



Munich Personal RePEc Archive

An Empirical Model of Learning under Ambiguity: The Case of Clinical Trials

Fernandez, Jose

University of Louisville

3 April 2008

Online at <https://mpra.ub.uni-muenchen.de/8621/>

MPRA Paper No. 8621, posted 07 May 2008 04:52 UTC

An Empirical Model of Learning under Ambiguity: The Case of Clinical Trials

Jose M. Fernandez*

University of Louisville

April 18, 2008

Abstract

In this paper, I present an empirical model of learning under ambiguity in the context of clinical trials. Patients are concerned with learning the treatment effect of the experimental drug, but face the ambiguity of random group assignment. A two-dimensional Bayesian model of learning is proposed to capture patients' beliefs on the treatment effect and group assignment. These beliefs are then used to predict patient attrition in clinical trials. Patient learning is demonstrated to be slower when taking into account group ambiguity. In addition, the model corrects for attrition bias in the estimated treatment effect.

I. Introduction

In this paper provides an empirical examination of learning under ambiguity in the context of clinical trials. Bayesian models have been widely used to capture consumer learning for experience goods [see Erdem and Keane (1996), Ackleberg (2006), Crawford and Shum (2006), Coscelli and Shum (2004)]. Learning is accomplished by observing a series of signals from a single known source and updating one's beliefs according to Bayes' Rule. In the case of clinical trials, the issue of learning is complicated by patient randomization into experimental and control groups. Patients are uncertain as to their true group assignment (the source of the signals); therefore, changes in health cannot be immediately attributed toward the experimental treatment. I propose a two dimensional model of learning to explain how patient beliefs on drug quality and group assignment evolve. In addition to changes in health, the model uses side effects to capture both an explicit price for participating in the experiment and an implicit source of group information. Variation in the rate of side effects between treatment groups can be used by patients to infer group assignment. A secondary goal of this paper is to use the structural model, along with side effect experiences, to study patient attrition. The model is used to answer the question, "how do patients' beliefs affect attrition decisions and, consequently, treatment effect estimates?"

Clinical trials serve as a powerful setting to study learning for several reasons. First, randomization requires that treatment groups be composed of individuals with similar characteristics and beliefs. Randomization allows the investigators to extrapolate that differences in health seen between the two groups at later intervals of the trial are due to treatment effects and not to unobserved variations in the groups. By the same token, differences in patients' behavior observed in later stages of the trial could be used to infer patient learning. Second, clinical trials offer an exclusive good (the experimental drug). Patients can only gain experience with the product by participating in the trial. Observational data relies on consumers exclusively shopping at the firms available within the sample, which may not necessary hold. This property emphasizes the advantage experimental data has over obser-

vational data, where unobserved variables could potentially have a large effect on learning. Lastly, the common initial condition problem associated with dynamic models does not apply to a clinical trial. The parameters of the experiment are well defined at the onset of the trial. Patients are informed of the initial treatment probability, the possible length of treatment, and compensation prior to beginning the trial.

The model utilizes data from two clinical trials: (1) AZT as a treatment of HIV and (2) Topamax as a treatment for alcoholism. The first clinical trial is the AIDS Clinical Trials Group (ACTG) 019. The ACTG 019 was conducted during the early stages of AIDS research and is the primary clinical trial responsible for investigating the first treatment for HIV infection. The results of this clinical trial found AZT to be a viable treatment for short-run improvements in the health of HIV patients. Few researchers have questioned the efficacy of AZT to treat HIV, but some are concerned with the safety of AZT, as 80% of the patient sample in this experiment fail to complete the full 80 week treatment. The pattern of attrition is unusual because patients in the treatment groups are observed to leave the trial at a faster rate than patients in the placebo group. Figure 1 illustrates the survival rate of each treatment group over the duration of the trial.

The second dataset studies the effect of an anti-seizure medication, Topamax, on alcoholism. The medication is found to alleviate some dependency on alcohol at the cost of experiencing side effects. Unlike the side effects in the AZT trial, side effects associated with Topamax are generally mild and well tolerated. The attrition pattern in the Topamax study has patients in the placebo group leaving at a higher rate than patients in the treatment group. The second trial provides an interesting comparison between a life threatening disease, AIDS, and a long term addiction, alcoholism. The two medications also differ in the intensity of their side effects.

The results of this paper demonstrate that patient learning is slower than previously reported. The additional ambiguity faced by patients causes posterior treatment effect estimate to converge at a slower rate. The combination of side effects and a group assignment

uncertainty leads to early attrition in both trials, but later attrition is attributed to low posterior means of the treatment effect. The structural model also indicates that attrition bias is present in both experiments. The placebo health trend is found to be over estimated in the AZT trial and the treatment effect is found to be under-estimated in the Topamax study.

II. Literature Review

The relationship between attrition decisions and patient learning is initially studied by Philipson and DeSimone (PD, 1997). PD recognize that if attrition behavior is the result of an exogenous process, then randomization in RCT's should imply homogenous attrition rates across treatment groups. The authors reject the null hypothesis of homogenous attrition rates in a collection of clinical trials on substance abuse. This result motivates the study of attrition as a utility maximizing decision.

The proposed model is an extension to Chan and Hamilton (CH, 2006) who investigate attrition behavior in a clinical trial on HIV combination treatments (ddI, AZT+ddI, AZT+ddC) with an active control (AZT alone). The authors offer a framework for evaluating randomized experiments in light of subject learning. In this study, a structural model of attrition incorporating patient learning is constructed and estimated. Side effects enter as unobservable group characteristics leading to disutility, but do not affect patient learning. Subject learning is found to be rapid and significant; however early attrition behavior is inferred to be the result of unobserved negative side effects. The authors find evidence that when accounting for unobserved side effects, the treatment that maximizes utility is not always the most medically effective treatment.

The model of patient behavior described in this paper provides several extensions to CH. The most significant of these extensions is constructing a learning model to allow for uncertainty in both group assignment and treatment effect. A patient's belief on group assignment is allowed to affect her attrition decision. In Chan and Hamilton (2006), the

model of learning is one dimensional in that the patient is concerned about her health expectations and not concerned about her group assignment. This non-trivial extension captures changes in patient’s behavior associated with changes in the initial treatment probability. For example, a patient who is initially informed that her chance of receiving the experimental treatment is 90% is less likely to exit the trial than if she were given only a 10% chance. If a patient believes she may receive the experimental drug with a 90% chance, then she may remain in the trial longer for fear of forgoing actual treatment, even when initial health benefits appears to be small. Second, the model incorporates observable side effect experiences as informative signals of treatment group assignment and a source of disutility for participating in the trial. Some clinical trials have gone as far as to include an active ingredient in the placebo medication to mimic the side effect present in the experimental medication in order to reduce the likelihood of patients learning their treatment assignment¹.

III. Data

Data from two clinical trials are used to evaluate patient learning. The first dataset is an experiment testing the effectiveness of AZT to improve white blood cell counts in HIV patients. The second clinical trial tests the effect of the anti-seizure medication, Topamax, on alcohol addiction. These two datasets provide a good comparison in that both are double blinded placebo controlled trials, but the difference in side effect severity between trials is large. In both cases, side effects are modelled to provide information about group assignment, but in the HIV trial the side effects may provide a significant level of disutility forcing patients out of the trial regardless of the information content. The two experiments also differ in disease severity because HIV patients in the 1980s had life expectancies, after diagnosis, of 2.5 years, while alcoholics can have much longer expected life spans².

A. *ACTG 019*

The ACTG 019 is a double-blinded, placebo-controlled Phase III study carried out between 1987-1989³. The study evaluates the efficacy and safety of AZT at two daily doses (500mg and 1500mg) to prevent or delay the progression of AIDS in persons with asymptomatic HIV infection. The public measure of health is CD4 counts, which are white blood cells in the body that identify and defend against infection. The HIV virus inhibits the production of these white blood cells. A randomized controlled trial of 1,338 subjects was conducted in which patients were given free medical services as compensation for their participation in the trial. Patients are observed at baseline, 8, 16, 32, 48, 64, and 80 weeks. ACTG 019 is a multicenter clinical trial drawing patients from 11 different states across the nation⁴. Originally, 428 patients are randomized into the placebo group, 453 patients into the low dose group, and 456 into the high dose group⁵.

The several measures of health are recorded such as counts for hemoglobin, platelets, granulocytes, and CD4 white blood cells. In an effort to control for sample selection into the clinical trial, the investigators required that patients have a CD4 count ≤ 500 cells/mm³. At the onset of the trial, 12% of the patients are found to fit the AIDS criteria⁶. The pattern of CD4 progression is illustrated in Figure 2. Initially, both experimental groups, on average, have a greater improvement in health, but these improvements appear to dissipate towards the end of the trial leaving the placebo group with the greatest improvement in overall health.

The investigators also collect demographic information for each patient, which includes the variables of race, sex, and method of initial infection. The descriptive statistics for these measurements are found in Table 1. Given the randomization process in the clinical trial, it is not surprising to find similar measurements of the mean across treatment groups for demographic variables. An important demographic to highlight is the method of infection. The method of infection is classified into three groups: homosexual experience, bisexual experience, and intravenous (IV) drug use. Over 70% of infections occurred through a

homosexual experience between 1986-89. In 2006, the highest prevalence of HIV infections is among female heterosexuals (typically sex workers). Given the changes in the HIV population, the reader should treat any pharmacological results found in this study with some reserve.

Side effect experiences are recorded. The side effect measures are aggregated into two classifications: blood related side effects and liver related side effects. Blood related side effects include anemia, neutropenia, and thrombopenia. Liver side effects include hepatic toxins, elevated liver enzymes (SGOT, SGPT), and bilirubin. These side effects are measured via blood samples taken at fixed intervals throughout the trial. For the purpose of this study, side effects are defined as the occurrences of negative health shocks, either blood or liver, in any group. Traditionally, a side effect is measured as the *difference* in negative health shocks between groups. The definition used here allows for side effects to occur in the placebo group. Placebo group side effects are assumed to occur as natural events associated with the disease. Side effects are experienced by 30.4% of patients receiving treatment. Side-effects may convey information about treatment group assignment. To illustrate this point, consider the difference in anemia occurrences across treatment groups (see Table 1). Clearly, if a patient does experience anemia, then she could infer that she is not in the placebo group, as no one in the placebo group reports anemia as a side effect.

All treatment groups experience substantial attrition with less than 20% of the sample surviving to week 80. If attrition is an exogenous process, then attrition rates should be identical across groups. In this sample, attrition rates are heterogenous and follow a specific ordering from lowest to highest: placebo, low dose, and high dose. The average patient spell is 33 weeks. The average patient spell within a treatment group is 28 weeks in the high dose group, 34 weeks in the low dose, and 38 weeks in the placebo group.

B. Topamax

The second dataset is from a randomized double blinded study testing the use of Topamax to treat alcohol dependence. Topamax (Topiramate) was originally approved by the FDA in 1996 to treat seizures, but has recently been tested to treat various types of addiction including alcoholism. In this study, 150 patients participate in a 12 week study where half of the patients are randomly assigned to receive either placebo or Topamax. In addition to the study medication, patients also participate in therapy sessions, receive free medical services, and are given weekly compensation of \$20 per visit. Prior to consent, patients are informed of the possible risks including potential side effects associated with Topamax, which include drowsiness, dizziness, slurred speech, and slowing of motor skills.

The study collects both medical and demographic information about each patient. Demographic data includes drug use, employment status, race, age, height, weight, criminal record, and family history. Table 2 provides a comparison of means for some observable demographic variables. The average patient is a 42 year old male with a body mass index (BMI) of 26.12, 13.53 years of schooling, and an annual income of \$39,000. Most patients hold at least a high school diploma and have careers in various working sectors including Clerical and Sales (25%), Administration (24%), and Skilled Manual Labor (18%)⁷. The racial profile of patients in this experiment is white (60%), Mexican (29%) and other (11%). The distribution of race in this study is more a reflection of location, Texas, than of the population of alcoholics.

Alcohol dependence is monitored using a patient's Gamma-glutamyl transferase (GGT) level. The GGT test provides a better measure of alcohol dependence than a blood alcohol level (BAL) measure because a patient would need to abstain from alcohol for 4-5 weeks to reach normal GGT levels. A patient need only abstain from alcohol for several hours to affect the BAL measure. While GGT levels are initially higher in the experimental group than the placebo group, the largest decrease in GGT levels are observed in the experimental group. Baseline GGT values (81.8 Topamax and 65.3 Placebo) are compared with values

recorded during the twelfth visit (57.8 Topamax, 52.5 Placebo)⁸.

The study documents side effect events by recording the date of onset, duration, severity, and action taken for each type of side effect. The likelihood of experiencing at least one side effect is about the same in both groups, Topamax (47.33%) and Placebo (46%). The average number of side effects reported by patients in each group is 4.65 in Placebo, and 7.78 in Topamax. The most common side effects experienced in the study are parasthesia (18%), somnolence(16%), anxiety (13%), fatigue (12%), and weight loss (10%)⁹. Side effects may have both direct and indirect consequences on patient participation. The direct effect of side effects is disutility associated with pain or discomfort. An indirect effect associated with side effects is learning in that a patient may infer her group assignment through side effect experiences. For example, parasthesia is eight times more likely to occur in the treatment group than in the placebo group. Uninformative side effect signals include anxiety and somnolence, as they are equally likely to occur in either treatment group.

Overall, 44% of placebo patients and 28% of Topamax patients exit the trial. On average, patients are 9.7% more likely to leave the placebo group than the Topamax group during any given week. A formal test where the null hypothesis is homogenous attrition rates across groups is rejected at the 1% level using a paired Student t test ($T = -4.96$, $p\text{-value} = .001$). Two probable reasons for heterogenous attrition rates are side effects and learning. If patients are sensitive to side effects, then, holding all else equal, as side effects increase, so should the attrition rates of patients in the group experiencing side effects more often. If side effects are the predominant factor of attrition, then one would expect to find higher attrition rates in the Topamax group instead of the placebo, but attrition is observed to be higher in the placebo group. The second reason for heterogenous attrition rates is learning. Patients receive signals on the treatment effect of the drug through observed changes in health. If a patient does not experience a significant improvement in health, then she may infer that her group assignment is placebo, and would exit the trial in favor of the outside medical option. Attrition as a result of learning would lead to higher attrition rates in the placebo group, as

is observed in the data.

IV. Structural Model

A. *Experimental Design*

This section describes the organization of a clinical trial, defines patient preferences, and describes the mechanism of Bayesian Learning used to update beliefs. Clinical trials are a type of experiment where participants are randomly selected into either a control or treatment group, but are blinded (uncertain) to their actual group assignment. Randomization creates two groups, which are statistically identical, allowing researchers to attribute future changes in health outcomes to treatment efficacy, and not to differences in group characteristics. In this setting, treatment efficacy is measured as the mean difference in health outcomes between the treatment group and the placebo group. The term "double-blinded" implies that both physicians and patients are unaware of group assignment (treatment or placebo). In principle, double blinding ensures that neither physicians, nor patients may bias health outcomes through private motivations. Blinding introduces the ambiguity of group assignment to the patient, which makes patient learning potentially valuable. Two types of controls are commonly used in clinical trials: placebo and conventional control. A placebo contains no active ingredient, while a conventional control contains an active ingredient that is the approved standard of treatment outside of the trial. Therefore, patients who participate in a placebo controlled trial are more motivated to learn their treatment assignment than those who participate in standard control trials for fear of forgoing actual treatment available outside of the trial.

At the onset of the trial, patients are made aware of the number of possible treatments, the probability of receiving each treatment, treatment duration, and the type of compensation. Patients receive free medical consultations and/or monetary allowances as compensation for participating in the experiment. Once the trial has commenced, measurements on patient health and side effect experiences are taken at fixed time intervals. At

the end of each period, a patient may display one of two types of behavior: (1) exiting the trial (attrition) or (2) remaining in the trial and consuming the prescribed amount of the assigned drug (compliance)¹⁰.

B. Patient Preferences

Consider a clinical trial in which there are G treatment groups. Patients are randomly assigned to treatment group g with a fixed probability, $1/G$. In each treatment period t , patient i observes her health H_{igt} , two side effect measures $(S1_{igt}, S2_{igt})$, and the price of the outside medical option, p_{it} . Health and side effects are assumed to be distributed joint normal.

$$(1) \quad \begin{bmatrix} H_{igt} \\ S1_{igt} \\ S2_{igt} \end{bmatrix} \rightsquigarrow N \left(\begin{bmatrix} H_{i0} + \theta_{it} + \theta_{gt} \\ \bar{s}1_{igt} \\ \bar{s}2_{igt} \end{bmatrix}, \Omega = \begin{bmatrix} \sigma_\varepsilon^2 & \rho_1 \sigma_{u1} \sigma_\varepsilon & \rho_2 \sigma_{u2} \sigma_\varepsilon \\ \rho_1 \sigma_{u1} \sigma_\varepsilon & \sigma_{u1}^2 & \rho_3 \sigma_{u1} \sigma_{u2} \\ \rho_2 \sigma_{u2} \sigma_\varepsilon & \rho_3 \sigma_{u1} \sigma_{u2} & \sigma_{u2}^2 \end{bmatrix} \right)$$

Patients are assumed to know certain components of their health equation. In particular, a patient knows her initial level of health, H_{i0} , the progression of the disease without treatment, $\theta_{it} = \theta_1 + \theta_2(t-1) + e_i$, but is uncertain about the group specific experimental effect, $\theta_{gt} = \theta_3 d_g + \theta_4(t-1) 1_{g=\text{treatment}}$. After observing prices, health, and side effects in a given period, the patient decides whether to remain in the trial or consume the outside option.

Patients are assumed to be forward looking agents and make the discrete choice of attrition by maximizing the expected discounted sum of future utility flows. A clinical trial represents a finite horizon dynamic problem where decisions are made conditional on a patient's information set, I , at time (t)

$$(2) \quad \max_{T^*} E \left[\sum_{\tau=t}^{T^* \leq T} \beta^{\tau-t} \sum_{g=1}^G \delta_{igt} U_{ig\tau} + \frac{\beta^{T^*}}{1-\beta} V^0(T^*, \delta_{igt}) \mid I \right]$$

where a patient's choice of attrition time is $T^* \leq T$ (T represents the patient remaining

in the trial), δ_{igt} is patient i 's belief on being in treatment group g at time t , and β is the discount factor. The outside option is defined as

$$(3) \quad V^0(t, \delta|I) = c_1 + c_2t + \eta_1\delta + \eta_2t\delta + v_{iot}$$

where c , is a set of parameters to be estimated, capturing the utility of alternative medical options available outside of the trial during each week . The value of the outside option is also allowed to vary based on a patient's group assignment beliefs through the parameters η . The term $(c_2 + \eta_2\delta)t$ captures long term effects of participating in the trial. After observing health and side effects in a given period, the patient decides whether to remain in the trial or consume the outside option¹¹.

Let the single period utility function for a patient in group (g) be defined as

$$(4) \quad U_{igt} = -\frac{1}{\gamma} \exp(-\gamma H_{igt}) - a_1 S1_{igt} - a_2 S2_{igt} + \xi_i + v_{igt}$$

which is a constant absolute risk aversion (*CARA*) utility function with a linear portion capturing side effects and the opportunity cost of forgoing the outside option^{12, 13, 14}. The parameters γ and a represent the patient's sensitivity to risk and side effects, respectively. The coefficient of absolute risk aversion is given by the parameter $\gamma > 0$. Risk is an important aspect of learning models as there is a trade-off to waiting an additional period and forgoing the outside option to learn more about product quality. The outside option represents a certainty equivalent offered to the patient. When the certainty equivalent is at least as large as the expected value of the utility flows, it is optimal for a risk averse patient to choose the outside option. There are two sources of unobserved heterogeneity, ξ_i and v_{igt} . The first error term, ξ_i , is a normally distributed person specific error with mean zero that captures a patients unobserved value for participating in the trial. This error is observed by the patient but unobserved by the econometrician. The second error term, v_{igt} , is assumed to be distributed Type I extreme value and captures unobserved changes on the

patient's outside option. This error is revealed to the patient at time t , but unobserved to the econometrician.

A patient's information set is defined as $I_{it} = \{\mathcal{H}_{it}, \mathcal{S}_{it}, t\}$, where $\mathcal{H}_{it} = [H_{i1} - H_{i0} - \theta_{i1} - e_i, \dots, H_{it-1} - H_{i0} - \theta_{it-1} - e_i]$ is a patient's health history and $\mathcal{S}_{it} = [S1_{i1}, \dots, S1_{it}, S2_{i1}, \dots, S2_{it},]$ is a patient's side effect history. A patient learns her specific experimental effect, μ_{igt} , and treatment group assignment, δ_{igt} , through an application of Bayes' Law on the variables in her information set.

C. Patient Learning Process

At the beginning of the trial, patients construct prior beliefs about drug quality.¹⁵ The experimental effect prior, μ_{ig0} , is assumed to be distributed bivariate normal

$$(5) \quad \mu_{ig0} = \begin{bmatrix} \lambda 1_{i0} \\ \lambda 2_{ig0} \end{bmatrix} \sim N \left(\underline{\mu}_{-g0} = \begin{bmatrix} \underline{\lambda 1}_0 \\ \underline{\lambda 2}_{g0} \end{bmatrix}, \Sigma = \begin{bmatrix} \sigma_{\lambda 1}^2 & 0 \\ 0 & \sigma_{\lambda 2}^2 \end{bmatrix} \right)$$

where $\underline{\lambda 1}_0$ is the prior dose effect and $\underline{\lambda 2}_{g0}$ is the treatment group specific time trend. Patients are informed on the initial randomization probability, $\delta_{g0} = 1/G$, which serves as their prior on treatment group assignment. Formally, $\delta_{igt} = \Pr(g = g_i | I_{it})$ where $g \in \{\text{Placebo, Low Dose, High Dose}\}$ in AZT study and $g \in \{\text{Placebo, Treatment}\}$ in the Topamax study. The variable g_i is a patient's actual treatment group. Patients are assumed to be Bayesian learners and use these priors to update their beliefs. Let posterior beliefs be indicated with a superscript ($'$).

If the treatment effects for each group are known by the patients, but treatment assignment is unknown, then Bayes' Law can be used to learn group assignment. Let $f(\cdot)$ be the probability density function of health for patients in group g . A patient updates her belief on group assignment according to Bayes' Law

$$(6) \quad \delta'_j = \Pr(g = g_j | x) = \frac{\delta_j f(x | g = g_j)}{\sum_{i=1}^G \delta_i f(x | g = g_i)}$$

where the posterior belief on group assignment, δ'_j , is dependent on both the relative likelihood of observing a series of signals, x , given the patient is in group g and the patient's prior belief of being in group g . Given the discrete nature of group assignment and the continuous nature of health, patients initially assume health is drawn from a normal mixture distribution. The normal mixture distribution is the weighted sum of G distinct normal distributions

$$(7) \quad L(H_{it}, S_{it} | \mu_{igt}(\mathcal{H}_{it}, \mathcal{S}_{it}), \delta_{igt}(\mathcal{H}_{it}, \mathcal{S}_{it}), \Sigma, \bar{s}1_g, \bar{s}2_g, \Omega) \\ = \sum_{j=1}^{G=3} \delta_{igt}(\mathcal{H}_{it}, \mathcal{S}_{it}) \int_{\lambda \in \Lambda} \phi(H_{it}, S_{it} | \lambda, \Omega, \bar{s}1_g, \bar{s}2_g) \phi(\lambda | \mu_{igt}(\mathcal{H}_{it}, \mathcal{S}_{it}), \Sigma, g = j) d\lambda$$

where Λ is the support of possible experimental effect values (λ), $L(\cdot)$ is the likelihood function over the unknown parameters, $\phi(\cdot)$ is the normal probability density function, and treatment group beliefs are restricted to satisfy $\sum_{j=1}^G \delta_{igt} = 1$. The normal mixture model captures three concepts: [1] health within a group is normally distributed; [2] the patient is uncertain as to the parameter values of the normal distribution (ie mean and variance of health for different treatment groups); [3] the patient recognizes that there are G possible processes generating her health outcomes.

Given the patient's likelihood function on health, one would normally proceed by calculating the posterior distribution on experimental effects and treatment group assignment conditional on an observed value of health. Unfortunately, a closed form solution to the posterior distribution of equation (7) does not exist¹⁶. A continuous random variable requires a conjugate prior for a closed form solution of the posterior distribution, but the posterior distribution is always defined for discrete random variables. Therefore, I discretize the probability state space defined by equation (7) for a finite number of health, side effect, and experimental effect states conditional on group assignment and a patient's priors¹⁷.

Once the probability space is discretized, an application of Bayes' Law provides the joint

posterior distribution on beliefs.

$$(8) \quad \Pr\left(\mu'_{igt}, \delta'_{igt} | \mathcal{H}_{it}, \mathcal{S}_{it}, \Sigma, \bar{s}1_g, \bar{s}2_g, \Omega\right) = \frac{w_{igt}(\lambda = \mu_{igt})}{\sum_G \sum_\Lambda w_{igt}(\lambda)}$$

$$w_{igt}(\lambda) = \Pr(\mathcal{H}_{it}, \mathcal{S}_{it} | \lambda, \bar{s}1_g, \bar{s}2_g, \Omega, \delta_{igt-1} = 1) \Pr(\lambda | \underline{\mu}_{g0}, \Sigma, \delta_{igt-1} = 1) \Pr(\delta_{igt-1} = 1)$$

For each patient, I store a set of probabilities corresponding to each $(\mu'_{igt}, \delta'_{igt})$ beliefs pair. The set of probabilities provides a finite approximation to the continuous joint distribution described in (7). As the number of pairs tends to infinity, the discrete approximation converges towards the continuous distribution. Given the joint posterior distribution on beliefs, a patient's expected utility is defined as

$$(9) \quad E_t [U | I_{it}] = \sum_G \sum_\Lambda U(\mu'_{igt}, \delta'_{igt}) \Pr\left(\mu'_{igt}, \delta'_{igt} | \mathcal{H}_{it}, \mathcal{S}_{it}, \Sigma, \bar{s}1_g, \bar{s}2_g, \Omega\right)$$

where the expectation is taken over all the possible treatment group and experimental effect states conditional on a patient's information set at time t. Patients must also consider future values on utility. Therefore, a predictive posterior distribution is required to evaluate the expected value of utility at time t+k for k>0. The k period predictive posterior distribution is

$$(10) \quad E_t [U | I_{it+k}] = \sum_G \sum_\Lambda U(\mu'_{igt+k}, \delta'_{igt+k}) (V_{igt+k}) \pi(\mathcal{H}_{it+k}, \mathcal{S}_{it+k})$$

$$V_{igt+k} = \left[\sum_{\mathcal{H}} \sum_{\mathcal{S}} \Pr\left(\mu'_{igt+k}, \delta'_{igt+k} | \mathcal{H}_{it+k}, \mathcal{S}_{it+k}, \Sigma, \bar{s}1_g, \bar{s}2_g, \Omega\right) \right]$$

$$\pi(\mathcal{H}_{it+k}, \mathcal{S}_{it+k}) = \Pr(\mathcal{H}_{it+k}, \mathcal{S}_{it+k} | b_g, l_g, \Omega, \mu'_{igt+k-1}, \delta'_{igt+k-1}) \Pr\left(\mu'_{igt+k-1}, \delta'_{igt+k-1}\right)$$

where $\pi(\mathcal{H}_{it+k}, \mathcal{S}_{it+k})$ is the distribution of future health outcomes conditional on having the belief pair of $(\mu'_{igt+k-1}, \delta'_{igt+k-1})$ and V_{igt+k} is the probability of observing the belief pair of $(\mu'_{igt+k-1}, \delta'_{igt+k-1})$ conditional on past health outcomes. A patient iterates the predictive posterior distribution to evaluate the expected value of future utility k periods into the

future.

Illustration.—Consider the following example of learning under ambiguity. Assume a hypothetical clinical trial with two groups, two experimental effects, and two outcomes for each observable health variable. The health variable, A , is equal to 1 if the patient has AIDS, and zero otherwise; S is equal to 1 if the patient experiences a side effect, and zero otherwise. Initially, patients are randomized into either the placebo (P) or experimental (E) group with $Pr(P) = 0.5$. There are two possible treatment types within each group, $u_{low} < u_{high}$. The probability of observing a low experimental effect in the placebo (experimental) group is $Pr(u = u_{low}|P) = .8$ ($Pr(u = u_{low}|E) = .2$). These probabilities serve as the patient’s initial prior on the distribution of groups and treatment types. Lastly, the likelihood of observing any pair of health outcomes conditional on treatment type and group is given by a discrete distribution found in each column within the table below¹⁸.

	Pr(P)=.5		Pr (E)=.5		Total
	Pr($u_{low} P$)=.8	Pr($u_{high} P$)=.2	Pr($u_{low} E$)=.2	Pr($u_{high} E$)=.8	
A=0 S=0	(.1) (.8) (.5)	(.2) (.2) (.5)	(.2) (.2) (.5)	(.2) (.8) (.5)	.16
A=1 S=0	(.6) (.8) (.5)	(.5) (.2) (.5)	(.1) (.2) (.5)	(.1) (.8) (.5)	.34
A=0 S=1	(.1) (.8) (.5)	(.1) (.2) (.5)	(.5) (.2) (.5)	(.6) (.8) (.5)	.34
A=1 S=1	(.2) (.8) (.5)	(.2) (.2) (.5)	(.2) (.2) (.5)	(.1) (.8) (.5)	.16
Total	.4	.1	.1	.4	1

A patient improves her utility by increasing her health and decreasing her uncertainty. A patient is uncertain as to which group and treatment type she is (i.e. in which cell of the matrix she is). She resolves this uncertainty by using Bayes rule to update her beliefs about treatment type and group¹⁹. The probability of a given (treatment type, group) pairing is captured by a binomial distribution. The prior variance over beliefs is $var(G, u) = .66$ ²⁰. If the patient observes the health outcome $\{A=1, S=0\}$, then her posterior beliefs on the joint probability of being both in the placebo group and experiencing a low health return is $Pr(u_{low}, P) = \frac{Pr(A=1, S=0|u_{low}, P)}{Pr(A=1, S=0)} = 0.7059$. The posterior variance over

beliefs becomes $var(G, u) = .4654$. There are a few noteworthy points about this exercise. First, the exercise illustrates that learning can decrease patient uncertainty, thus improving patient utility. Second, the exercise reinforces the need to model the joint distribution of treatment types and group assignment. Initially, treatment types and group assignments are distributed independently (as in the table above), but after receiving the first signal the posterior probability of a low treatment type is dependent on the posterior group assignment belief (the table below).

	Placebo		Experimental		Total
	u_{low}	u_{high}	u_{low}	u_{high}	
A=0 S=0	(.1) (.7059)	(.2) (.1470)	(.2) (.0295)	(.2) (.1176)	.1294
A=1 S=0	(.6) (.7059)	(.5) (.1470)	(.1) (.0295)	(.1) (.1176)	.5118
A=0 S=1	(.1) (.7059)	(.1) (.1470)	(.5) (.0295)	(.6) (.1176)	.1706
A=1 S=1	(.2) (.7059)	(.2) (.1470)	(.2) (.0295)	(.1) (.1176)	.1882
Total	.7059	.1470	.0295	.1176	1

D. Value Function

Given the learning model described in the previous section, the attention is returned to the patient's dynamic discrete choice problem. The experimental design of a clinical trial lends itself to the use of a dynamic discrete choice framework to model attrition behavior. Patients are informed of the start and end dates of the experiment. Patients are informed of the type of compensation they may receive and the discontinuance of compensation if they are to exit the trial without the option to return. Each period, the patient may take one discrete action, compliance or attrition. Given the patient's information set, $X_{it} = [\mu'_{igt}, \sigma^2_{igt}, \delta'_{igt}, t, e_i, \epsilon_i]$, a patient's value function for remaining in the trial at treatment period t can be represented by the following value function $V_t(X_{it}) = \max_{D_{T^*} \in \mathcal{D}_t} V(X_{it}, D_{T^*})$ where $D_{T^*} = [d_t = 0, d_{t+1} = 0, ..d_{T^*} = 1, ..d_T = 1]$ is the unique sequence of discrete choices such that patient i expects to exit the trial at time $T^* \leq T$. The sub-value function,

$V(X_{it}, D_{T^*})$, is defined

$$(11) \quad V(X_{it}, D_{T^*}) = E_t[U|X_{it}] + \beta E_v \left[\max \left(E_x [V(X_{it+1}, D_{T^*}) | X_{it}], V^0(t+1, \delta' | I) \right) \right]$$

and the value function in the last period, T , is $V(X_{iT}, D_T) = E_t[U|X_{iT}]$. The discount factor $\beta \in [0, 1)$ represents the patience of a patient between periods, where individuals with low values of β place greater weight on the current level of utility rather than future levels of utility. The case of $\beta = 0$ is the myopic patient who only maximizes current utility flows. The case of $\beta = 0.98$ is the "forward looking" patient who values learning more than the myopic patient. A forward looking patient is willing to trade low health returns in the early stages of the trial for more health signals. These additional signals are used to update beliefs and reduce uncertainty, thus increasing utility. In this paper, I examine learning under both specifications of the discount factor. The expectations in equation (11) are first taken over the distribution of beliefs, and then taken over the future values of v_{igt} ²¹. Recalling that $v_{igt} \rightsquigarrow EV\left(0, \frac{\tau^2 \pi^2}{6}\right)$, a closed form solution for the expected value of the future period's value function exists as

$$(12) \quad \begin{aligned} & E_v \left[\max \left(E_x \left[\bar{V}(X_{it+1}, D_{T^*}) - v_{igt+1} | X_{it} \right], 0 \right) \right] \\ &= \tau \left[\Gamma + \ln \left(1 + \exp \left[\frac{E_x \left[\bar{V}(X_{it+1}, D_{T^*}) | X_{it} \right]}{\tau} \right] \right) \right] \end{aligned}$$

where $V(X_{it+1}, D_{T^*}) = \bar{V}(X_{it+1}, D_{T^*}) - v_{igt+1}$, and Γ is the Euler constant²². The optimal value function can be solved by first taking the expectation with respect to beliefs, and then using (12) to solve (11) backward recursively for each value of T^* . Finally, the value of T^* that maximizes $V(X_{it}, D_{T^*})$ is the optimal value function, $V_t(X_{it})$.

V. Econometric Specification

A. Basic model

This section presents the econometric method used to estimate the attrition model. The likelihood function is comprised of two parts: attrition decisions and the distribution of outcome measurements. In the previous section, a patient's value function is defined as $V(X_{it}, D_{T^*}) = \bar{V}(X_{it}, D_{T^*}) - v_{igt}$, where $X_{it} = [\mu'_{igt}, \sigma^2_{igt}, \delta'_{igt}, t, e_i, \xi_i]$, $V(\cdot)$ is a latent variable, and $d_{it} \in D_{T^*}$ is the observed binary choice, where $d_{it} = 1$ when the patient exits the trial, and 0 otherwise. The error term, v_{igt} , is assumed to be distributed Type I extreme value. The error terms e_i and ξ_i are patient specific errors known to the patient, but unknown to the econometrician. These errors are assumed to be distributed iid normal with variance equal to σ_e^2 and σ_ξ^2 , respectively.

Given the distribution of the errors, the conditional survival probabilities at time t are

$$(13) \quad \begin{aligned} \Pr(d_{it} = 0 | d_{it-1} = 0, e_i, \xi_i) &= \frac{\exp(\bar{V}_{it}/\tau)}{1 + \exp(\bar{V}_{it}/\tau)} \\ \Pr(d_{it} = 1 | d_{it-1} = 0, e_i, \xi_i) &= \frac{1}{1 + \exp(\bar{V}_{it}/\tau)} \\ \Pr(d_{it} = 0 | d_{it-1} = 1, e_i, \xi_i) &= 0 \\ \Pr(d_{it} = 1 | d_{it-1} = 1, e_i, \xi_i) &= 1 \end{aligned}$$

The "no re-entry" policy of an RCT is captured by the last two equations. The unconditional probability of attrition at time t is then given by equation (13)

$$(14) \quad \Pr(d_{it} = 1 | e_i, \xi_i) = \frac{1}{1 + \exp(\bar{V}_{it}/\tau)} * \left[\prod_{j=1}^{t-1} \frac{\exp(\bar{V}_{ij}/\tau)}{1 + \exp(\bar{V}_{ij}/\tau)} \right]$$

The second part of the likelihood function is the distribution of health and side effects. Define the vector of deviations between observed and predicted outcomes as $Z_{it} =$

$[e_i + \varepsilon_{ig1}, e_i + \varepsilon_{ig2}, \dots, e_i + \varepsilon_{igt}, u1_{ig1}, \dots, u1_{igt}, u2_{ig1}, \dots, u2_{igt}]'$. The deviations are stacked in the following order: health, blood related side effects, and liver related side effects. Given the assumption of joint normality between health and side effects, the probability density function of patient i 's outcome measurements is given by

$$(15) \quad f(Z_{it}, \Omega) = \frac{1}{(2\pi)^{3t/2} |\Omega|^{1/2}} \exp \left[-\frac{1}{2} Z'_{it} \Omega^{-1} Z_{it} \right]$$

where t is the period when patient (i) exits the trial, and Ω is the health covariance matrix. The total number of observations for patient i is equal to $3t$ as there are three measures of outcomes taken over the t treatment periods.

Define the set of parameters to be estimated as $\Psi = \{\theta, b_g, l_g, \Omega, \Sigma, \mu_g^p, a_1, a_2, \gamma, \eta, \sigma_\xi^2, \tau\}$. The likelihood contribution of the i 'th patient conditional on the unobservable errors is

$$(16) \quad L_i(\Psi | d_{it}, \mathcal{H}_{it}, \mathcal{S}_{it}, e_i, \xi_i) = f(Z_{it}, \Omega | d_{it} = 1, \theta, b_g, l_g) * \Pr(d_{it} = 1 | e_i, \xi_i, \Psi)$$

To remove the unobserved heterogeneity, an expectation of the likelihood contribution is taken over the distributions of the unobserved heterogeneity errors.

$$(17) \quad \begin{aligned} & E(L_i(\Psi | d_{it}, \mathcal{H}_{it}, \mathcal{S}_{it}, e_i, \xi_i)) \\ &= f(Z_{it}, \Omega | d_{it} = 1) \int \left[\int \Pr(d_{it} = 1 | e_i, \xi_i, \Psi) f(e_i) de_i \right] f(\xi_i) d\xi_i \end{aligned}$$

While no closed form solution exists for the expected value in equation (17), I employ simulation methods as suggested by Stern (1994) to integrate out the unobserved patient specific heterogeneity. The simulated value of the log likelihood function is given by

$$(18) \quad \log \widehat{L}_i(\Psi | d_{it}, \mathcal{H}_{it}, \mathcal{S}_{it}) = \log \left(\frac{1}{R^2} \sum_{r=1}^R \sum_{k=1}^R L_i(\Psi | d_{it}, \mathcal{H}_{it}, \mathcal{S}_{it}, e_r, \xi_k) \right)$$

and the parameters of the dynamic model are estimated by $\max_{\Psi} \log \sum_{i=1}^n \widehat{L}_i(\Psi | d_{it}, \mathcal{H}_{it}, \mathcal{S}_{it})$.

B. Identification

The identification of the structural parameters can be divided into three parts: treatment effect parameters, learning parameters, and utility parameters. The variables of interest (CD4 and GGT) and side effects are assumed to be distributed joint normal²³. Treatment effect parameters are identified through variation in observed health and side effect measures conditional on the patient remaining in the trial. The first moments of these variables identify the treatment effect and mean side effect levels. The placebo or baseline health parameters are identified jointly through the health equation and patient attrition decision. A necessary exclusion restriction is to assume that patients know their baseline health parameters, $\theta_{it} = \theta_1 + \theta_2(t - 1) + e_i$. Further, the placebo treatment effect, $\theta_{placebo,t} = 0$, is set to zero. The model can only identify relative differences in health between groups. By normalizing the placebo treatment effect to zero, the absolute experimental treatment effect is identified instead of the relative treatment effect. The treatment effect parameters, $\theta_{g,t}$, only enter the health equation (treatment *beliefs* enter into the attrition decision), but the baseline health parameters, θ_{it} , and unobserved health errors enter into both the attrition model and the health equation.

Identification of the learning parameters relies on the randomization process found in the experiment design. Patient randomization creates statistically identical treatment groups in both demographics and beliefs. In the absence of learning and controlling for side effects, attrition behavior should remain the same across groups. Therefore, learning is identified via heterogeneous attrition behavior between groups after controlling for side effects. The prior constant and prior time treatment effect are separately identified. A patient needs at least two signals to update the prior time treatment effect (two signals are needed to calculate a slope), thus attrition occurring after receiving the first signal (week 3), but before the second signal (week 6) identifies the time treatment effect. The remaining residual in attrition rates across all time periods identifies the prior constant.

The parameters in the utility function are identified by co-variation in observed health,

side effects, and attrition decisions. The parameter of constant absolute risk aversion, γ , is identified by changes in attrition associated with a patient’s posterior beliefs (treatment effect and group assignment) and the functional form of the utility function. The value of the outside option is identified by variation in the external price of treatment and attrition decisions affecting all treatment groups equally. The variance of the extreme value error, v , is separately identified only in the dynamic model, $\beta > 0$, via the functional form of the continuation value in the value function of remaining in the trial (see equation 12). In the static model, $\beta = 0$, the variance of the extreme value error is not separately identified from the utility function parameters. In this case, the variance term is set equal to the estimated value of the variance found from the dynamic model. The variances of the unobserved heterogeneity in utility, ξ , is identified by variation in unobserved patient persistence to either remain or exit the trial, which is constant across time periods.

VI. Results

A. AZT Results

Health Outcomes.—Structural parameter estimates are found in Table 3 and treatment effect estimates are found in Table 4. In both models, the estimates on health outcomes are roughly the same. Liver related side effects appear to occur with the same frequency across groups at an average rate of 10%. On the other hand, the instance of blood related side effects does vary across treatment groups. Patients in the low and high dose groups are 4% and 11% more likely to experience a blood related side effect than patients in the placebo group, respectively. I would expect blood related side effects to have a relatively stronger impression on learning than liver related side effects due to the difference in occurrence across treatment groups. Another interesting outcome is the treatment effect of AZT. The treatment effect is found to be positive and statistically significant, but the magnitude is small. The dose-effect leads to a 2% improvement in CD4 counts and the time effect is half a percent increase in each subsequent period. There appears to be a particular match quality

associated with AZT. The majority of the treatment variation is captured in the patient specific unobserved heterogeneity. Patient unobserved heterogeneity constitutes 71.1% of the unobserved health error. When comparing the OLS and structural model estimates, the largest deviation in estimates is within the placebo health coefficients. OLS states the improvement in health from placebo is 18% with a time trend that is not statistically different from zero. The structural model estimates of placebo health are a 9% improvement in health during the first eight weeks and a 1.2% decrease in health every 8 weeks thereafter. The treatment effect estimate across the models is relatively the same, but it should be noted that the structural model corrects for attrition bias by decreasing the estimated treatment effect.

Prior Distribution.—In both the static and dynamic model, patients’ prior on the treatment effect is initially a negative return followed by an increase in health over time. The static model provides a very pessimistic view of the treatment effect by suggesting that AZT would lead to a 350% reduction in CD4 counts with a 1.1% increase over time. The dynamic model gives an optimistic view with CD4 counts falling by 17% initially, then rising 215% over time. Given these priors, positive health returns are experienced at week 9. Though neither model provides estimates of the prior which are close to the actual treatment effect, both models convey information that patients believe they would become sicker before they became healthier. The prior variance captures the level of confidence patients have in their prior means. The estimated prior variance terms are small relative to the health variance, suggesting that patients’ posterior treatment effect beliefs would converge slowly to the actual treatment effect.

Utility Parameters and the Outside Option.—The difference in utility parameters between the two models is small. The coefficient of absolute risk aversion, γ , is slightly larger in the static model (1.2511) than in the dynamic model (1.014). Yet when comparing risk between patients in the AZT study and the Topamax study, I find HIV patients to be more risk averse. This result highlights the difference in severity between the two diseases.

Though the AZT trial is longer than the Topamax trial, the effect each drug has on utility is considered over the lifetime of the patient. On average, an HIV patient's lifetime is shorter than that of an alcoholic; therefore, risk is more concentrated for an HIV patient within a given time frame than that of an alcoholic with the same time frame. Side effect estimates in the AZT study are very large. The marginal disutility associated with blood related side effects is a decrease of 166 utils. Liver decreased utility by 64 utils. Experiencing a side effect essentially pushes patients out of the study. The parameter estimate for "finishing reward" is 3.2862 utils in the dynamic model and 7.3841 utils in the myopic model. A positive coefficient for this parameter captures the curvature of the survival function with respect to time and suggests the survival function is decreasing at a decreasing rate with time. This result implies that the likelihood of exiting the trial next period decreases the longer a patient remains in the trial.

For HIV patients the outside option is sparse. In the late 1980's, the only treatment available was AZT at 100mg doses. The cost of treatment was estimated to be \$10,000 annually and AZT could only be obtained with a diagnosis of AIDS. The number of individuals with AIDS was increasing over this time period, which increased the supply of AZT and helped improve availability. In order to capture some of these dynamics in the outside option, I include a set of dummy variables capturing a patient's AIDS status and the calendar year. Patients who have AIDS at the onset of the trial have higher values of the outside option than HIV patients. The value of the outside option improves over time in a non-linear fashion. The outside option experiences a small increase in 1988 of .31 utils followed by a much larger increase in 1989 of 1.89 utils when the FDA "fast tracks" the approval of AZT for public consumption. The outside option is also comprised of two additional components: baseline parameter values and belief dependent parameter values. The baseline parameter values capture changes in the outside option that are independent of patient treatment beliefs, such as improvements in medical technology. The baseline parameter values are initially negative (-1.4567) relative to the value of participating in the

experiment, but improve at a rate of 0.5089 every 8 weeks thereafter. The belief dependent component interacts a patient's posterior belief on group assignment, $E_{it}(G = \text{Treat})$, with the constant and time trend of the outside option. Patients who place a higher probability of being in one of the treatment groups find the outside option less appealing over time. Holding all things constant, the outside option would imply that the rate of attrition is decreasing as the patient's posterior probability of receiving treatment rises.

Goodness of Fit.—The predicted versus observed survival probability is plotted in Figure 3. The structural model captures 96% of the variation in attrition choices overall. Over the first 32 weeks of the trial, the model underestimates the survival probability, but it performs well thereafter. The model provides a modest fit when evaluating survival probabilities within treatment groups. Survival probabilities are overestimated (underestimated) if the points lie above (below) the x-axis. On average, the model appears to over state survival probabilities for individuals in the treatment groups and to underestimate survival probabilities for patients in the placebo group.

Learning.—Figure 4 displays the evolution of posterior beliefs by treatment group. Patients in all groups increase their posterior belief of being in the low dose group throughout the trial. The largest increase is seen by patients in the high and low dose treatment groups. The rate of increase differs across groups until week 32 by which time each patient has received three signals. While it should be intuitive to observe posterior probabilities increasing for patients in experimental groups, I am puzzled by the consistent increase made by patients in the placebo group. The same type of behavior is displayed when modelling posterior probabilities of being in the high dose group. In week 8, 20% of patients believe they are in the placebo group. This percentage decreases to under 10% by week 16, and close to zero by week 32. One explanation for these results is the initial strong priors on the treatment effect. Even in the presence of low returns, patients remain confident that the treatment effect would be large over time. When the expected large increase in health does not materialize, patients eventually start to redistribute their beliefs downward toward

the low dose group.

Treatment effect learning is observed to be slow. Note in figure 5, the posterior means do approach the actual treatment effect estimates, but they arrive after week 48. Learning, in the early stages of the trial, is faster for patients in the high dose group, followed by the low dose group. The posterior treatment effect mean falls relatively faster for patients in the high dose group, which indicates a faster rate of learning. Overall, learning in this trial is slow. At the end of the trial, patients in all treatment groups place about a 50% chance of being in either the high or low dose groups and the posterior treatment effect mean dose not significantly vary between groups.

To further explore the slow learning rate, the structural dynamic model is used to simulate the evolution of patient beliefs in a fictitious clinical trial lasting 288 weeks (instead of 80 weeks as in the original trial). Extending the length of the trial to 288 weeks provides patients with more than three times the number of possible signals generated in the original trial. As in the shorter experiment, treatment effect beliefs tended to rise above the estimated treatment effects early in the trial. After week 32, patients adjusted their beliefs downward and eventually "over shoot" the estimated treatment effect. After week 96, treatment effect beliefs increased slowly towards the estimated treatment effect and remained positive for all treatment groups²⁴. Coincidentally, patient group beliefs in the simulated experiment *do not converge* to their actual treatment assignments as would be predicted with a typical model of Bayesian Learning. Instead, posterior group beliefs remain constant with beliefs being equally divided between the low and dose groups.

In the AZT experiment, the estimated treatment effect is close to zero making it difficult for patients to distinguish between placebo and AZT generated health signals. For this reason, a patient's posterior group belief converges slowly (if at all) to the patients actual treatment group assignment. It is possible in this model for posterior treatment effect beliefs to converges towards the within group treatment effect without converging in group assignment beliefs. When this occurs the treatment effect becomes a weighted average

of health signals and priors from all the treatment groups.

Decomposing Patient Attrition.—The model of patient behavior allows for changes in side effects, learning, and the value of external options to affect patient attrition. In this section, I isolate the contribution towards attrition made by these three components. In the model, side effects have both an implicit and explicit effect on attrition. Implicitly, side effects help patients to infer their true group assignment. Explicitly, side effects may be considered the price or direct cost of consuming the experimental drug. I evaluate these two components separately. To capture the explicit cost of side effects on attrition, patient behavior is simulated restricting the parameters on side effect sensitivity, $a_1 = 0$ and $a_2 = 0$, to zero. From these results, I compare hazard rates between the restricted outcome and the unrestricted structural parameters (see table 5). Side effects account for a 19% increase in attrition during the first eight weeks, 14% the next eight weeks, but reduces down to 1% by week 64. The diminishing effect of side effects on attrition is consistent with having fewer unhealthy patients in the trial as time increases.

The same method is used to find the effect of the outside option on attrition. The outside option parameters (except for those interacted with group beliefs) are set to zero and hazard rates are computed. The value of the outside option includes increased access to AZT, lower cost of care, and improvements in medical technology. On average, the outside option reduces the sample size by about 3% every eight weeks. There is one notable deviation at week 16 associated with the approval of AZT as a treatment for individuals with AIDS.

The effect of learning on attrition is found by differencing the overall hazard rate with the side effect and the outside option hazard rates. Initially, learning decreases attrition. The benefit received from reducing uncertainty is greater than the cost of remaining in the trial. These benefits peak in week 16 where learning decreases attrition by 5.4%. The returns from learning begin to diminish in the later periods where the marginal benefit of an additional signal is decreasing. Learning becomes the dominant effect causing attrition in

the last two periods of the experiment. Table 5 displays how hazard rates change throughout the trial across the three components.

The effect of learning can further be disaggregated into learning from health signals and side effects. In clinical trials, the practice of altering the placebo to display side effects is becoming more accepted. The motive is to minimize patient learning derived from differences in side effect intensity between the control and experimental drug. To separate the explicit and implicit effect of side effects on attrition, I simulate patient behavior under the restriction that placebo side effects are equal to the low dose treatment side effects²⁵. Hazard rates are then calculated using the previously described method. The results suggest that if the difference in side effect intensity is removed, then more learning increases attrition. The mean hazard rate of learning on attrition without side effects is a 4.3% decrease in the sample each evaluation period (8 weeks). The difference in the learning hazard rate with and without side effect is on average a 3.5% decrease in attrition. Therefore, the implicit effect of side effects is a 0.8% decrease in attrition. Clinical trial designers must then face a trade-off between heterogenous attrition rates when the placebo is not altered, and increased overall attrition when the placebo is altered.

B. Topamax Results

Health Outcomes.—In this section, I present estimates of the structural parameters for the Topamax study found in Tables 6 and 7. Two models of behavior are estimated to provide a comparison between the "forward looking" patient and the myopic patient. In the dynamic (static) model, the weekly discount factor is fixed at $\beta = .98$ ($\beta = 0$). To allow for comparisons between models, the static model of attrition is estimated as a standard logit, where the variance of the extreme value error is set equal to the estimated variance of the extreme value error in the dynamic model. First consider the health outcome parameter estimates. In both models, the estimates on health outcomes are roughly the same. Individuals in the treatment group are 4% more likely to experience fatigue (side effect 1)

and 10% more likely to experience parathesia (side effect 2) than individuals in the placebo group. Given these estimates on side effect experiences, I would expect parathesia to better convey group assignment, as the difference in means between treatment groups is greater than the side effect of fatigue.

The estimated health coefficients represent the marginal effect of each explanatory variable on the natural logarithmic change in health ($\Delta \log(GGT)$)²⁶. Initially, patients in the Topamax group experience a 6% (TE) increase in GGT levels relative to patients in the placebo group, but in each subsequent period, GGT levels fall 7% (TE Time) faster in the Topamax group than in the placebo group. Unobserved heterogeneity in health between patients is primarily driven by the patient specific error. Patient unobserved heterogeneity constitutes 70.5% of the unobserved health error. When comparing the OLS and structural model estimates, the largest difference in coefficients is the treatment-time effect. The reduced form model is found to underestimate the time-treatment effect component, but the difference is not statistically significant. The point estimate for the treatment-time effect in the structural model is twice the size of the OLS estimate. The estimated treatment-time effects by model are OLS -0.0276 ($-0.0548, -0.0004$) and structural model -0.074 ($-0.1159, -0.0321$)²⁷. The lack of statistical support of attrition bias may be attributed to a relatively small sample size (148 patients in the Topamax study).

Prior Distribution.—The estimated prior distribution of the treatment effect provides insight on a patient’s level of optimism (or pessimism) for participating in the trial. In the dynamic model, the estimated health priors suggest patients believe the experimental drug will lead to an initial 13% increase in GGT levels, followed by a 69% decrease in each subsequent week when compared to GGT levels in the placebo group. Given these estimates, patients in the experimental group *expect* an improvement in their health status after 1.17 weeks of participating in the trial. On the other hand, the estimated treatment effect, found using observed health outcomes, suggests health status improves after 1.86 weeks of participation for patients in the experimental group. This result implies patients are initially

optimistic about the potential health benefits associated with the experimental drug.

Utility Parameters and the Outside Option.—The largest differences between the dynamic and myopic models are found in the estimates for the utility parameters and the outside option. The utility function parameters capture a patient’s sensitivity to changes in risk, income, and side effect experiences. Sensitivity to risk is captured by the coefficient of absolute risk aversion, γ . Patients are less risk averse in the dynamic model ($\gamma = 0.1240$) than in the myopic model ($\gamma = 0.3634$). This result is consistent with economic theory in that the forward looking patient spreads her risk over current and future periods. The myopic patient concentrates risk in the current period. In the same light, the side effect parameters are smaller in the dynamic setting than in the myopic setting. The direct effect of side effects is concentrated in the current period for the myopic patient, but spread across current and future periods for the forward looking patient. As previously argued in the introduction of this paper, side effects experienced in the Topamax study are better tolerated than side effects experienced in the AZT study. Evidence of this is found in the sign of the coefficient of parathesia (side effect 2). Parathesia leads to an increase in utility of 9.4456 utils in the dynamic model and 9.969 utils in the myopic model. I suspect that this parameter is adjusting for spillover learning effects not already captured within the model. Fatigue (side effect 1) leads to a disutility of 4.6825 utils in the dynamic model and 6.892 utils in the myopic model.

The model captures a patient’s response to compensation by estimating a patient’s income elasticity. In the dynamic setting, a patient’s utility is inelastic to changes in income. A 1% increase in income leads to a 0.59% increase in utility. In the myopic setting, the effect of compensation on attrition is strong. A 1% increase in income leads to a 7.9% increase in utility. The large difference in income elasticity between the two models displays the sensitivity of parameter estimates with respect to the discount factor²⁸.

Lastly, I include an additional parameter in the last period utility function, $U_T(\cdot)$, to capture any unobserved benefits (or costs) from completing the trial. In some experiments,

patients are rewarded for completing the trial by receiving a free supply of the actual experimental treatment regardless of group assignment. The parameter estimate for "finishing reward" is 17.2976 utils in the dynamic model and 4.4037 utils in the myopic model. A positive coefficient for this parameter captures the curvature of the survival function with respect to time and suggests the survival function is decreasing at a decreasing rate with time. This result implies that the likelihood of exiting the trial next period decreases the longer a patient remains in the trial.

The outside option is an important component of attrition because it captures the patient's opportunity cost of remaining in the trial. The outside option can be decomposed into two parts: external options and longer term cost of participation. The outside option can include a large range of external options including detoxification clinics, support meetings, or drinking. The outside option also internalizes long term effects of the experimental drug on health even after the patient has stopped taking the drug. The study dataset does not contain information on patient choices after exiting the trial. Therefore, the model takes a simplistic approach in handling external options, but the second component of the outside option is captured by allowing a patient's group assignment belief to change her valuation of the outside option. The outside option constant represents the benefit measured in utils of exiting the trial the first period. In both the dynamic model and myopic model, the value of the outside option constant is initially negative (dynamic model = -0.0905 and myopic model = -1.1549), but the outside option improves quickly in subsequent periods (dynamic model = 0.0501 and myopic model = 1.7417). When the outside option parameters are interacted with the patient's subjective probability of being in the treatment group, $E_{it}(G = \text{Treat})$, then the initial value of the outside option increases (dynamic model = 0.1734 and myopic model = 0.3699). The time-belief component of the outside option differs in the two models. The dynamic model finds a decreasing effect on the outside option (-0.3014) as beliefs and time increase. Patients who believe they are actually receiving treatment are *less* likely to exit the trial. The myopic model finds an increasing effect on the outside

option (1.304) as belief and time increase. Patients who believe they are actually receiving treatment are *more* likely to exit the trial. Again, this divergence in the estimated parameters is tied to the discount factor. Including patient demographics in the outside option and using a heterogenous discount factor may alleviate this problem.

Goodness of Fit in Survival Rates.—The dynamic model’s ability to explain patient behavior is measured by using the estimated structural parameters to simulate attrition decisions, then compare these outcomes with observed decisions in the experiment. First, I use the original GGT values for each patient found at the beginning of the trial as a starting point on health histories. To capture unobservable health shocks, I draw 10 patient specific errors using the estimate of unobserved health heterogeneity and a four random health shocks from the estimate of the regression/learning parameter. Conditional on a patient’s original GGT value at the start of the trial, I use these simulated unobservable errors in conjunction with the patient’s group assignment, estimated treatment effects, and health-side effect covariance matrix to complete a patient’s health and side effect history²⁹. Next, I simulate a set of unobservable utility shocks. The utility shocks consist of a patient specific error and the extreme value error. I simulate 10 patient specific utility errors from a normal distribution with variance equal to the estimated variance of the unobserved heterogeneity in the utility function. The second utility shocks are drawn independently across patients and time from an extreme value distribution with variance equal to $(\hat{\tau}\pi)^2/6$. The simulated health histories and unobservable errors are then used to solve each patient’s value function. If a patient’s value function is > 0 , then the patient remains in the trial; otherwise, the patient exits the trial. Figure 6 exhibits the observed survival rate versus the predicted survival rate for each period. The dynamic model explains 98.4% of the variation in overall survival (attrition) rates, but precision is lost when evaluating predicted survival rates by group. The model captures 92% of the variation in survival rates within the placebo group, but only 56% for the experimental group³⁰. The model underpredicts survival rates in the experimental group and overpredicts survival rates in the placebo group. I suspect that the

lack of precision on survival probabilities may be related to non-compliant behavior. While consuming less pills reduces the effect on health, it also mitigates the disutility from side effects. Ignoring non-compliant behavior may overstate the effect of side effects on attrition; thus, leading to the divergence between predicted and observed survival probabilities.

Learning .—The model captures learning in two parts: (1) Posterior Group Assignment Beliefs and (2) Posterior Expected Treatment Effect. Simulation methods are used to quantify the two parts. Using the estimated structural parameters from the dynamic model, I simulate patient behavior and record the evolution of patient beliefs in each period. Figure 7 demonstrates the evolution of group assignment beliefs and treatment effects for both placebo and experimental patients. The vertical axis captures the mean Bayesian probability of being in the treatment group and the x axis measures treatment periods. Initially, Bayesian beliefs on receiving active treatment decrease in both groups. This result is consistent with having optimistic beliefs. The decrease in subjective treatment probability is greater in the placebo group than the experimental group. While both groups are disappointed with their health outcomes, the experimental group receives a modest return, thus placing a higher probability of being in the experimental group than the placebo patients.

The second component of learning is the evolution of the posterior treatment effect. The posterior treatment effect is $E_{it}(\theta_{gt}) = \sum_{j=1}^h \mu H_j f(d_{it}, \mathcal{H}_{it}, \mathcal{S}_{it})$ where μH is a $1 \times h$ vector of health states and $f(d_{it}, \mathcal{H}_{it}, \mathcal{S}_{it})$ is the discrete posterior distribution of health states, which is dependent upon a patient's health history, side effect history, and prior on group assignment. The posterior mean represents the weighted mean of posterior treatment effect beliefs where a patient's belief on treatment group assignment serve as the weights. In period zero, the posterior mean for both groups is equal to $(.5)(\text{prior TE constant}) + (.5)0 = 0.06635$ where the placebo treatment effect is fixed at zero. I include the estimated treatment effect (TE) for the experimental group as a guide (the "treatment effect" in the placebo group is zero). Ideally, learning would produce posterior treatment effect beliefs close to the estimated treatment belief in the experimental group and posterior beliefs close to zero in

the placebo group. Patients in the experimental group are observed to have posterior means relatively closer to the estimated experimental treatment effect after receiving three signals, but placebo patients have posterior means relatively far from zero after three signals.

There are several points to be made from this result. First, the two dimensional learning model proposed in this paper can capture slow learning. The traditional model of Bayesian learning where each signal is equally weighted by construction forces the first signal to have the largest marginal effect on posterior beliefs. To a naive reader, it may be easy to confuse faster learning with a result that is actually driven by functional form. The two dimensional model can capture slow learning by allowing signals to be weighted using posterior beliefs on group assignment. In this study, the largest marginal improvement in information occurs in the second signal for patients in the experimental group and the last signal for placebo patients. Next, it is observed that patients in the placebo group are relatively faster learners about group assignment, but slower learners about treatment effects. I suspect that this result is driven by the statistically significant pharmacological effect of the experimental drug. The experimental drug generates a strong enough effect, allowing patients to separately identify the treatment effect from random health shocks.

Convergence: Topamax.—As described in the previous section, the additional uncertainty associated with random group assignment causes patient learning to converge slowly toward the estimated treatment effect. To illustrate this point, the structural dynamic model is used to simulate the evolution of patient beliefs in a fictitious clinical trial lasting 30 weeks (instead of 12 weeks as in the original trial). Extending the length of the trial to 30 weeks provides patients with more than three times the number of possible signals generated in the original trial. Figure 8 displays the results of this experiment. Patients in the placebo group do not receive an active drug, therefore the treatment effect in this group is zero. The placebo group does converge towards a treatment effect of zero, but does not converge on group assignment. In week 3 the posterior difference in log health is -0.22, but the posterior treatment effect mean plateaus at week 12 and remains constant

thereafter at -0.044. The posterior probability of being in the placebo group for placebo patients reaches 75% in week 9, but decreases at a decreasing rate through the remainder of the experiment with a minimum of 40%. The divergence in the posterior group belief is related to the movement of the posterior treatment effect for placebo patients towards zero. As the posterior treatment effect approaches zero, it becomes more difficult for the patient to distinguish between signals generated from a non-active drug (placebo) or an experimental drug that has a low return (TE close to zero).

Topamax patients experience convergence in both the treatment effect and group assignment. On the posterior probability of being in the experimental group, Topamax patients reach a low of 31% in week 12, but posterior probability reaches a maximum of 73% in week 30³¹. On the posterior treatment effect, the largest difference between the estimated treatment effect and the Bayesian posterior mean occurs in week 3 when the estimated effect is a 1.2% decrease in GGT and the posterior treatment effect is a 22% decrease in GGT. The posterior mean begins to converge towards the estimated treatment effect at week 12. In the last week of the experiment, the difference between the estimated and posterior treatment effect is reduced to 7 percent. These results reinforce the findings from the previous section that treatment group uncertainty persist, thus leading to a slower rate of learning than would be found when treatment assignment is ignored.

Decomposing Patient Attrition.—Using the same methods described in the AZT section, I decompose the effects of side effects, learning, and the value of external options on patient attrition (see table 8). Two types of side effects are considered in the Topamax study, parathesia and fatigue. Both side effects are found to have a trivial explicit effect on attrition. The mean hazard rate associated with side effects is a 0.1% decrease in attrition. On the other hand, the outside option is found to have a strong negative effect on attrition rates. With the exception of week 9, the outside option decreases attrition by an average of 4%. The effect of learning on attrition is found by differencing the overall hazard rate with the side effect and the outside option hazard rates. Learning is the dominant component of

attrition in the Topamax trial. Initially, learning increases attrition by 18.5%, but the effect of learning on attrition is diminishing. In week 12, the learning hazard rate has decreased to an 11% increase in attrition rates. The learning hazard rates suggest that patients are more likely to exit the trial early when returns are low instead of waiting for more signals to reduce uncertainty.

To separate the explicit and implicit effects of side effects on attrition, I simulate patient behavior under the restriction that placebo side effects are equal to treatment side effects. Hazard rates are then calculated using the previously described method. The mean hazard rate of learning on attrition without side effects is a 13.5% decrease in the sample each evaluation period (3 weeks). Learning without side effects causes more early attrition, but the mean effect over the duration of the trial is virtually the same as the learning hazard with side effects (13.51% w/o side effects and 13.49% with side effects).

VII. Limitations

In this section, I identify limitations of the model in capturing all observable patient behavior. These limitations can be categorized into two groups: model limitations and data limitations. The structural model imposes several assumptions about patient behavior and the pharmacology of the experimental drug. First, patients are assumed to be risk averse with a specific type of utility function. The constant absolute risk aversion utility function provides a convenient analytical specification, but tends to overstate risk aversion. The consequence is an overestimate of attrition rates. Second, estimation of dynamic models are computationally expensive. To alleviate the computational burden, I restricted both the number of side effects used to two and the severity of these side effects to a binary variable. The trade-off of a small state space is loss of statistical efficiency. Third, a linear function (containing a constant and a slope term) is adopted to capture the pharmacological effect of the experimental drug. A more flexible function would capture more of the variation in health outcomes, but would require the model to capture patient learning over these

additional parameter values, which would further increase the computation burden.

With respect to data limitation, I focus on limitations associated with experimental design and survey collection. In experimental design, clinical trials commonly use small samples (typically fewer than 100 observations), which leads to poor estimates of the structural parameters. Further, a sample selection problem arises at the onset of the trial. The subset of patients from the population are generally not representative of the population. Instead, these patients are typically poor or very unhealthy. For this reasons, the estimated treatment effect should be viewed as the improvement in health for this subset of patients and not the general public. In survey design, there exist several variables of interest to an economist, which may not typically be recorded by the health researcher. Two examples are the availability of alternative treatments and the cost of care to treat the disease in the absence of the clinical trial. These variables define the patient's opportunity cost for participating in the trial. A set of variables capturing the choice of entering the trial would alleviate the self selection problem and improve estimates on patient's priors. Information pertaining to how the patient becomes aware of the trial would also be useful to estimate a patient's prior. A patient who is introduced to the experiment by her physician may have higher priors than if she receives the same information from a newspaper ad. Lastly, collecting subjective group probability from each patient throughout the trial could be used to measure the "goodness of fit" in the model.

VIII. Conclusion

In this paper, I provide an empirical examination of learning under ambiguity in the context of clinical trials. A structural model of patient behavior is adapted to include a two dimensional model of Bayesian learning. The model captures patient learning of the experimental drug's treatment effect despite ambiguity of treatment group assignment. The structural model is estimated using data from two clinical trials, a Topamax study and an AZT study. These clinical trials offer a comparison between a drug with strong side

effects used to treat HIV and a drug with few side effects used to treat alcoholism. In both studies, patient learning is found to be slow. In neither study did the patients' posterior on group assignment converge to the actual assignment by the end of the trial. Posterior treatment effects appear to initially diverge from estimated treatment effects in both trials, but slowly converge to the estimated treatment effect in subsequent periods. The fastest rate of learning is found in the treatment group of the Topamax study. After 12 weeks in the study, the posterior treatment effect of patients in the treatment group converges to the estimated treatment effect, but uncertainty about treatment assignment persists.

The prior distribution on treatment effects reveals that patients are initially optimistic about the experimental drug. In the Topamax study, the prior distribution implies that positive returns are expected in week 1.17, but in reality the positive returns occur at 1.87 weeks. In the AZT study, the prior distribution indicates a short run decrease in health followed by a strong long run increase in health. The estimated prior distribution implies that patients view AZT as a cure for HIV. Patients expect a positive returns in week 9 of the study. The estimated treatment effect indicates that positive returns occur immediately, but the magnitude of the treatment effect is much smaller than the prior treatment effect mean implies.

Attrition bias is present in both experiments. In the Topamax study, OLS coefficient estimates under-state the treatment effect. The estimated treatment effect in the structural model is more than two times the magnitude of the OLS estimates. In the AZT study, the least squares estimates on placebo health are underestimated when compared to the structural results. The least squares estimate of placebo health, an 18% improvement in health, is twice as large as the structural model estimate, a 9% improvement. To mitigate some of the problem associated with attrition, I perform a counter-factual experiment that increases the probability of being randomized into a treatment group. Patients in the Topamax study are insensitive to changes in the randomization probability. Increasing the randomization probability by 0.1 leads to a 0.3% decrease in overall attrition. On the other

hand, attrition in the AZT study decreases by 1.1% with a 0.1 increase in the randomization probability. Patients in the placebo group are the most responsive in this regard because attrition rates are decrease by 1.7% for the same change in the randomization probability.

The proposed model can be modified to capture other aspects of learning. Specifically, a patient's prior distribution on the treatment effect can be modelled to vary with patient demographics. The prior distribution may also vary by different types of advertisement used to recruit patients. Patients who are recruited via physician's advice may have much higher priors about the treatment effect than those patients recruited via newspaper ads. In addition, the subject of learning under ambiguity is not limited to clinical trials, but may be applied in many other topics of interest for the economist. In the area of education, the study of student attrition from college courses could be viewed as learning ones ability (treatment effect) to perform well in a class given the ambiguity of a teacher's quality (group assignment).

IX. Bibliography

Ackerberg, D., "Advertising, Learning, and Consumer Choice in Experience Good Markets: A Structural Examination," *International Economic Review*, 44, 1007-1040 (2003).

Berkovec, James and Steven Stern (1991). "Job Exit Behavior of Older Men." *Econometrica*. 59 (1): 189-210.

CDC Staff, "The Cost of AIDS and Other HIV Infections: Review of the Estimates." mimeo (1987).

Crawford, G. and Shum, M. (2005). "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, forthcoming.

Chan T. and Hamilton B. "Learning, Private Information and the Economic Evaluation of Randomized Experiments," (forthcoming, *Journal of Political Economy*, 2006).

DeGroot, M.H. . *Optimal statistical decisions*. McGraw-Hill, 1970

Erdem, T. and Keane, M. (1996). "Decision-Making Under Uncertainty: Capturing Dy-

namic Brand Choice Processes in Turbulent Consumer Goods Markets,” *Marketing Science*, 15, 1-20.

Hausman, J. and Wise, D. (1979) "Attrition Bias in Experimental and Panel Data: The Gary Income Maintenance Experiment," *Econometrica*, Vol. 47, No. 2. (Mar., 1979), pp. 455-474.

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.

Hrobjartsson, A. Gotzsche, P. C. (2001). "Is the Placebo Powerless? An Analysis of Clinical Trials Comparing Placebo with No Treatment," *New England Journal of Medicine* VOL 344; PART 21, pages 1594-1602

Kircsh, I., Moore, T., and Nicholas, S., "The Emperor's New Drugs," *Prevention & Treatment*, Volume 5, Article 23, posted July 15, 2002

Malar, C. "The Estimation of Simultaneous Probability Models," *Econometrica* Vol 45, No. 7 (Oct. 1977)

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

Malani, A., "Patient Enrollment in Medical Trials: Selection Bias in a Randomized Experiment," mimeo (October 2005)

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

A.H.P. Paes, A. Bakker and C.J. Soe-Agnie, "Measurement of patient compliance," *Pharm World Sci* 20(2): 73-77 (1998).

Stern, S., "Two Dynamic Discrete Choice Estimation Problems and Simulation Method Solutions," *Review of Economics and Statistics*, November 1994, 76(4): 695-702.

Volberding et. al., (1990). "Zidovudine in Asymptomatic Human Immunodeficiency Virus Infection." *New England Journal of Medicine*, Vol. 322, No. 14 (April). p. 941 - 949.

Figure 1

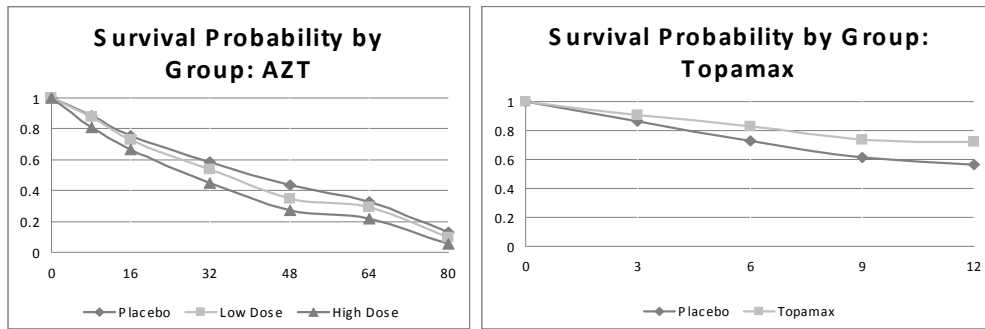


Figure 2




Figure 3: Goodness of Fit - AZT




Figure 4

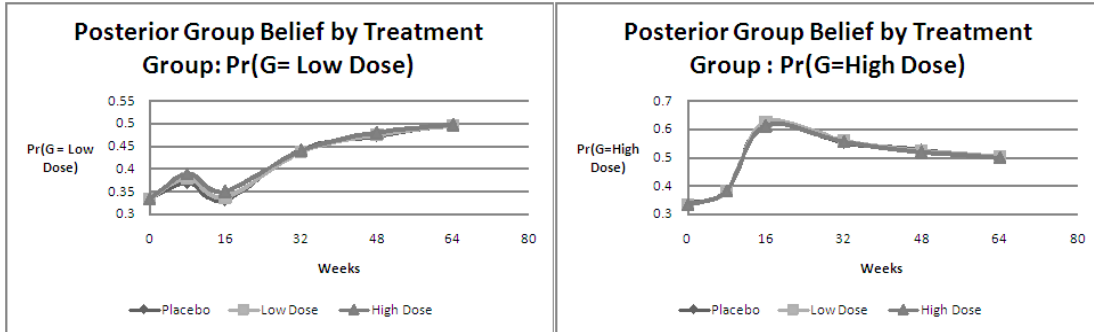



Figure 5



Posterior Treatment Effect Mean by Group Posterior TE Mean by Group: Simulation

Figure 6: Goodness of Fit - Topamax




Figure 7

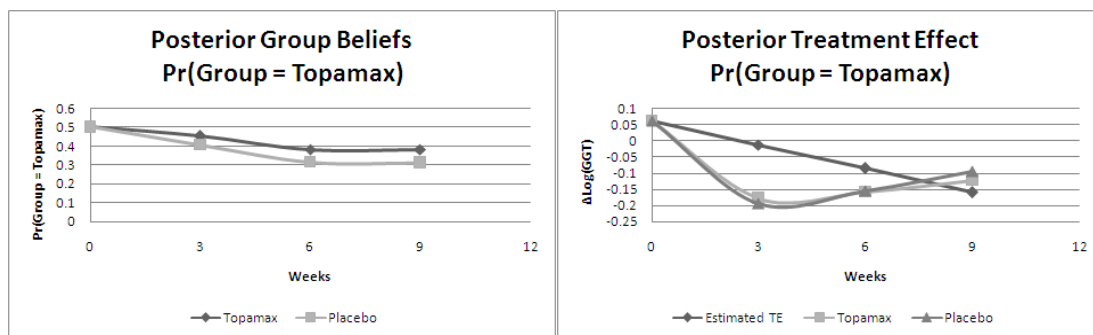
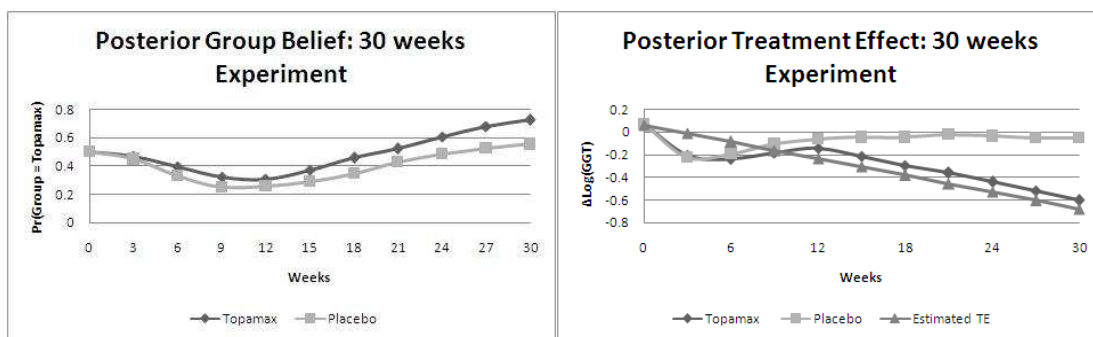


Figure 8



Tables

Table 1: ACTG 019 Descriptive Statistics

Variables	Placebo		Low Dose		High Dose	
	Mean	STD	Mean	STD	Mean	STD
Black	.07	.26	.10	.30	.13	.34
Female	.09	.29	.07	.26	.08	.27
Hispanic	.10	.30	.09	.29	.10	.30
Age	34.74	7.77	34.93	8.23	34.72	8.08
Weight	166.57	24.91	164.62	26.86	167.34	30.09
Hemoglobin	14.67	1.1	14.58	1.12	14.69	1.18
Platelets	225.46	59.94	219.53	53.73	223.20	55.76
Granulocytes	2947.49	1229.32	2820.63	1042.17	2928.29	1059.3
Homosexual	0.73	0.44	0.72	0.45	0.72	0.45
Bisexual	0.14	0.35	0.21	0.41	0.18	0.38
IV Drug Use	0.11	0.31	0.06	0.24	0.07	0.26
CD4	373.48	163.66	383.61	172.8	379.88	163.92
Blood	0.0167	0.3389	0.0183	0.2952	0.068	0.5646
Liver	0.0365	0.5151	0.0464	0.5787	0.0257	0.5064
AIDS	0.1023	0.3031	0.142	0.3491	0.1148	0.3189
Duration (<i>weeks</i>)	37.63	27.31	34.10	26.58	28.70	25.42
N	428		453		456	

Table 2: Patient Characteristics

Variable	Placebo	Topamax	ALL
AGE	42.07	41.51	41.79
BMI	26.57	25.67	26.12
Female	0.27	0.31	0.29
Yrs. School	13.55	13.51	13.53
Income last 30 Days	3279.65	3391.43	3335.54
Base line GGT	65.31	81.8	73.55
End GGT	52.52	57.83	55.18
AVG Survival (weeks)	5.92	6.72	6.32
N	75	75	150

Table 3: Structural Parameters: ACTG 019 Study

Prior Distribution of Treatment Effects (TE)	Dynamic	Static	Utility Parameters	Dynamic	Static
Prior TE Dose	-0.1767 ^a (0.0165)	-3.5022 ^a (0.0205)	Utility: Unobserved Heterogeneity	0.003 ^a (4.09E-06)	0.0030 ^a (3.7E-05)
Prior TE Time	2.1596 ^a (0.0162)	0.0111 (0.0127)	Coef. of Absolute Risk Aversion	1.014 ^a (0.0123)	1.2511 ^a (0.0151)
Prior Variance (TE dose)	0.0048 ^a (8.2E-05)	0.2893 ^a (0.0035)	Side Effect 1 Sensitivity	166.1292 ^a (0.0122)	169.6744 ^a (0.0121)
Prior Variance (TE Time)	0.2564 ^a (0.0054)	0.7454 ^a (0.0090)	Side Effect 2 Sensitivity	64.2802 ^a (0.0122)	82.8041 ^a (0.0135)
			Finishing reward	3.2862 ^a (0.0122)	7.3841 ^a (0.0226)
			Tau	57.3458 ^a (0.711)	57.3458 (fixed)
Health Outcomes Parameters	Dynamic	Static	Outside Option Parameters	Dynamic	Static
Health Variance	0.0453 ^a (0.0009)	0.0449 ^a (0.0011)	Outside Option Constant	-1.4567 ^a (0.0163)	-87.8342 ^a (0.0128)
Health: Unobserved Heterogeneity	0.1114 ^a (0.0036)	0.1118 ^a (0.0019)	Outside Option Time	0.5089 ^a (0.0158)	0.0291 ^a (0.0134)
Blood Side Effects: Placebo	0.1435 ^a (0.0015)	0.1450 ^a (0.0063)	Outside Option E _{it} (Dose =1)	-0.3755 ^a (0.0139)	-2.8339 ^a (0.0122)
Liver Side Effects: Placebo	0.1048 ^a (0.0011)	0.1051 ^a (0.0035)	Outside Option E _{it} (Dose =1)*Time	-0.1265 ^a (0.0159)	-11.3141 ^a (0.0133)
Blood Side Effects (Dose =1)	0.1841 ^a (0.0018)	0.1805 ^a (0.0049)	Outside Option E _{it} (Dose =3)	0.3802 ^a (0.0124)	-10.2138 ^a (0.0121)
Liver Side Effects (Dose =1)	0.0971 ^a (0.001)	0.0981 ^a (0.0019)	Outside Option E _{it} (Dose =3)*Time	-1.3831 ^a (0.0139)	-3.5557 ^a (0.0183)
Blood Side Effects (Dose =3)	0.256 ^a (0.0025)	0.2547 ^a (0.0048)	AIDS Dummy	0.286 ^a (0.0177)	-0.9332 ^a (0.0121)
Liver Side Effects (Dose =3)	0.1171 ^a (0.0012)	0.1157 ^a (0.0054)	Year 88 Dummy	0.31 ^a (0.0143)	-66.5812 ^a (0.0151)
Variance Side Effect 1	0.154 (0.1044)	0.1530 ^a (0.0200)	Year 89 Dummy	1.86 ^a (0.0268)	1.2397 ^a (0.0255)
Variance Side Effect 2	0.0947 (0.2514)	0.0948 ^a (0.0174)	Discount Factor	B = .98 (fixed)	B = 0 (fixed)
Cov(Health, Side Effect 1)	-0.0087 (0.0058)	-0.0091 (0.0126)	Log Likelihood Value	-6785.6339	-6812.8377
Cov(Health, Side Effect 2)	-0.0029 (0.0049)	-0.0030 (0.0113)	F	54.4076 ^a	
Cov(Side Effect 1, Side Effect 2)	0.0554 ^a (0.105)	0.0553 ^a (0.0126)	N	1337	

standard errors are in parenthesis, each dose level is equal to 500mg of AZT

F test: Likelihood ratio test (dynamic vs static)

a= Significant at 99% b= Significant at 95% c= Significant at 90%

Table 4: AZT Health Parameters

VARIABLES	OLS	Structural
Constant	0.1829 ^a (0.0149)	0.0895 ^a (0.0133)
Time	-0.0092 ^a (0.0043)	-0.0121 ^a (0.0038)
Dose	0.0149 (0.0082)	0.0189 ^a (0.0063)
Time*1(Dose=500mg)	0.021 ^a (0.005)	-
Time*1(Dose=1500mg)	0.011 (0.0072)	-
Time*1(AZT)	-	0.0045 (0.0041)
AIDS	-0.6617 ^a (0.0228)	-

^a = 1%, ^b = 5% level of significance

Table 5: Hazard Rate - AZT

week	8	16	32	48	64	80
Side Effects	-0.1961	-0.1376	-0.0377	-0.0337	-0.0104	0.0599
Outside Option	-0.0352	-0.0842	-0.0285	-0.0294	-0.0122	-0.0247
Learning	0.0151	0.0544	0.0174	0.0127	-0.0434	-0.1074
Learning without Side Effects	-0.0643	0.0088	0.0105	-0.0119	-0.0686	-0.1364

Table 6: Structural Parameters: Topamax Study

Prior Distribution of Treatment Effects (TE)	Dynamic	Static	Utility Parameters	Dynamic	Static
	Prior TE constant	0.1327 ^a (0.0442)		-1.5686 ^a (0.1035)	Utility: Unobserved Heterogeneity
Prior TE Time	-0.6908 ^a (0.0482)	6.0883 ^a (0.0441)	Coef. of Absolute Risk Aversion	0.1240 ^a (0.0055)	0.3634 ^a (0.0996)
Prior Variance (TE constant)	0.2814 ^a (0.0125)	0.5287 ^a (0.0233)	Income Elasticity	0.5933 ^a (0.0477)	7.9407 ^a (0.5737)
Prior Variance (TE Time)	0.0614 ^a (0.0027)	0.4928 ^a (0.0297)	Side Effect 1 Sensitivity	4.6825 ^a (0.0442)	6.892 ^a (0.0497)
Health Outcomes Parameters	Dynamic	Static	Side Effect 2 Sensitivity	-9.4456 ^a (0.0441)	-9.969 ^a (0.0597)
Health Variance	0.0388 ^a (0.0018)	0.0389 ^a (0.003)	Finishing reward	17.2976 ^a (0.0441)	4.4037 ^a (0.1238)
Health: Unobserved Heterogeneity	0.0930 ^a (0.0044)	0.0935 ^a (0.0125)	Tau	18.4968 ^a (0.9137)	18.4968 (fixed)
Placebo Side Effect 1 Mean	0.1016 ^a (0.0040)	0.101 ^a (0.018)	Outside Option Parameters	Dynamic	Static
Placebo Side Effect 2 Mean	0.0427 ^a (0.0018)	0.0386 ^a (0.015)	Outside Option Constant	-0.0905 (0.0596)	-1.1549 ^a (0.462)
Treatment Side Effect 1 Mean	0.1487 ^a (0.0059)	0.1505 ^a (0.0098)	Outside Option Time	0.0501 (0.1082)	1.7417 ^a (0.3942)
Treatment Side Effect 2 Mean	0.1489 ^a (0.0056)	0.149 ^a (0.011)	Outside Option E _{it} (G = Treat)	0.1734 ^a (0.0641)	0.3699 ^a (0.0703)
Variance Side Effect 1	0.1103 ^a (0.0204)	0.1104 ^a (0.0205)	Outside Option E _{it} (G = Treat)*Time	-0.3014 ^a (0.0553)	1.304 ^a (0.4023)
Variance Side Effect 2	0.0866 ^a (0.0202)	0.0863 ^a (0.021)	Discount Factor	0.98 (fixed)	0 (fixed)
Cov(Health, Side Effect 1)	0.0034 (0.0056)	0.0033 (0.0073)	Log Likelihood Value	-514.1299	-516.1637
Cov(Health, Side Effect 2)	-0.0055 (0.0051)	-0.0057 (0.0067)	F	4.0676 ^b	
Cov(Side Effect 1, Side Effect 2)	0.0527 ^a (0.0193)	0.0528 ^a (0.0199)	N	148	

standard errors are in parenthesis

F test: Likelihood ratio test (dynamic vs static)

a= Significant at 99% b= Significant at 95% c= Significant at 90%

Table 7: Topamax Health Parameters

Variables	OLS	Structural
Constant	-0.1804 ^a (0.0508)	-0.1837 ^a (0.0379)
Time	-0.0035 (0.0105)	-0.0232 (0.0218)
Treatment Constant	0.0656 (0.0692)	0.0621 (0.075)
Treatment Time	-0.0276 ^b (0.0139)	-0.074 ^a (0.0214)

^a = 1%, ^b = 5% level of significance

Table 8: Hazard Rate - Topamax

week	3	6	9	12
Side Effects	-0.0011	0.0027	0.0089	-0.0047
Outside Option	0.0506	0.0402	0.0008	0.0623
Learning	-0.1855	-0.1292	-0.114	-0.1106
Learning without Side Effects	-0.2093	-0.1222	-0.1006	-0.1086

Appendix: Discretization Method

The learning model present in this paper requires that patient's beliefs on health outcomes follow a normal mixture model³². The model is a linear

$$\sum_{j=1}^{G=3} \delta_{igt} \int_{\lambda \in \Lambda} \phi(\mathcal{H}_{it}, \mathcal{S}_{it} | \lambda, \Omega, b_g, l_g) \phi(\lambda | \mu_{igt}, \Sigma, \delta_{igt} = 1) d\lambda \text{ s.t. } \sum_{j=1}^{G=3} \delta_{igt} = 1$$

combination of G normal distributions where the linear weights, δ_{igt} , must sum to one. Unfortunately, a closed form solution to the posterior distribution of the unknown distribution parameters in a normal mixture model does not exist. Still, the posterior distribution may be approximated by discretizing the probability states space conditional on a patient's prior beliefs. A patient's prior on treatment group assignment is $\delta_{ig0} = 1/G$, which is provided to the patient by the clinical trial investigators. The patient also forms priors on the experimental effect. These priors are assumed to be normally distributed.

$$(A2) \quad \mu_{ig0} = \begin{pmatrix} \lambda_{1_{ig0}} \\ \lambda_{2_{ig0}} \end{pmatrix} \sim N \left(\underline{\mu}_{g0} = \begin{bmatrix} \underline{\lambda}_{1_0} \\ \underline{\lambda}_{2_{g0}} \end{bmatrix}, \Sigma = \begin{bmatrix} \sigma_{\lambda_1}^2 & 0 \\ 0 & \sigma_{\lambda_2}^2 \end{bmatrix} \right)$$

Initially, the priors on the dose effect, $\lambda_{1_{ig0}}$, and the treatment time trend, $\lambda_{2_{ig0}}$, are assumed to be independent. Therefore, A2 can be decomposed into two univariate normal distributions. I then select a set of discrete values for λ_1 and λ_2 , $\Lambda = [\lambda_{-2}, \dots, \lambda_0, \dots, \lambda_2]$, where $\lambda_i = \underline{\lambda}_{1_0} + (i) \sigma_{\lambda_1}$. Each discrete value of the experimental effect is separated by one standard deviation. The probability mass function of Λ is defined as

$$\text{Pr}(\lambda_i | \underline{\lambda}_{1_0}, \sigma_{\lambda_1}) = \frac{\Phi\left([\lambda_i + \frac{\sigma_{\lambda_1}}{2} - \underline{\lambda}_{1_0}] / \sigma_{\lambda_1}\right) - \Phi\left([\lambda_i - \frac{\sigma_{\lambda_1}}{2} - \underline{\lambda}_{1_0}] / \sigma_{\lambda_1}\right)}{\sum_{\Lambda} \Phi\left([\lambda_j + \frac{\sigma_{\lambda_1}}{2} - \underline{\lambda}_{1_0}] / \sigma_{\lambda_1}\right) - \Phi\left([\lambda_j - \frac{\sigma_{\lambda_1}}{2} - \underline{\lambda}_{1_0}] / \sigma_{\lambda_1}\right)}$$

where $\Phi(\cdot)$ is the normal cumulative density function. An analogous equation defines the probability mass function for λ_2 . These probability mass functions discretize a patient's

prior belief on the experimental effect and their treatment group assignment.

Given the set of possible experimental effects, equation A1 can be discretize over health outcomes state space. A patient's health is defined as a vector containing CD4 counts, blood related side effects, and liver related side effects. Initially, the side effect state space takes on a simple definition: 1 if a side effect is experienced and 0 otherwise³³. Both blood and liver side effects are defined in this manner. CD4 counts are discretize into 10 equispace values over the sample range of CD4 counts. Therefore, the health outcomes state space for a given value of λ , side effect means (b_g, l_g) , and treatment group assignment is $S = (10) (2) (2) = 40$. The rectangular area defined by $\Phi(S|\lambda, b_g, l_g)$ is found utilizing the multivariate normal cumulative density function in the MATLAB programming language. The joint distribution of health outcomes and beliefs is then found by evaluating the 40 health outcome states for each experimental effect and treatment group state. The total probability space is then $(\# \text{ of health states}) \times (\# \text{ of experimental effect states}) \times (\# \text{ of Groups}) = 3000$, which must be carried in memory for each patient.