whether the subject had been received to prior antiretroviral therapy (which again may suggest a greater spread of 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 in order 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 over a three year period in the early 1990s, during which time new AIDS treatments were becoming available. The vector $x_{i,t}$ 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 likely to be less stable and follow up in the trial (the other categories are haemophiliac and homosexual). It may be the case that subjects with more financial resources are 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 so we include these variables as proxies.

Side effects are specified to be a function of treatment assignment, and, as shown in equation (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. We now turn to the discussion of the parameter estimates from the myopic and forward-looking learning models.

5.1 RESULTS FOR THE STRUCTURAL ATTRITION MODELS

The parameter estimates from the myopic and forward-looking models of the impact of the factors that affect the decision to remain in ACTG 175 or seek alternative treatment are presented in columns (1) and (2) of Table 4. Positive coefficients indicate a higher probability of remaining in the trial. The forward-looking model appears to provide a better fit to the data, although the two models yield similar parameter estimates in many cases, so in the subsequent discussion we focus on the results from Column (2).

The positive estimate of γ shown in column (2) of Table 4 implies that subjects with higher

expected CD4 counts are more likely to remain in the trial. In Figure 6 we plot 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 1000. The figure implies that for individuals with CD4 counts less than 200, indicating serious immune system damage and high risk of opportunistic infection, slight declines in CD4 counts substantially decrease the utility associated with remaining in the trial, leading to dropout. Conversely, for individuals with CD4 counts greater than 500, which are generally in 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 the RE than those receiving AZT alone.

With regard to the coefficient estimates for the other important covariates, subjects who had prior antiretroviral therapy are more likely to remain in the trial, perhaps because they have more experience with the side effects induced by these types of drugs. Symptomatic individuals (symptoms include diarrhoea, minor skin or oral conditions, lack of energy, swollen glands, etc.) are experiencing health problems similar to those induced by treatment side effects, and so these latter factors may 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.

Before proceeding, we consider the overall fit of the forward-looking model by comparing the predicted probabilities of dropout generated by the coefficient estimates with the observed frequencies in ACTG 175. For each subject, we use the estimates in column (2) to construct the probability of dropout in weeks 8, 20, 32, ..., 104 and then construct sample averages. Figure 7 shows that the model appears to fit the data quite well. For the most part, the predicted probabilities of dropout are quite close to the actual survival probabilities. The main difference occurs in week 8, in which the model slightly overpredicts the survival probability. This likely reflects the common priors concerning treatment effectiveness as well as our assumption of time-invariant side effects. However, the discrepancy between the observed and predicted week 8 dropout is not more than six percentage points, suggesting that the additional complication of a richer side-effect specification may not yield substantial extra benefits.

5.2 HOW FAST DO SUBJECTS LEARN?

The estimates of the standard deviations of the prior intercept and time slope parameters in the bottom rows of Table 4 indicate substantial variation in subjects' prior beliefs concerning the effectiveness of trial treatment. In order to assess implications of these estimates for the speed with which expected beliefs concerning the impact of the trial drugs on CD4 counts converge to the actual value for a particular subject, we constructed the expected CD4 count in each period, equal to $H_{i0} + \mu_{it}^t$, using equation (13) and the rational expectations assumption for representative subjects, assuming that they have a baseline CD4 count of 350.

Figure 8 presents the expected ("Belief") and actual CD4 counts for subjects receiving AZT and AZT+ddC, and exhibits a number of notable features.¹⁵ First, because subjects have common priors concerning treatment effectiveness and do not know the treatment they are receiving, at baseline both subjects described in Figure 8 expect their CD4 count to increase from 350 to 383 by

week 8. For the AZT patient, this belief overstates the actual change in his CD4 counts, while the expected change in CD4 counts slightly understates the actual change for the AZT+ddC subject. Moreover, the common prior beliefs concerning week 8 CD4 counts resulting from randomization and subject blinding implies that the observed difference in actual CD4 counts in week 8 will have no impact on the difference in dropout behavior across treatment arms. Initial differences in attrition across trial arm are therefore likely to result from the side effects of the different treatments.

As the subjects accumulate information on their CD4 counts over the course of the trial, Figure 8 shows that their beliefs generally converge toward the actual values, implying that the greater effectiveness of AZT+ddC (and the AZT+ddI and ddI treatments not shown in the figure) eventually leads to reduced dropout compared to those receiving AZT alone. Learning behavior may therefore explain the fanning out of the survival curves observed in Figure 4. The graph for the AZT+ddC recipient in the bottom half of the figure suggests that beliefs understate actual CD4 counts when improvement in CD4 counts is observed between weeks 44 and 56. This fairly sharp change is difficult to immediately capture with prior beliefs described by an intercept and time trend. Moreover, the subject is unlikely to initially believe that the sharp increase at week 56 is permanent. However, as the improvement continues in the subsequent weeks, the posterior beliefs converge on the actual CD4 count by week 104.

The posterior variance described in equation (13), $\sigma_{\mu,t}^2$, describes the degree of uncertainty regarding treatment effectiveness perceived by trial subjects. Figure 9 describes the evolution of the posterior standard deviation, $\sigma_{\mu,t}$, over the course of the trial. The figure suggests that while subjects initially experience substantial uncertainty regarding effectiveness, by the end of the trial the posterior standard deviation is approximately one-half the prior standard deviation. From equation (14) and the positive estimate of γ , reduction in uncertainty regarding treatment effectiveness

¹⁵ Similar features are observed for subjects receiving AZT+ddI and ddI.

increases expected utility. Forward-looking subjects incorporate the reduction of this uncertainty in the future when making their decision to remain in the trial in the current period. Overall, Figures 8 and 9 suggest that while subjects are initially uncertain regarding the impact of the trial on health status, expected effectiveness generally converges to its observed value over the course of the trial, and patient uncertainty declines substantially.

5.3 SIDE EFFECT DISTRIBUTIONS

In this section we examine the side effects of the alternative treatments and their impact on dropout. To provide an indication of how patients value side effects, the results from column (2) of Table 4 suggest that the average AZT patient with a CD4 count of 350 would be willing to trade a 3.7% reduction in his CD4 count for a 1% improvement in the side effects associated with AZT.¹⁶ Comparing the treatments, the coefficient estimates indicate that patients receiving AZT+ddC are more likely to drop out of the trial, after accounting for expected impact on CD4 counts, than are other subjects, suggesting that this therapy has negative side effects relative to other treatments. On the other hand, the positive and significant estimate for ddI suggests that subjects receiving this therapy appear to experience the fewest side effects. However, the subject level variation in side effects suggests that many individuals in this group may still experience greater side effects than those receiving AZT, AZT+ddC, or AZT+ddI. Plots of the subject-level side effect distributions for the four treatment groups shown in Figure 10 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, but there is still substantial overlap of the random effect distributions. One interpretation of our 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 of ACTG 175 due to the immediately

¹⁶ Because of the declining marginal utility of H_{it} shown in Figure 6, an AZT patient with a CD4 count of 200 (500) would accept a 2.4% (6.7%) reduction in his CD4 count in exchange for a 1% improvement in side effects.

perceived higher side effects associated with combination therapy. Over time, the fact that AZT subjects experience greater declines their CD4 counts offsets the lower side effects associated with the treatment relative to AZT+ddC, leading to increased attrition among patients receiving this therapy.

We have interpreted the difference in dropout probabilities across the treatment arms of ACTG 175 as reflecting the difference in side effects, conditional on expected CD4 count. Because side effects are not observed directly, one might argue that the results reflect other factors, such as differences in the outside treatment options available to subjects in the two groups. However, randomization implies that the average outside option will not differ across the two groups at baseline due to unobserved characteristics.

5.4 THE IMPACT OF ALTERNATIVE TREATMENTS ON CD4 COUNTS

Table 5 presents the estimates for the change in CD4 outcome equation (8) associated with the forward-looking learning model 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 5 suggest that AZT+ddC subjects have the highest CD4 count, on average, and AZT patients the lowest. Note, however, that there is substantial subject-level heterogeneity in CD4 profiles, particularly for patients receiving AZT+ddI. In contrast, AZT appears to generate 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 with CD4 counts indicates that the ranking of the ACTG 175 therapies on each of these measures did not necessarily coincide. For example, AZT+ddC was found to have the most negative side effects, while it is clearly superior to AZT when examining CD4 counts. We now turn to an assessment of 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 been free to do so.

Using the structural estimates of the subject utility function shown in Table 4 and the CD4 profile estimates in Table 5, we construct the expected discounted values at baseline of the stream of utilities associated with each of the four ACTG 175 treatments for each of the 2200 subjects in the trial. The simulation is conducted under the assumption 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 “type” (i.e., his 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 four therapeutic alternatives is recorded.

One difficulty in conducting a simulation 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, since 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 4 and 5. Table 6 presents the “market shares” for each of the treatments 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)).

The preference data presented in Table 6 indicates three notable findings. First, either 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 treatments may be preferred by many patients. This is borne out by the finding from the table that no treatment has a market share of more than 35% in any of the scenarios, with the exception of the simulation assuming perfect positive correlations in column (6), where ddI has 53% market share. The generally 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 6 is that a non-negligible fraction of subjects prefer AZT alone to the other treatments, despite the fact that AZT had the least average 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 any positive impact of the therapy on

CD4 counts. 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, some 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 6 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 of trial subjects. In particular, we ask the following set of questions: How would the choices of subjects differ from those shown in Table 6 if subjects have only one period (i.e., week 8 of the trial) of information on the side effects and effectiveness of each of the four drugs? How would the choices differ if the subjects have two periods (i.e., week 20) of information? Nine periods (i.e., through week 104) of information? To answer these questions, we simulate subject choices under each of these information scenarios using the parameter estimates from Tables 4 and 5 and the Bayesian updating formulas in Equation (13), and compare those choices to those made under perfect certainty as shown in Table 6. We assume that the correlation of side effects and direct effects across treatments is 0.5 (corresponding to column (2) of Table 6) and again assume that subjects cannot choose the outside option.

Figure 11 presents the results of these simulations for the four treatments under different levels of information. On the vertical axis we plot the fraction of subjects choosing the drug in question who would have chosen a different drug if they had full information. We denote these

subjects as “mismatched.” For example, at week 8, approximately 35% of the subjects who chose AZT based on one period of information would have chosen a different treatment if they had full information regarding their level of side effects and the direct effect on their CD4 count as in Table 6. The figure indicates that learning appears to occur fairly quickly. For AZT subjects, the second period of information sharply reduces mismatch because at week 8, subjects only knew side effects, but had common prior beliefs concerning drug effectiveness. At week 20, subjects gain their first piece of information concerning the impact of the treatment on CD4 counts and infer that AZT is less effective. For AZT+ddI and ddI patients, three periods of information (week 32) is enough to reduce mismatch to almost zero. Learning is somewhat more gradual for AZT+ddC patients. This reflects the fact that initially, the negative side effects of this treatment dominate the longer term effect on CD4 counts in subject choices, leading subject to choose alternative treatments. As subjects gain additional information on the longer term effects of AZT+ddC, particularly regarding the time trend shown in Table 5, treatment choice approximates that under perfect certainty. The generally rapid learning shown in Figure 11 indicates that market shares for the four treatments by week 56 are virtually identical to those simulated under full information in column (2) of Table 6. 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 so significantly enhances patient welfare.

7. CONCLUSION

This paper presents an alternative 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 of the RE that have typically been the focus of the literature. The standard approach does not capture a variety of features often observed in REs, 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 RE. Our evaluation framework incorporates these factors by viewing subjects as utility maximizing agents who decide each period whether to drop out or remain in the RE by comparing 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 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 RE. We are then able to evaluate treatment effectiveness in the RE 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 the AIDS randomized clinical trial ACTG 175, a well known longitudinal randomized clinical trial. 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 20%), AZT alone yields higher utility than the other treatments. Moreover, AZT+ddC appears to have 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 REs, such as partial or non-compliance with experimental treatments or crossover behavior if subjects are allowed to choose an alternative experimental treatment at some point during the RE. Such features may provide greater information concerning correlations in side effects or direct effects across treatments. We have also assumed that the role of the experiment's investigators is neutral, so that dropout decisions are made by the subject. However, in some REs, the investigators

may drop subjects from the experiment.¹⁷ In this case, one could incorporate a utility maximizing model of investigator behavior into the evaluation framework and view exit from the RE as a competing risk. 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 the alternative 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. Due to the lack of subject blinding in that case, separate prior beliefs would likely have to be constructed for each treatment.

¹⁷ See Philipson and Hedges (1998) for a discussion of investigator behavior. Recall that in ACTG 175, we found no evidence that the trial investigators dropped subjects from the experiment.

TABLE 1
OLS ESTIMATES OF THE IMPACT OF ALTERNATIVE TREATMENTS
ON CD4 COUNTS

Dependent Variable is Change in CD4 Count Relative to Baseline

Variables	Specification	
	(1)	(2)
Constant	-9.798 (4.223)	-13.481 (7.890)
AZT+ddI	67.799 (5.924)	67.354 (5.853)
AZT+ddC	24.251 (5.982)	24.924 (5.910)
ddI	50.140 (5.829)	49.965 (5.759)
<i>t</i>	-4.479 (0.850)	-4.797 (0.841)
AZT+ddI* <i>t</i>	1.816 (1.169)	2.296 (1.155)
AZT+ddC* <i>t</i>	5.329 (1.186)	5.386 (1.172)
ddI* <i>t</i>	1.278 (1.154)	1.504 (1.140)
Includes Covariate Controls?	No	Yes
Includes Indicators for Week of Dropout?	No	No

Notes: Standard errors in parentheses. Based on 12520 subject-week observations. Excluded category is AZT. Covariate controls include age, gender, race, symptomatic HIV infection, screening CD4 count, and prior antiretroviral therapy.

TABLE 2
DISCRETE HAZARD ESTIMATES OF THE PROBABILITY OF ATTRITION,
BY WEEK AND TRIAL ARM

Variable	Trial Arm			
	AZT (1)	AZT+ddI (2)	AZT+ddC (3)	ddI (4)
Week 8	-1.148 (0.068)	0.017 (0.096)	0.096 (0.095)	-0.247 (0.103)
Week 20	-1.170 (0.073)	-0.232 (0.111)	-0.060 (0.106)	-0.203 (0.108)
Week 32	-1.307 (0.084)	-0.126 (0.122)	-0.119 (0.123)	-0.227 (0.124)
Week 44	-1.264 (0.086)	-0.178 (0.127)	-0.102 (0.125)	-0.255 (0.127)
Week 56	-1.246 (0.090)	-0.154 (0.130)	-0.087 (0.130)	-0.168 (0.128)
Week 68	-1.321 (0.099)	-0.153 (0.143)	-0.187 (0.147)	-0.084 (0.137)
Week 80	-1.246 (0.100)	-0.338 (0.152)	-0.178 (0.146)	-0.096 (0.138)
Week 92	-1.286 (0.108)	-0.269 (0.158)	-0.316 (0.164)	-0.189 (0.152)
Week 104	-1.137 (0.106)	-0.234 (0.150)	-0.123 (0.149)	-0.300 (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 3
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 4
PARAMETER ESTIMATES FOR PROBABILITY OF REMAINING IN TRIAL

Variable		Model	
		Myopic Learning (1)	Forward-Looking Learning (2)
	γ	0.0035 (0.0014)	0.0068 (0.0016)
AZT	Mean	1.074 (0.252)	0.849 (0.257)
	S.D.	1.069 (0.151)	1.360 (0.152)
AZT+ddI	Mean*	0.120 (0.101)	0.153 (0.112)
	S.D.	1.136 (0.159)	1.470 (0.169)
AZT+ddC	Mean*	-0.074 (0.096)	-0.078 (0.102)
	S.D.	1.166 (0.155)	1.394 (0.159)
ddI	Mean*	0.229 (0.099)	0.270 (0.112)
	S.D.	0.973 (0.158)	1.336 (0.155)
<hr/>			
	Age	0.026 (0.004)	0.027 (0.005)
	Male	0.145 (0.104)	0.114 (0.112)
	White	0.013 (0.073)	0.069 (0.079)
	Screening CD4 Count	0.00002 (0.0004)	-0.00005 (0.0004)
	Prior Antiretroviral Treatment	0.276 (0.066)	0.262 (0.072)
	Symptomatic	-0.096 (0.083)	-0.079 (0.089)
	IV Drug User	-0.515 (0.100)	-0.588 (0.111)
	Homosexual	0.079 (0.086)	0.085 (0.091)
	Calendar Quarter Since 4/91	-0.075 (0.022)	-0.153 (0.023)
<hr/>			
Prior Intercept	S.D.	57.661 (54.274)	66.330 (25.409)
Prior Time Slope	S.D.	109.156 (42.458)	58.663 (10.824)
<hr/>			
	Log-Likelihood	-3528.85	-3567.56

Note: Standard errors in parentheses. Estimates based on the 2200 subjects in ACTG 175. The forward-looking learning model in column (2) assumes a discount factor of 0.95.

*Parameter estimate of mean drug effect relative to AZT.

TABLE 5
PARAMETER ESTIMATES FOR CD4 EQUATION
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* <i>t</i>	Mean	-8.810 (0.848)
	S.D.	9.699 (0.735)
AZT+ddI* <i>t</i>	Mean	-8.158 (0.972)
	S.D.	14.397 (0.733)
AZT+ddC* <i>t</i>	Mean	-1.875 (0.898)
	S.D.	11.871 (0.673)
ddI* <i>t</i>	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 6
SIMULATIONS OF PATIENT PREFERENCES FOR ACTG 175 THERAPIES

Correlation between Side Effects	0	0.5	-0.5	0.5	-0.5	1	-1
Correlation between Drug Effectiveness	0	0.5	-0.5	-0.5	0.5	1	-1
ACTG 175 Therapy	(1)	(2)	(3)	(4)	(5)	(6)	(7)
AZT	0.21 [0.15, 0.28]	0.16 [0.06, 0.32]	0.21 [0.18, 0.27]	0.20 [0.18, 0.27]	0.21 [0.19, 0.26]	0.06 [0, 0.42]	0.21 [0.18, 0.26]
AZT+ddI	0.28 [0.24, 0.32]	0.31 [0.21, 0.40]	0.30 [0.26, 0.33]	0.33 [0.28, 0.36]	0.30 [0.27, 0.33]	0.34 [0.13, 0.54]	0.31 [0.27, 0.34]
AZT+ddC	0.22 [0.19, 0.25]	0.18 [0.12, 0.24]	0.22 [0.20, 0.25]	0.22 [0.19, 0.24]	0.23 [0.20, 0.25]	0.08 [0.02, 0.20]	0.22 [0.20, 0.24]
ddI	0.28 [0.25, 0.33]	0.35 [0.24, 0.46]	0.26 [0.23, 0.30]	0.25 [0.20, 0.31]	0.26 [0.23, 0.29]	0.53 [0.23, 0.74]	0.26 [0.23, 0.29]

Note: Simulations based on 2200 subjects from ACTG 175. Table entries are the fraction of patients preferring the indicated treatment. Entries in brackets are 5th and 95th percentiles of patients preferring the treatment, respectively. Table entries are conditional on not choosing the outside option. 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," *Econometrica*, 73, mimeo, Department of Economics, University of Minnesota.
- Crawford, G. and Shum, M. (2005). "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 73, 1137-1174.
- Currie, G. and Park, S. (2000). "The Effects of Advertising and Consumption Experience on the Demand for Antidepressant Drugs," mimeo, Dept. of Economics, University of Calgary.
- 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.
- Ferrall, C. (2002). "Estimation and Inference in Social Experiments," mimeo, Department of Economics, Queen's University.
- 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.
- Greenberg, D. and Shroder, M. (1997). *The Digest of Social Experiments*, 2nd ed. The Urban Institute Press: Washington, D.C.
- 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.

Malani, A. (2003). "Estimating Treatment Response in Clinical Trials when Placebo Effects are Present," mimeo, University of Virginia Law School.

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 1

A. Bayesian Updating of Beliefs

We assume

$$(11) \quad H_{it} \sim N(H_{i0} + \mu_{it}, \sigma_v^2)$$

and we assume $\mu_{it} = \mu_{0i} + \mu_{1i}t$. We also assume prior beliefs as

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

Let

$$A_{t-1} = \begin{pmatrix} 1, 1, \dots, 1 \\ 1, 2, \dots, t-1 \end{pmatrix},$$

and

$$H_{i,t-1} = (H_{i,1} - H_{i,0}, \dots, H_{i,t-1} - H_{i,0})'.$$

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

$$(A.1) \quad E \left[\begin{pmatrix} \mu_{1i} \\ \mu_{2i} \end{pmatrix} \middle| I_{it} \right] = Dd,$$

$$(A.2) \quad \text{var} \left[\begin{pmatrix} \mu_{1i} \\ \mu_{2i} \end{pmatrix} \middle| I_{it} \right] = D,$$

where

$$D = \left(\frac{A_{t-1}' A_{t-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \right)^{-1},$$

and

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

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

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

To derive L_i in (16) and (17) 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 let's define a discrete decision process $D_{i,t,u} = \{d_{i,t} = 1, \dots, d_{i,t+u} = 1, d_{i,t+u+1} = 0, \dots, d_{i,T} = 0\}$, where $d_{i,s}$ is a discrete decision at period s , with $d_{i,s} = 1$ represents subject i stays in the experiment and $d_{i,s} = 0$ represents drops-out. $D_{i,t,u}$ represents the decision of subject i stays in the experiment until $t+u \leq T$, and drops out in period $t+u+1$.

Define

$$(A.3) \quad \hat{V}_i^T(D_{i,t,u}; I_{i,t}, t) = E[U_{i,t}^T - U_{i,t}^0 | I_{i,t}] + \beta E[U_{i,t+1}^T - U_{i,t+1}^0 | I_{i,t}] \\ + \dots + \beta^u E[U_{i,t+u}^T - U_{i,t+u}^0 | I_{i,t}]$$

Then we have

$$(A.4) \quad V_i(I_{it}, t) = \max\{V_i^T(D_{i,t,0}; I_{it}, t), V_i^T(D_{i,t,1}; I_{it}, t), \dots, V_i^T(D_{i,t,T-t}; I_{it}, t), 0\}$$

Analogous to equation (14), we have

$$(A.5) \quad E[U_{it}^T - U_{it}^0 | I_{it}] \equiv \hat{\mu}_{it} \\ = -\exp(-\gamma * (H_{i0} + \mu_{it}^t)) + \frac{\gamma^2}{2}(\sigma_v^2 + \sigma_{\mu,t}^2) - x_{Si}\alpha_i - \varepsilon_{iSt},$$

where μ_{it}^t and $\sigma_{\mu,t}^2$ are defined as in equation (13).

Based on (A.5) and using iterative expectation we will obtain

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

since $E[\varepsilon_{iS,t+1} | I_{i,t}] = 0$ by assumption. After some manipulation we obtain

$$\begin{aligned}
(A.6) \quad & E[U_{i,t+1}^T - U_{i,t+1}^0 | I_{it}] \equiv \hat{u}_{i,t+1} \\
& = -\exp\left\{ -\gamma^* [H_{i0} + \begin{pmatrix} 1 \\ t+1 \end{pmatrix}' D_{t+1} \cdot (\frac{A_{t-1}' H_{i,t-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \cdot \mu_0)] \right. \\
& \quad + \frac{\begin{pmatrix} 1 \\ t+1 \end{pmatrix}' D_{t+1} \cdot \begin{pmatrix} 1 \\ t+1 \end{pmatrix}}{\sigma_v^2} \cdot \begin{pmatrix} 1 \\ t \end{pmatrix}' D_t \cdot (\frac{A_{t-1}' H_{i,t-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \cdot \mu_0)] \\
& \quad \left. + \frac{\gamma^2}{2} \left[\left(\sigma_v^2 + \begin{pmatrix} 1 \\ t+1 \end{pmatrix}' D_{t+1} \cdot \begin{pmatrix} 1 \\ t+1 \end{pmatrix} \right) + \left(\frac{\begin{pmatrix} 1 \\ t+1 \end{pmatrix}' D_{t+1} \cdot \begin{pmatrix} 1 \\ t+1 \end{pmatrix}}{\sigma_v^2} \right)^2 \cdot \left(\begin{pmatrix} 1 \\ t \end{pmatrix}' D_t \cdot \begin{pmatrix} 1 \\ t \end{pmatrix} + \sigma_v^2 \right) \right] \right\} \\
& \quad - x_{Si} \alpha_i,
\end{aligned}$$

where $\Sigma_{\mu 0}^{-1}$, A_{t-1} , and $H_{i,t-1}$ are defined as above, $\mu_0 = \begin{pmatrix} \mu_{10} \\ \mu_{20} \end{pmatrix}$, and

$$D_t = \left(\frac{A_{t-1}' A_{t-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \right)^{-1}, \quad D_{t+1} = \left(\frac{A_t' A_t}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \right)^{-1}.$$

We can use the same method and iterate for k periods. By method of induction we will get

$$\begin{aligned}
(A.7) \quad & E[U_{i,t+k}^T - U_{i,t+k}^0 | I_{it}] \equiv \hat{u}_{i,t+k} \\
& = -\exp\left\{ -\gamma \cdot \left[H_{i0} + \psi_{t,k}' \tilde{D}_{t,k} \cdot \left(\frac{A_{t-1}' H_{i,t-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \cdot \mu_0 \right) \right] + \frac{\gamma^2}{2} \left[(\psi_{t,k})^2 (\phi_{t,k} + \sigma_v^2) \right] \right\} - x_{Si} \alpha_i,
\end{aligned}$$

where

$$\begin{aligned}
\tilde{D}_{t,k} &= \begin{pmatrix} \tilde{D}_{i,t+k} \\ \tilde{D}_{i,t+k-1} \\ \dots \\ \tilde{D}_{i,t} \end{pmatrix}; \quad \tilde{D}_{i,t+l} = \begin{pmatrix} 1 \\ t+l \end{pmatrix}' D_{t+l} = \begin{pmatrix} 1 \\ t+l \end{pmatrix}' \left(\frac{A_{t+l-1}' A_{t+l-1}}{\sigma_v^2} + \Sigma_{\mu 0}^{-1} \right)^{-1}; \\
\phi_{t,k} &= \begin{pmatrix} \sigma_{\mu,t+k}^2 \\ \sigma_{\mu,t+k-1}^2 \\ \dots \\ \sigma_{\mu,t}^2 \end{pmatrix}; \quad \sigma_{\mu,t+l}^2 = \tilde{D}_{i,t+l}^2 \cdot \begin{pmatrix} 1 \\ t+l \end{pmatrix}; \\
\psi_{t,k} &= \begin{pmatrix} 1 \\ \hat{\psi}_{t,k}[1] \\ \hat{\psi}_{t,k}[1] + \hat{\psi}_{t,k}[2] \\ \dots \\ \sum_{l=1}^k \hat{\psi}_{t,k}[l] \end{pmatrix}; \quad \psi_{t,k} = \begin{pmatrix} \tilde{\psi}_{t,k}[1] \\ \tilde{\psi}_{t,k}[1] \times \tilde{\psi}_{t,k}[2] \\ \dots \\ \tilde{\psi}_{t,k}[1] \times \tilde{\psi}_{t,k}[2] \times \dots \times \tilde{\psi}_{t,k}[k] \end{pmatrix}; \quad \tilde{\psi}_{t,k} = \begin{pmatrix} \sigma_{\mu,t+k}^2 / \sigma_v^2 \\ \sigma_{\mu,t+k-1}^2 / \sigma_v^2 \\ \dots \\ \sigma_{\mu,t+1}^2 / \sigma_v^2 \end{pmatrix},
\end{aligned}$$

and $[l]$ denotes the l -th element in a $k \times 1$ vector.

From above we will get

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

and therefore

$$(A.8) \quad V_i(I_{i,t}, t) = \max \left\{ \max \left\{ \begin{aligned} & W_i^T(D_{i,t,0}; I_{i,t}, t) - x_{S,i} \alpha_i, W_i^T(D_{i,t,1}; I_{i,t}, t) - (1 + \beta) x_{S,i} \alpha_i, \\ & \dots, W_i^T(D_{i,t,T-t}; I_{i,t}, t) - \frac{1 - \beta^{T-t+1}}{1 - \beta} x_{S,i} \alpha_i \end{aligned} \right\} + \varepsilon_{iSt}, 0 \right\}$$

We can evaluate the probability of drop-out or staying in the experiment of subject i at each period t .

FIGURE 1
AVERAGE CD4 COUNT PROFILES FOR HYPOTHETICAL RCT

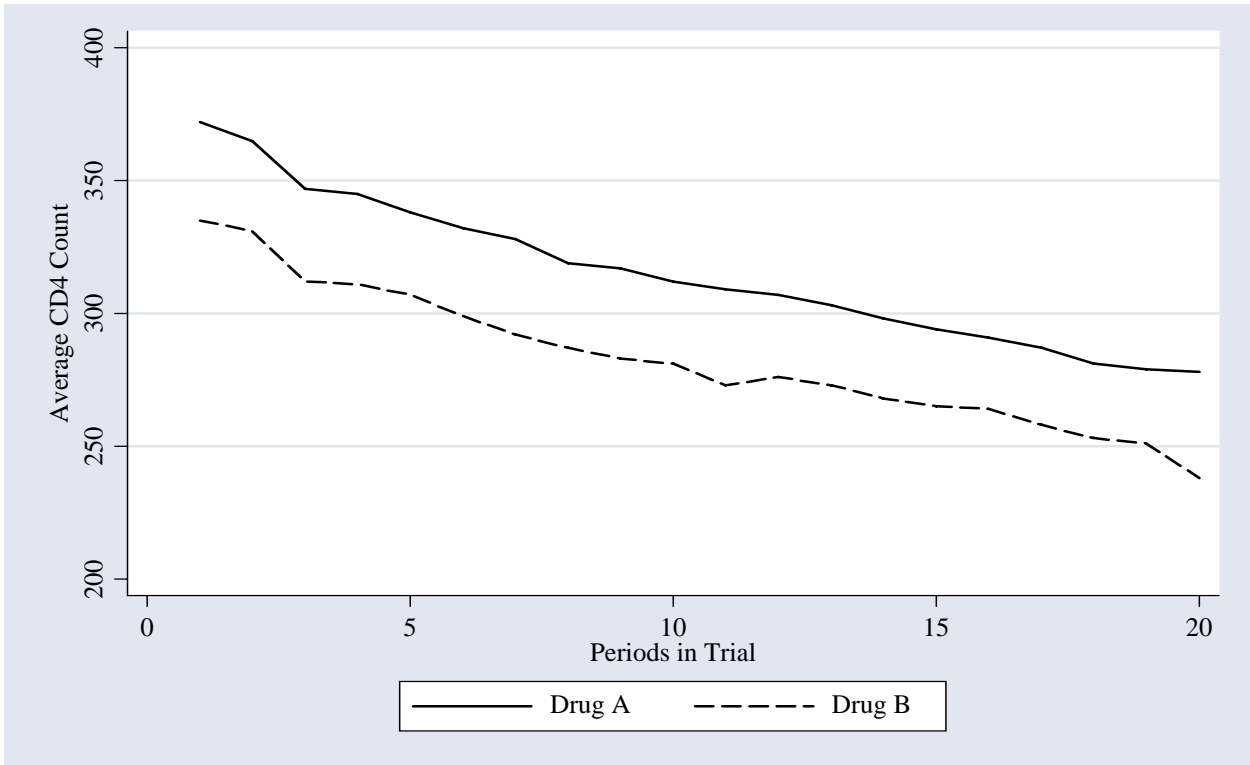


FIGURE 2
SURVIVOR FUNCTIONS FOR HYPOTHETICAL RCT

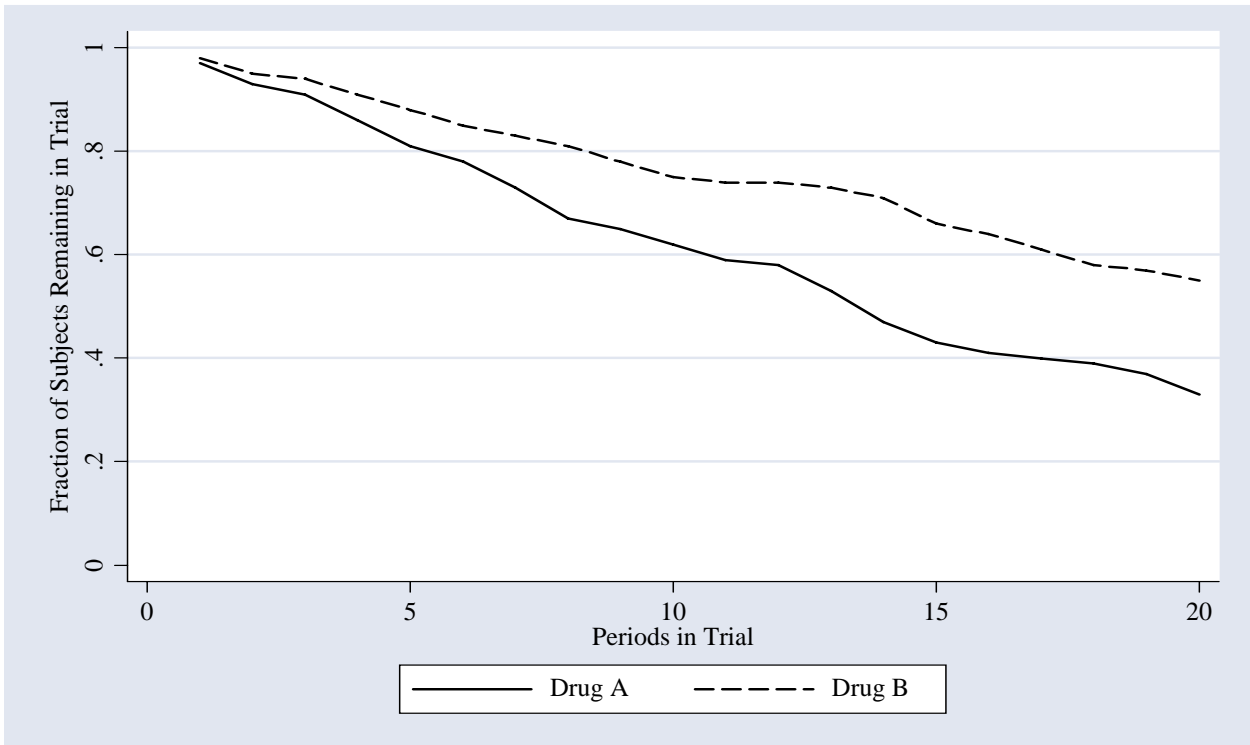


FIGURE 3
CD4 COUNT PROFILES, BY TREATMENT ARM

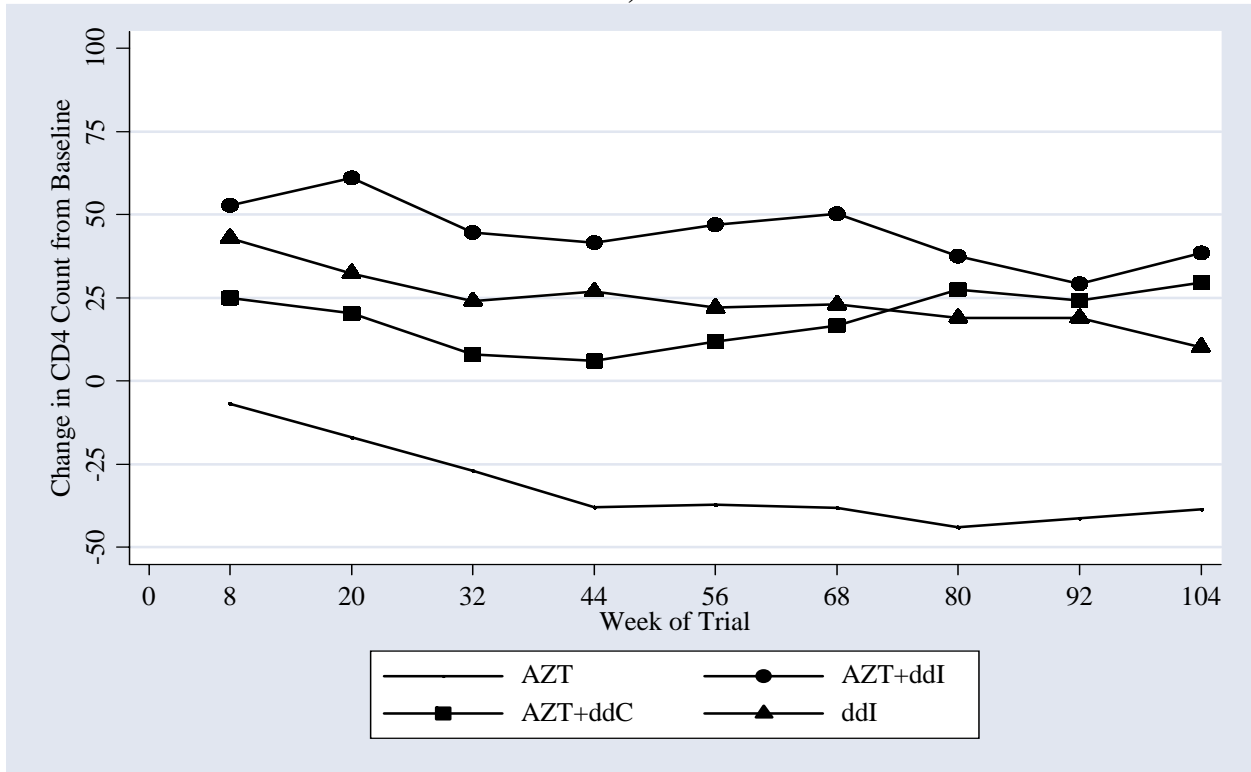


FIGURE 4
SURVIVOR FUNCTIONS, BY TREATMENT ARM




FIGURE 5
CD4 COUNT PROFILES, BY TREATMENT AND ATTRITION GROUP

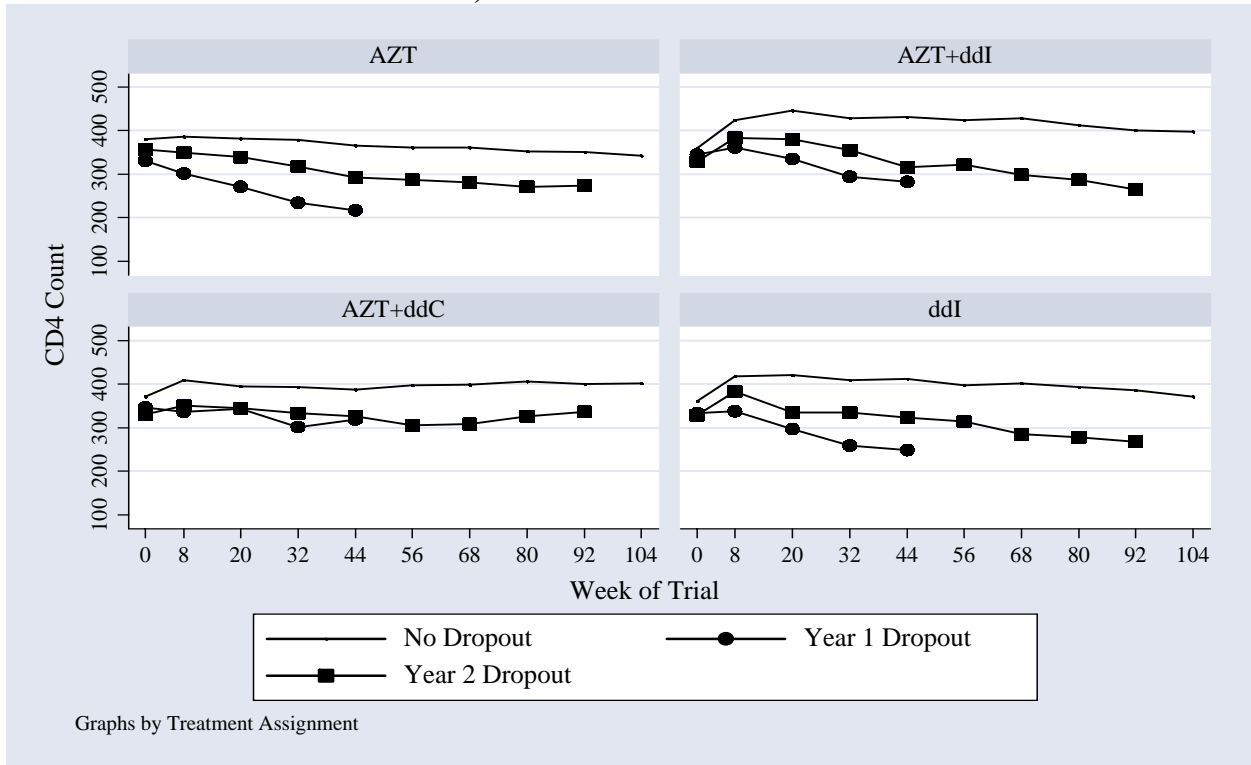
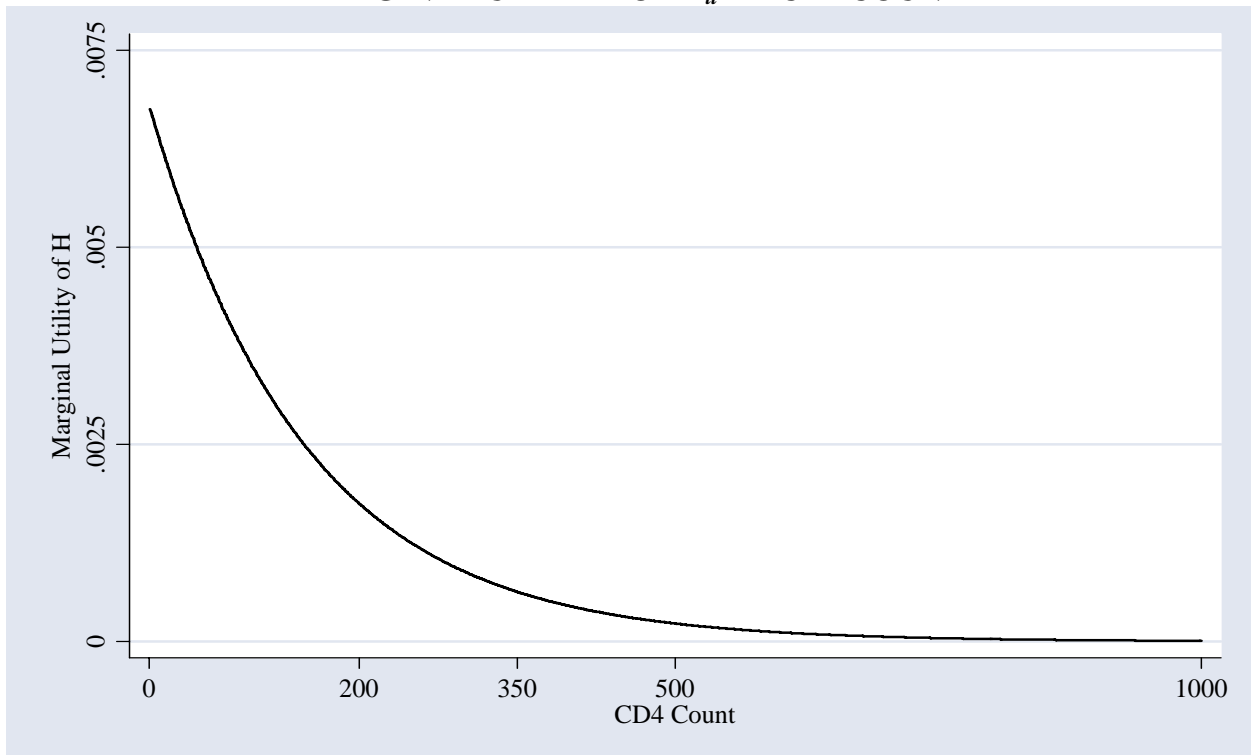
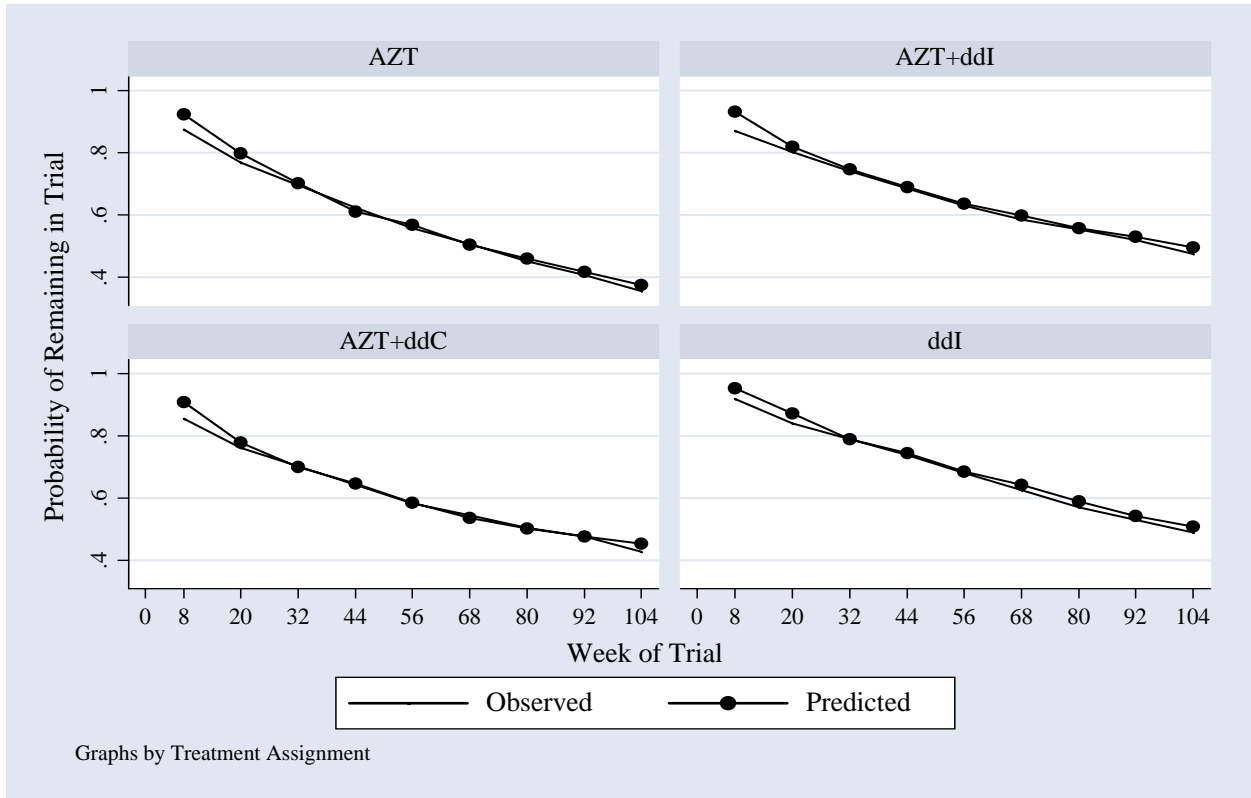


FIGURE 6
MARGINAL UTILITY OF H_{it} BY CD4 COUNT



**FIGURE 7
OBSERVED AND PREDICTED DROPOUT PROBABILITIES**



**FIGURE 8
CONVERGENCE BETWEEN ACTUAL AND EXPECTED CD4 COUNTS**

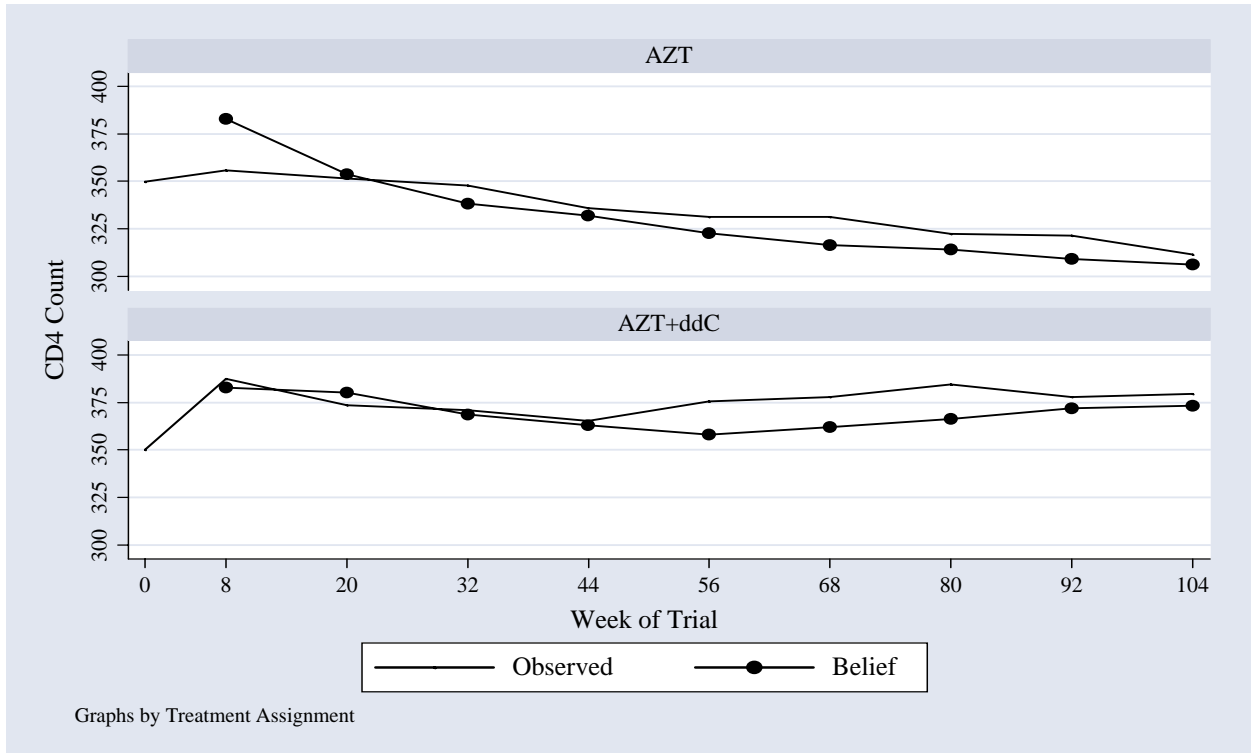


FIGURE 9
EVOLUTION OF UNCERTAINTY

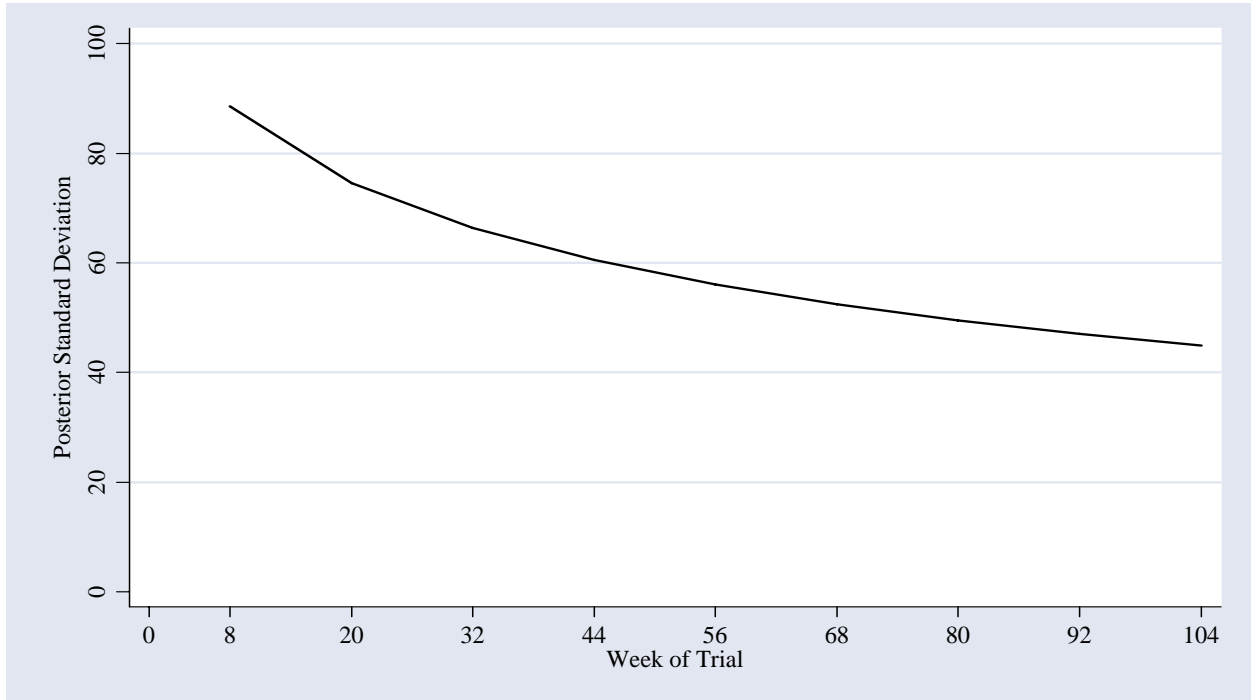


FIGURE 10
DISTRIBUTION OF SUBJECT SIDE EFFECTS, BY TREATMENT GROUP

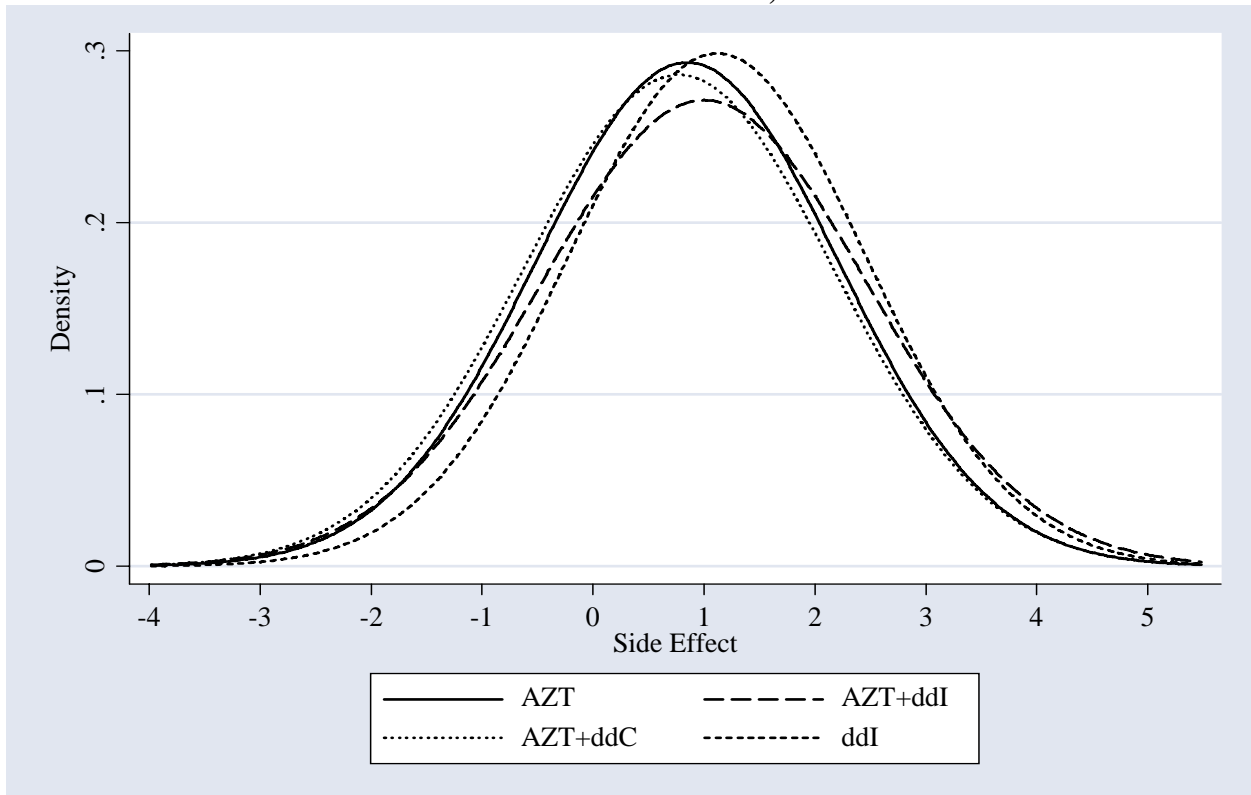


FIGURE 11
SIMULATED IMPACT OF INFORMATION ON OPTIMAL TREATMENT CHOICE

