Experimenting within Experiments: Patient Learning and Experimentaiton in Clinical Trials

Jose Manuel Fernandez*

The University of Louisville

Abstract

This paper proposes a structural model to study if patient non-compliance is the result of strategic experimentation or side effect avoidance. The model is estimated using data from a clinical trial studying the effects of Topamax on alcoholism. I find empirical evidence supporting the hypothesis of "forward looking" patients who use non-compliant behavior to increase their learning rates. On average, dose consumption is 6.5% higher in the Topamax group and 9.2% higher in the placebo group when comparing "forwarding looking" versus myopic patients. The dynamic model predicts that patients in the experimental group who experiment with dose converge faster to their actual group assignment than if they were to maximize static utility only. Lastly, the estimated treatment effect found in the structural model is nearly twice the magnitude of the OLS estimate.

Introduction

This paper proposes a structural model of patient behavior to study if patient non-compliance is the result of strategic experimentation or side effect avoidance. In a clinical

*Acknowledgements: I would like to thank John Pepper, Amalia Miller, Steven Stern, Michael Conlin, Leora Friedberg, Anup Malani, and Marc Santugini for their helpful comments. I gratefully acknowlege the financial support from the Bankard Fund for Political Economy at the University of Virginia. All remaing errors are my own. Email: jose.fernandez@louisville.edu

trial, a patient is uncertain along two dimensions: her treatment group assignment and the experimental drug's treatment effect. In an effort to resolve these uncertainties, patients may vary dose consumption (display noncompliant behavior) to accelerate the learning process. For example, a patient may forgo all her pills in the first period, then attribute any changes in her health over this period as part of the natural progression of her disease. In the following period, she may then choose to act compliant and again observe changes in her health. If these changes do not deviate sufficiently from her previous experience, then she can infer either the treatment effect is trivial or she has been randomized into the placebo group. On the other hand, if the change in health is large, then she updates her beliefs by increasing her posterior belief of being randomized into the treatment group.

Patient learning is modelled within a dynamic discrete choice (DDC) framework. A patient observes changes in her health (both side effects and health status), then uses these observed changes to update her beliefs on the treatment effect and treatment group assignment simultaneously using Bayes' Law. Next, the patient chooses whether to remain in the trial and if so how much to consume. The key insight of the model is the information gathered from side-effect experiences. Although side effects do present a direct source of disutility, patients may use side effect experiences as an instrument to identify their group assignment. Consider a randomized placebo controlled trial where the experimental treatment is known to present side effects in the form of changing one's hair green. If a patient were to experience such a side effect, though unpleasant, the patient could infer that she has been assigned to the treatment group and would be able to update her beliefs on the treatment effect without the additional uncertainty of possibly consuming a placebo. Further, if changes in health and side-effects are increasing with dose, then learning may be accelerated by consuming a higher dose than the optimal myopic level. Mirman, Grossman, and Kihlstrom (1977) describe a model of strategic experimentation with pharmaceutical drugs where "forward-looking" agents are found to accelerate learning by consuming more of the drug than the optimal myopic dosage.

The structural model is estimated using data from a clinical trial, which tests the effects of Topamax, an anti-seizure medication, on alcohol consumption. The estimated model finds empirical evidence supporting the hypothesis of "forward looking" patients who use non-compliant behavior to increase their learning rates. On average, dose consumption is 6.5% higher in the Topamax group and 9.2% higher in the placebo when comparing forwarding looking versus myopic patients. This result is consistent with theoretical models of learning with strategic experimentation. Further, the dynamic model predicts that patients in the experimental group who experiment with dose converge faster to their actual group assignment than if they were to maximize static utility only.

Literature Review

Self selection models are widely used to control for attrition and non-compliance [Hausman (1979), and Heckman, Hohmann, and Smith (2000)]. Self -selection assumes that a patient makes a sequential choice: (1) "should I remain in the trial, and (2) if I remain, then how much should I consume?" A first stage regression on survival probabilities is estimated. These predicted survival probabilities are then used to correct for attrition bias in second stage regressions on health outcome measurements. Sample selection models highlight the effect of attrition on drug efficacy estimates; however, these methods can neither account for learning, nor provide a framework in which to mitigate attrition bias by providing alternative experimental specifications.[1]

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. Heterogeneous attrition rates are

[1] Malani (2005) further extends the discussion of selection by demonstrating self-selection in clinical trials for treatment of ulcers is positively correlated with the inital treatment probability. The author states that as the probability of treatment increases the enrollment of less optimistic (healthier) patients increases, thus leading to lower estimated values of the treatment effect (smaller bias).

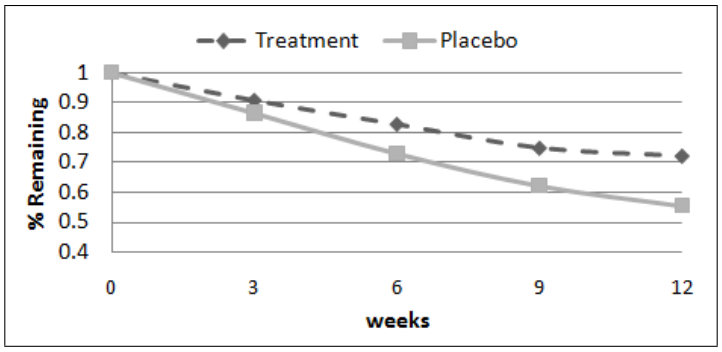

Figure 1: Survival Function by Treatment Group

also a feature of the Topamax clinical trial data, where attrition rates are consistently higher in the placebo group than in the treatment group (see Figure 1). Moreover, the difference in attrition rates between treatment groups increases between weeks 3 and 12. This result motivates the study of attrition as a utility maximizing decision.

The model presented in this paper is closely related to the growing empirical literature on the demand for experience goods using Bayesian learning within a DDC framework [Erdem and Keane (1996), Ackleberg (2006), Crawford and Shum (2006), Coscelli and Shum (2004)]. In these studies, an agent learns a product's quality through repeated exposure (purchase and/or advertising). These models utilize observational consumer level data to infer product quality learning from a consumer's purchase history. Chan and Hamilton (2006) investigates 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. Side effects enter as unobservable group characteristics leading to disutility, but do not affect patient learning. The authors find evidence that when accounting for unobserved side effects, the treatment that maximizes utility is not always the most medically effective treatment.

Unlike these previous models, which primary study switching behavior (attrition), this study focuses on quantity experimentation. The role of patient noncompliance has been widely studied in the medical field [see Efron and Feldman (1991), Goetghebeur and Lapp

(1997), Fischer and Goetghebeur (2004)], but the issue has received much less attention from the economics literature. Ellickson, Stern, and Trajenberg (1999) are the first to develop a model of prescription drug adherence. The study builds on physician's motives when prescribing drugs and patient adherence associated with these drugs. Although, the authors provide an estimation strategy to measure patient welfare, the study lacks appropriate data to estimate the model of behavior. Lamiraud and Geoffard (2007) does estimate a model of patient welfare concentrating only on drug adherence. The authors view noncompliance as the result of a utility maximizing decision in a static discrete choice model. Side effects are a source of disutility and may lead to non-compliance. The authors estimate their model using data from a randomized clinical trial that compares the efficacy of 2 tritherapy for treatment of HIV and find that for two drugs, which demonstrate the same level of clinical efficacy and toxicity, a higher adherence level is associated with higher patient welfare. These reduced form results highlight the importance of non-compliance when evaluating the welfare effects of pharmaceutical drugs, but these techniques do not explicitly control for learning or experimentation.

This paper extends the current literature by modelling and estimating patient noncompliant behavior in a dynamic discrete choice model with Bayesian learning. In the previous models, agents are only uncertain about product quality, we consider a model where a patient's uncertainty extends along two dimensions: drug quality and treatment group assignment.[2] Therefore, patients cannot immediately attribute observed changes in health to a treatment effect. Following Fernandez (2009), a patient's belief on treatment group assignment is allowed to affect her attrition decision and her level of compliance. 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% may be less likely to exit the trial than if her

[2]The proposed model can be used in other fields of economics. In the economics of education, college students face uncertainty in teacher quality and in their own ability when choosing to drop a class. In labor economics, firms face uncertainty in demand and worker productivity when making worker attrition decisions.

chances are only 10%. 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 appear to be small.

Second, the model incorporates observable side effect experiences as informative signals of treatment group assignment and a potential source of disutility caused by participating in the trial. Further, this source of disutility may intensify with an increase in dose consumption. Some clinical trials do include an active ingredient in the placebo medication to mimic the side effect present in the experimental medication thereby reducing the likelihood of patients learning their treatment assignment.[3] Lastly, financial compensation for participating in the trial may be used to estimate welfare effects associated with side effect experiences, uncertainty, and demand for clinical trial participation. The additional structure provides a framework to analyze patient behavior under different experimental designs, such as different initial treatment probabilities or compensation amounts.

Data

Data from a randomized double blinded placebo controlled study testing the use of Topamax to treat alcohol dependence is used to evaluate patient learning and experimentation. 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.

[3] A clinical trial studying the effects of acyclovir for the prevention of recurrent herbes simplex virus eye disease use a lactose filler that served to mimic the gastrointestinal side effects of the acyclovir. (New England Journal of Medicine, July 1998)

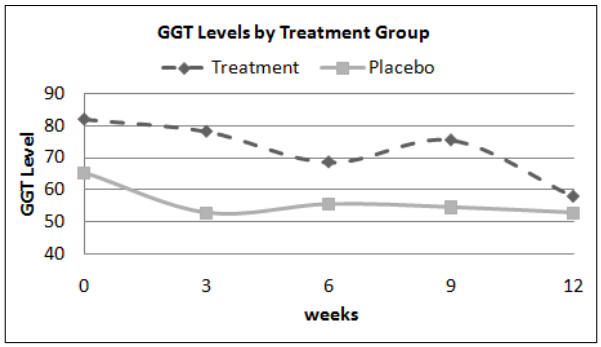

Figure 2: Health Status by Treatment Group

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 1 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%). The racial distribution among patients is white (60%), Mexican (29%) and other (11%).[4]

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).[5]

The study documents side effects by recording the date of onset, duration, severity, and action taken. The likelihood of experiencing at least one side effect is about the same in

[4]The distribution of race is more a reflection of the study location, Texas, than of the population of alcoholics.

[5]Though unsual, it is possible for randomization to yield unequal baseline measures in GGT.

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%).[6] 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-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, then so should the difference in attrition rates between treatment groups. 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. An alternative 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.

[6] Parasthesia is a numbing sensation felt along the extreminities. Sombelance can be described as extreme fatigue.

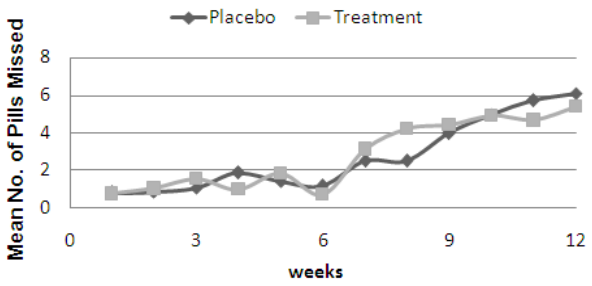

Figure 3: Noncompliance by Treatment Group

An important characteristics of the data is the presence of patient non-compliance. Non-compliance is the event when a patient does not consume her prescribed amount of the study medication. Non-compliance is observed by the econometrician through pill consumption. At each visitation, a patient must return any unused tablets prior to receiving additional medication. Missed doses are calculated as [# of pills prescribed] − [# of pills consumed] = Δd. Missed doses serve as the best measure of non-compliance, but this measure is not without flaws. A patient may choose to lie and always report she has consumed all her prescribed pills when she has not. Alternatively, a patient may be non-compliant because she has lost some tablets, but the measure of missed dose would observe her as being compliant.[7] For simplicity, it is assumed that patients are honest and never lose their medication. Figure 3 illustrates non-compliant behavior by graphing the average level of Δd by group. Non-compliant behavior increases for both groups at the same average rate (7% - 8%) throughout the study, but there is greater variation in the level of non-compliance in the treatment group.

I propose two reasons to observe non-compliant behavior. First, dose is chosen as the result of maximizing current utility. A patient views improved health as the benefit of increasing her dose, but faces the cost of an increased risk of side-effects. If side effects are significantly different between groups, then one would expect higher rates of non-compliance

[7] Paes et al (1998) studied the compliance patterns of 91 diabetic patients using oral antidiabetics. Using Medication Event Monitoring System data as a standard, the results show that pill count and refill data overestimate the compliance of this group of patients.

in the experimental group than in the control group.

A second reason is the patient varies dose to learn her group assignment. In an extreme example, a patient may choose to consume none of her prescribed pills in the first period and consume all her pills in the following period. If the patient observes a significant improvement in health between the two periods, then she may infer she has been randomized into the experimental group, else her beliefs will adjust towards the placebo group. A different channel of learning is through side effects. Again, if a patient does not experience higher than normal side effects when taking all her pills, then she may infer her group assignment is the placebo group because placebo side effects are independent of dose. Patients can improve their rate of learning by becoming "active experimenters" through dose variation. While a non-learner is only concern with maximizing current utility. An active experimenter is a dynamic optimizer who recognizes that dose choice today will affect both health tomorrow and their rate of learning. The study focuses on these two sources of patient non-compliance to identify if clinical trial patients are active experimenters.[8]

The Topamax study used a progressive dose schedule as found in the table below. The amount of active Topamax given to the patient each week would increase at a rate of 25mg per day for the first four weeks and 50mg per day thereafter, but the number of pills given to the patients would not increase in the same monotonic fashion.[9] Failing to consume one tablet in the first period may be very different than not consuming a tablet in the last period, as the level of Topamax milligrams per tablet is different in each period.[10] The outcome variable of missed dose is made more uniform by modifying Δd to capture the number of milligrams not consumed each period, ΔdMG_{it}.[11]

[8] Non-compliant behavior caused by a patient forgetting to take her pills is viewed as a random event. On average, the percentage of non-compliant patients who forget to take their pills is expected to be the same in both groups.

[9] For anti-seizure medication, it is common to use a progessive treatment plan.

[10] A copy of the consent form used for this trial is available in the appendix. The consent form does inform the patient that she will receive an increasing amount of Topamax over the course of the trial, which is not to exceed 300mg of per day.

[11] The amount of active ingredient is assumed to be distributed uniformly among the pills within a given week. Therefore, the amount of Topamax per pill is the ratio of (mg per day)/(pills per day).

Dosing Schedule

Visit	1	2	3	4	5	6	7	8	9	10	11	12
mg per day	25	50	75	100	150	200	250	300	300	300	300	300
pills	14	21	21	28	21	14	28	42	42	42	42	42
pills per day	2	3	3	4	3	2	4	6	6	6	6	6

The following regression equation is used to identify the most common characteristics of non-compliant patients $\Delta dMG_{it} = X\beta + e_{it}$ for all $\Delta d_{it} > 0$ where β is a set of parameters to be estimated, X is a matrix of patient characteristics including time, gender, age, BMI, treatment group, occupation skill level, race, religion, and income. The error term, e_{it}, is an unobserved normally distributed random variable with mean zero and finite variance.

About 67% of the patients in the sample are found to be compliant $(\Delta d_{it} = 0)$. Due to the large mass of observations at zero, a Tobit model is appropriate to estimate the parameters. The maximum likelihood estimates of the parameters β are found in Table 3 . A statistically significant difference in patient compliance is detected between treatment groups. Non-compliance increases throughout the trial for both groups, but on average the placebo group *is more* non-compliant than the experimental group. Non-compliant behavior decreases with age and years of schooling. Patients who identify themselves as Christians are more compliant than non-Christians. Individuals in skilled professions are more non-compliant than unskilled patients. Finally, no statistically significant difference in compliance exists with respect to gender, BMI, race, and income. These results raises the questions of why are patient in the experimental group, who presumably are more likely to experience side effects, are more compliant than their peers in the placebo group?

These results serve as motivation to address the issue of non-compliant behavior using a structural approach. Unlike the "black box" approach of linear regression, a structural model has the advantage of explicitly stating how patients choose their level of dose and when to leave the trial. The structural model can provide insight on patient learning, dose behavior, attrition behavior, and the effect of attrition on the estimated treatment effect.

Structural Model of Attrition and Learning with Experimentation

Consider a clinical trial in which there are G treatment groups. Patients are randomly assigned to a treatment group with a fixed probability, $1/G$. In each treatment period t, patient i observes her health H_{igt}, and two types of side effects, $S1_{igt}$ (parathesia) and $S2_{igt}$ (fatigue). After observing these signals, the patient chooses a dose level, d_t, that maximizes her expected utility and then decides to remain in or exit the trial. Health and side effects are assumed to be distributed joint normal

$$(1) \qquad \begin{bmatrix} H_{igt} \\ S1_{igt} \\ S2_{igt} \end{bmatrix} \rightsquigarrow N\left(\begin{bmatrix} H_{i0}+\theta_{it}+\theta_{gt}(d) \\ s1_{igt}(d) \\ s2_{igt}(d) \end{bmatrix}, \Omega = \begin{bmatrix} \sigma_{\varepsilon}^2 & \rho_1\sigma_{u1}\sigma_{\varepsilon} & \rho_2\sigma_{u2}\sigma_{\varepsilon} \\ \rho_1\sigma_{u1}\sigma_{\varepsilon} & \sigma_{u1}^2 & \rho_3\sigma_{u1}\sigma_{u2} \\ \rho_2\sigma_{u2}\sigma_{\varepsilon} & \rho_3\sigma_{u1}\sigma_{u2} & \sigma_{u2}^2 \end{bmatrix} \right)$$

where d represents the sum of previous and current dose decisions, $d = \sum_t d_t$. Patients are assumed to know certain components of their health equation. In particular, a patient knows her initial health stock (log GGT), 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, $\theta_{gt}(d) = [\theta_3 d_g + \theta_4] 1_{g=treatment}$. Patients are assumed to know the conditional mean for both side effects and each side effect mean has the following functional form.

$$s_{igt}(d) = \begin{array}{ll} \frac{\exp(\varphi_1+\varphi_1 d)}{1+\exp(\varphi_1+\varphi_1 d)} & g = \text{treatment} \\ \frac{\exp(\varphi_1)}{1+\exp(\varphi_1)} & g = \text{placebo} \end{array}$$

Typically, pharmaceutical drugs go through three phases of testing before receiving approval from the Federal Drug Administration.[12] The second phase focuses on safety where side effects are analyzed in a larger sample setting. This information is then provided to participant in Phase 3 clinical trials. Therefore, it is reasonable to assume that patients are

[12]The first phase conducts a small sample test on efficacy. The second phase focuses on safety where side effects are analyzed in a larger sample size. The third phase increases the sample size dramatically by conducting the experiment simultaneously over several medical centers.

informed about the potential side effects. After observing prices (compensation to remain in the trial and the cost of alternative treatments), health status, and side effects in a given period, the patient decides whether to remain in the trial or consume the outside option.

Patient's Preferences

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 over a finite horizon, conditional on their information set, I, at time t:

$$\max_{A_\tau \in \mathcal{A},\, D \in \mathcal{D}} E\left[\sum_{\tau=t}^{T^* \leq T}\left[\beta^{\tau-t}(1-a_\tau)\sum_{g=1}^{G}\delta_{igt}U_{ig\tau}\left(d_{gt}\right)\right] + \frac{\beta^{T^*}}{1-\beta}V^0\left(T^*, g_{it}\right)|I\right] \tag{2}$$

where $a_t \in \{0,1\}$ is the patient's choice of attrition at time t (0 represents the patient remaining in the trial), $A_t = [a_t, a_{t+1}, ...a_T]$ is a particular sequence of attrition decisions from the set $\mathcal{A}$, $D = [d_t, d_{t+1}, ...d_T]$ is a sequence of dose choices from the set $\mathcal{D}$, δ_{igt} is patient i's belief on being in treatment group g at time t, and $\beta \in [0,1)$ is the discount factor. Once a patient has exited the trial, $\tau > T^*$, she may not return; therefore, if patient i chooses $a_t = 1$ then $a_{t+x} = 1$ for all $x \geq 0$.

The outside option is defined as $V^0\left(t, \delta|I\right) = c_1 + c_2 t + \eta_1\delta + \eta_2 t\delta + v_{iot}$ where c, is a set of parameters to be estimated that capture the utility associated with alternative medical options available outside of the trial. The outside option is also varies based on a patient's group assignment belief through the parameters η. Patients placing a high probability on being in the experimental group may be less attracted to the outside option. The term $\left(c_2 + \eta_2\delta\right)t$ captures long term effects of participating in the trial such as prolong side effects even after exiting. A patient observes health status and side effects each period, then decides whether to remain in the trial or consume the outside option. If she decides to remain in the trial, then she chooses an optimal dose amount subject to the number of pills prescribed.

The single period utility function for patient (i) in group (g) at time (t) is defined as

$$(3) \qquad U_{igt} = -\frac{1}{\gamma}\exp\left[\gamma H_{igt}\left(d\right)\right] - \alpha_1 S1_{igt}\left(d\right) - \alpha_2 S2_{igt}\left(d\right) + \alpha_3 \ln(M + pt) + \xi_i + v_{igt}$$

which is a constant absolute risk aversion $(CARA)$ utility function.[13, 14] The vector of parameters α represent the patient's sensitivity to side effects and changes in income. The coefficient of absolute risk aversion is given by the parameter $\gamma > 0$. Risk is an important aspect of learning models as there exists a trade-off to waiting an additional period and forgoing the outside option to learn more about product quality. 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 patient's unobserved value for participating in the trial. The unobservable could capture variation in health insurance coverage among the participants. 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.

A patient's information set is defined as $I_{it} = \{\mathcal{H}_{it}, \mathcal{S}_{it}, t\}$, where $\mathcal{H}_{it} = [H_{i1} - H_{i0} - \theta_{i1} - e_i, ..., H_{it-1} - H_{i0} - \theta_{it-1} - e_i]$ is a patient's health history net their baseline health and $\mathcal{S}_{it} = [S1_{i1,}..., S1_{it,}S2_{i1,}..., S2_{it,}]$ is a patient's side effect event history. A patient learns her specific treatment effect and treatment group assignment through an application of Bayes' Law on the variables in her information set. The following section demonstrates how patients update their beliefs.

[13] The expected value of (3) conditional on a patient's information set, I_{it}, has the following closed form solution when health is normally distributed

$$\begin{aligned} E_t\left(U_{igt}|I_{it}\right) &= -\frac{1}{\gamma}\exp\left(\gamma E_t\left(H_{igt}|I_{it}\right) + \frac{\gamma^2}{2} Var_t\left(H_{igt}|I_{it}\right)\right) \\ &\quad -a_1 S1_{igt} - a_2 S2_{igt} - a_3 \ln(M + tp_{it}) + \xi_i + v_{igt} \end{aligned}$$

where expected utility is increasing in the expected value of health and decreasing in the variance of health.

[14] Both Chan and Hamilton (2006) and Crawford and Shum (2006) define utility in this fashion. Erdem and Keane (1996) use a monotonic transformation of a CARA utility function.

Patient Learning Process

Patient are assumed to use Bayesian updating to infer treatment effects, μ_{igt}, and treatment group assignment, δ_{igt}. The prior distribution of the treatment effect, μ_{ig0}, is assumed to be distributed bivariate normal

$$\mu_{ig0} = \begin{pmatrix} \lambda 1_{i0} \\ \lambda 2_{ig0} \end{pmatrix} \rightsquigarrow N\left(\underline{\mu}_{g0} = \begin{bmatrix} \underline{\lambda 1}_0 \\ \underline{\lambda 2}_{g0} \end{bmatrix}, \Sigma\right) \tag{4}$$

where $\underline{\lambda 1}_0$ is the prior dose effect and $\underline{\lambda 2}_{g0}$ is the treatment group constant. Patients know the initial randomization probability, $\delta_{g0} = 1/G$, which serves as the prior on treatment group assignment. Formally, $\delta_{igt} = \Pr(g = g_i | I_{it})$ where $g_i \in \{\text{Placebo, Experimental}\}$.

Given the assumptions that health is distributed normal within a treatment group and there are G distinct groups, patients' beliefs over health outcomes are generated by a normal mixture distribution. The normal mixture distribution is the weighted sum of G distinct normal distributions

$$L = \sum_{j=1}^{G} \delta_{igt} \int_{\lambda \in \Lambda} \phi\left(H_{it}, S_{it} | \lambda, \Omega, s1_g, s2_g, d_{it}\right) \phi\left(\lambda | \mu_{igt}, \Sigma, D_{it-1,} \delta_{igt} = 1\right) d\lambda \tag{5}$$

where $L(\cdot)$ is the likelihood function over the possible health states, Λ is the support of possible treatment effect values (λ), $\phi(\cdot)$ is the normal probability density function, and treatment group beliefs are restricted to satisfy $\sum_{g=1}^{G} \delta_{igt} = 1$. The normal mixture model captures three concepts: [1] a patient's belief that health is normally distributed; [2] the patient is uncertain as to the parameter values of the normal distribution; [3] the patient recognizes that there are G possible processes generating her health outcomes (e.g. $G = 2$ in a two armed placebo controlled trial).

Given the patient's likelihood function on health, one would normally proceed by calculating the posterior distribution on treatment effects and treatment group assignment conditional on an observed value of health. Unfortunately, a closed form solution to the pos-

terior distribution of equation (5) does not exist.[15] 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 the components of (5) for a finite number of health, side effect, and treatment effect states conditional on group assignment and treatment priors.[16] Specifically, a patient's health corresponds to 10 discrete values of log GGT, H_{it}, and two discrete values on two side effect measures, S_{it}. The treatment effect is discritized into five dose coefficients, $\underline{\lambda 1}_0$, and five treatment intercepts, $\underline{\lambda 2}_{g0}$. For a given patient, health beliefs are summarized using a 40x25 matrix of probabilities conditional on dose.

Once the probability space is discretized, an application of the discrete Bayes' Law provides the joint posterior distribution on beliefs.

$$\Pr\left(\mu^{/}_{igt}, \delta^{/}_{igt} | \mathcal{H}_{it}, \mathcal{S}_{it}, \Sigma, \overline{s1}_g, \overline{s2}_g, \Omega, D_{it}\right) = \frac{w_{igt}\left(\lambda=\mu_{igt}\right)}{\sum_G \sum_\Lambda w_{igt}(\lambda)} \tag{6}$$

$$w_{igt}(\lambda) = \Pr\left(H_{it}, S_{it} | \lambda, \overline{s1}_g, \overline{s2}_g, \Omega, d_{it}, \delta_{igt-1}\right) \Pr\left(\lambda | \underline{\mu}_{g0}, \Sigma, D_{it-1}, \delta_{igt-1}\right) \delta_{igt-1}$$

For each patient, the set of probabilities corresponding to each $\left(\mu^{/}_{igt}, \delta^{/}_{igt}\right)$ beliefs pair is stored in memory. The set of probabilities provides a finite approximation to the continuous joint distribution described in (5). 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

$$E_t\left[U | I_{it}\right] = \sum_G \sum_\Lambda U\left(\mu^{/}_{igt}, \delta^{/}_{igt}\right) \Pr\left(\mu^{/}_{igt}, \delta^{/}_{igt} | \mathcal{H}_{it}, \mathcal{S}_{it}, \Sigma, \overline{s1}_g, \overline{s2}_g, \Omega\right) \tag{7}$$

where the expectation is taken over all possible treatment groups and treatment effect states conditional on a patient's information set at time t. Patients also consider future values on

[15] Expected - Maximization algorithm provides an approximation of the posterior distribution for a normal mixture model when using all observed data points. Patients may only use her own values of health to update beliefs. Therefore, posterior values on treatment effects and treatment group assignments cannot be identified with the EM algorithm for a single patient.

[16] See appendix for more information on discretizing the probability space

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

$$
\begin{aligned}
(8) \qquad E_t\left[U|I_{it+k}\right] &= \sum_{G}\sum_{\Lambda} U\left(\mu^{/}_{igt+k},\delta^{/}_{igt+k}\right)\left(V_{igt+k}\right)\pi\left(\mathcal{H}_{it+k},\mathcal{S}_{it+k}\right) \\
V_{igt+k} &= \left[\sum_{\mathcal{H}}\sum_{\mathcal{S}} \Pr\left(\mu^{/}_{igt+k},\delta^{/}_{igt+k}|\mathcal{H}_{it+k},\mathcal{S}_{it+k},\Sigma,\overline{s1}_g,\overline{s2}_g,\Omega\right)\right] \\
\pi\left(H_{it+k},S_{it+k}\right) &= \Pr(H_{it+k},S_{it+k}|\overline{s1}_g,\overline{s2}_g,\Omega,\mu^{/}_{igt+k-1},\delta^{/}_{igt+k-1})\Pr\left(\mu^{/}_{igt+k-1},\delta^{/}_{igt+k-1}\right)
\end{aligned}
$$

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

Value Function

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 when the trial will start and end. Patients are informed of the type of compensation they may receive during the trial, the discontinuance of compensation if they are to exit the trial, and the inability to return upon exiting the trial. Each period, the patient may take one discrete action: compliance, non-compliance at a specific dose level, or attrition. Given the patient's information set, I, a patient's value function for remaining in the trial at treatment period t can be represented by the following Bellman Equation $V^{CT}\left(I\right) = \max_{A_{T^*}\epsilon\mathcal{A}_t}\left[V\left(I\left(\widehat{d}\right)|A_{T^*}\right)\right]$ where $\widehat{d}$ is the optimal dose sequence and $A_{T^*} = [a_t = 0, a_{t+1} = 0, ..a_{T^*} = 1, ..a_T = 1]$ is the sequence of discrete choices such that patient i expects to exit the trial at time $T^* \leq T$. The sub-value function $V\left(I|A_{T^*}\right)$ is defined

$$
(9) \quad V\left(I\left(\widehat{d}\right)|A_{T^*}\right) = \max_{D_{T^*}\epsilon\mathcal{D}} E_H\left(U\left(I\left(d\right)\right) + \beta E_v\left[\max V\left(I'\left(d'\right)|A_{T^*}\right), V^0\left(I'|A_{T^*}\right)|I\left(d\right)\right]\right)
$$

where the operator $E_H\left(\cdot\right)$ is the expected value of the objective function with respect to the distribution of health and the operator E_v is the expected value of the continuation value.

The value function in the last period T is $V\left(I\left(\widehat{d}_T\right)\right) = \max_{d_T} E_H\left[U\left(I\left(d_T\right)\right)\right]$. 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. Learning and experimentation are examined under two specifications of the discount factor: $\beta = 0$ and $\beta = .98$. The expectations in equation (9) are first taken over the distribution of beliefs, and then taken over the future values of v_{igt}.[17] 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 continuation function exists

$$(10)\quad E_v\left[\max\left(E_H\left[\overline{V}\left(I'|A_{T^*}\right) - v_{igt+1}|I\right],0\right)\right] = \tau\left[\Gamma + \ln\left(1+\exp\left[\frac{E_H\left[\overline{V}\left(I'|A_{T^*}\right)|I\right]}{\tau}\right]\right)\right]$$

where $V\left(I'|A_{T^*}\right) = \overline{V}\left(I'|A_{T^*}\right) - v_{igt+1}$ and Γ is the Euler constant.[18] The function $\overline{V}\left(I'|A_{T^*}\right)$ is the portion of the value function that is calculated using observable measures of health, side effects, and dose. The optimal value function can be solved using the following steps: [1] the expectation of the value function with respect to beliefs and dose sequences is taken; [2] for a given attrition sequence, $A_{T^*}\epsilon\mathcal{A}$, the optimal dose sequence, $D_{T^*}\epsilon\mathcal{D}$, is found using (10) to solve (9) backward recursively; [3] the attrition sequence A_{T^*} that maximizes the value function, $V\left(I_{it}|A_{T^*}\right)$, is the optimal value function, $V^{CT}\left(I_{it}\right)$.

Econometric Specification

This section presents the econometric method used to estimate the model. The likelihood function is comprised of two parts: patient choices and the distribution of health outcomes. Each period a patient chooses a dose level and attrition choice pair, $\langle d_t, a_t\rangle$. The econometrician observes all choice pairs made by each patient in every period. To compare these choices, the econometrician calculates the value function for each choice pair condi-

[17] See Chan and Hamilton (2006) for a discussion on the expected value of the value function conditional on the distribution of beliefs.

[18] See Berkovec and Stern (1991) and Rust (1987)

tional on a patient's information set, I_{it}, and the distribution of unobservables. Potentially, an infinite amount of dose choices exists; therefore, dose levels are discretized into a finite set of values, which eases the computation of the value function for each pair.[19] A patient's value function for any choice pair $\langle d_t, a_t \rangle$, is

$$V\left(I\left(a_t, d_t\right)\right) = \max\left[V^{CT}\left(I\left(d_t, a_t\right)\right) - V^{0}\left(I|a_t = 1\right), 0\right] \tag{11}$$

The value function is solved recursively for all possible combinations of the discrete dose values. The dose level found to maximize the value function given a set of parameter is noted by $\widehat{d}_t$. A patient remains in the trial if $V\left(I\left(\widehat{d}_t, a_t\right)\right) > 0$. The probability of surviving within a given period is

(12)

$$\Pr\left(V\left(I\left(\widehat{d}_t, a_t\right)\right) > 0\right) = \Pr\left(\overline{V}_{igt}^{CT} - \overline{V}_{igt}^{0} > v_{igt} - v_{i0t}\right) = \Pr\left(\overline{V}_{it}\left(I\left(\widehat{d}_t, a_t\right)\right) > v_{igt} - v_{i0t}\right)$$

The error term v is assumed to be distributed Type I extreme value.

Given the assumption over the distribution of the errors, the probability of remaining in the trial is $\Pr\left(a_{it}{=}0|d_{it} = \widehat{d}, a_{it-1} = 0, e_i, \xi_i\right) = \frac{\exp\left(\overline{V}_{it}\left(I\left(\widehat{d}_t, a_t\right)\right)/\tau\right)}{1+\exp\left(\overline{V}_{it}\left(I\left(\widehat{d}_t, a_t\right)\right)/\tau\right)}$. The "no re-entry" policy of an RCT are captured by setting $\Pr\left(a_{it}{=}0|a_{it-1}{=}\,1\right) = 0$ and $\Pr\left(a_{it}{=}1|a_{it-1}{=}\,1, e_i, \xi_i\right) = 1$. The unconditional probability of attrition at time t is then given by equation (13)

$$\Pr(A = A_{T^*}, D = D_{T^*}|e_i, \xi_i) = \frac{1}{1+\exp\left(\overline{V}_{iT^*}\left(I\left(\widehat{d}_{T^*}, a_{T^*}\right)\right)/\tau\right)} \left[\prod_{t=1}^{T^*-1} \frac{\exp\left(\overline{V}_{it}\left(I\left(\widehat{d}_t, a_t\right)\right)/\tau\right)}{1+\exp\left(\overline{V}_{it}\left(I\left(\widehat{d}_t, a_t\right)\right)/\tau\right)}\right] \tag{13}$$

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 $Z_{it} = [e_i + \varepsilon_{ig1}, ...e_i + \varepsilon_{igt}, u1_{ig1}, ..., u1_{igt}, u2_{ig1}, ..., u2_{igt}]'$ where the deviations are stacked in the following order: health, side effect [1], and side effect [2]. Given the assumption of joint

[19] In practice, four discrete dose amounts are chosen for each period. The four discrete dose amounts are chosen by finding the quartilzes of dose consumption in each period.

normality between health and side effects, the probability density function of patient i's outcome measurements is given by $f\left(Z_{it},\Omega\right)=\frac{1}{(2\pi)^{3t/2}|\Omega|^{1/2}}\exp\left[-\frac{1}{2}Z_{it}'\Omega^{-1}Z_{it}\right]$ where t is the period when patient i exits the trial, and Ω is the 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. Note, each measure of health is dependent on the predicted dose amount at time t, $\widehat{d}_t$. Therefore, we incorporate the observed dose amounts with the predicted doses by including a probability density function for dose. As in the reduce form case, we assume dose is distributed normal $d_{it}\rightarrow N\left(\widehat{d}_t,\sigma_d^2\right)$. Additional covariates may be added to the dose density function, but to ease the computation burden no additional covariates are used in estimation.

Define the set of parameters to be estimated as $\Psi=\{\theta,\varphi,\Omega,\Sigma,\mu_g^p,\alpha_1,\alpha_2,\gamma,\eta,\sigma_\xi^2,\tau\}$. The likelihood contribution of the i'th patient conditional on the unobservable errors is

$$L_i\left(\Psi|a_{it},\mathcal{H}_{it},\mathcal{S}_{it},e_i,\xi_i\right)=f\left(Z_{it},\Omega|a_{it},\widehat{d}_t,\theta,\varphi\right)\Pr\left(a_{it}|\widehat{d}_t,e_i,\xi_i,\Psi\right)\Pr\left(d_{it}=\widehat{d}_t\right) \tag{14}$$

The expectation of the likelihood contribution is taken with respect to the unobserved errors, e_i and ξ_i.

$$E\left[L_i\right]=\int\int L_i\left(\Psi|a_{it},\mathcal{H}_{it},\mathcal{S}_{it},e_i,\xi_i\right)f\left(e_i\right)f\left(\xi_i\right)de_i d\epsilon_i \tag{15}$$

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 independent and distributed normal. The respective variance terms, σ_e^2 and σ_ξ^2, are parameters to be estimated. While no closed form solution exists for the expected value in equation (15), 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

$$\log\widehat{L}_i\left(\Psi|a_{it},d_{it},\mathcal{H}_{it},\mathcal{S}_{it}\right)=\log\left(\frac{1}{R^2}\sum_{r=1}^{R}\sum_{k=1}^{R}L_i\left(\Psi|a_{it},d_{it},\mathcal{H}_{it},\mathcal{S}_{it},e_r,\xi_k\right)\right) \tag{16}$$

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

Results

The structural parameter estimates are found in Table 5 and a set of reduced form results are found in Table 4 for comparison. The model is estimated twice under two different assumptions of patient behavior: "forward looking" and myopic. The second column of Table 5 contains the estimated coefficients for the dynamic model where the weekly discount factor is fixed at $\beta = 0.98$. The third column contains the estimated coefficients for the static (myopic) model, $\beta = 0$. To allow for comparisons between models, the variance of the extreme value error is restricted to be the same in both models.

First consider the health outcome parameter estimates as they pertain to side effects. On average, individuals in the experimental group are 3% more likely to experience fatigue (side effect 1) and 8% more likely to experience parathesia (side effect 2) than individuals in the placebo group. Parathesia conveys group information better than fatigue, because the difference in side effect frequency between treatment groups is greater for parathesia. Both variables are found to be statistically significant at the 1% level in both models.

GGT measures the level of liver enzymes associated with alcohol presence. High levels of GGT are associated with poor health outcomes. The estimated health coefficients represent the marginal effect of each explanatory variable on the natural logarithmic change in health ($\Delta \log(GGT)$).[20] In the absence of medication, patients in the dynamic model experienced a 14 percent decrease (placebo constant) in GGT levels during the first period, but experience a slow and statistically *insignificant* rise in GGT every period thereafter. The treatment effect in the experimental group is decomposed into two separate effects, a constant and a dose effect. Patients in the Topamax group experience a 45% (TE) increase in GGT levels relative to patients in the placebo group during the first period (constant). The effect of topamax in milligrams is captured as an elasticity. The natural log of Topamax consumption

[20] For small changes, the difference in $log(x + \Delta) - \log(x)$ represents the precent change in x.

measure in milligrams is the explanatory variable of interest. A 1 percent increase in the number of milligrams consumed leads to a .08% improvement in health. This estimate is roughly twice the magnitude found using ordinary least squares, which estimates the dose effect to be .047% (see table). Unobserved variation in health among patients is primarily driven by patient specific unobserved heterogeneity. Patient specific unobserved heterogeneity constitutes 76.5% of the unobserved health error.

The static model produce slightly different results. In the absence of medication, patients in the static model experienced a 25 percent decrease (placebo constant) in GGT levels during the first period. Health continues to improve in each subsequent period at rate of 1.6%, but this improvement is not statistically significant. Patients in the Topamax group experience a 45% (TE) increase in GGT levels relative to patients in the placebo group during the first period (constant). The static model estimate of the treatment dose effect is smaller in magnitude than the dynamic model. A one percent increase in milligram concentration yields a 0.05% improvement in health. Lastly, the standard errors in the dynamic model are strictly smaller in magnitude than the those found in the static model.

Prior Distribution

The estimated prior distribution of the treatment effect provides insight on a patient's level of optimism (pessimism) for participating in the trial. In the dynamic model, the estimated health priors suggest patients believe the experimental drug leads to an initial 11% decrease in GGT levels, followed by a 10% decrease in each subsequent week in compared to the progression of health without medication. The static model provides much more optimistic beliefs with patients believing in a 94% decrease in GGT levels followed by a more than 300% improvement in health each subsequent period. These estimates imply that patients *expect* an immediate improvement in health if assigned to the experimental group.

On the other hand, the estimated treatment effect, found from the health equation, suggests health status for patients in the experimental group improves only after consuming 42 mg of Topamax in the dynamic model and 44 mg in the static model. This result implies

patients' beliefs are initially optimistic about the potential health benefits associated with participating the experiment. The estimated prior variance terms in the dynamic model are *larger* than those in the static model, suggesting that patients have flatter priors (less confidence in their estimates) in the dynamic model (alternatively, myopic patients hold on to their prior beliefs longer than forward looking patients).

Utility Parameters

The largest differences between the dynamic and static models are found in the estimates of the utility parameters and the outside option. The utility 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. The coefficient of absolute risk aversion is found to be roughly the same in both models: $\gamma = 0.6773$ in the dynamic model and $\gamma = 0.6096$ in the static model. Patients can minimize risk in two dimensions when dose is offered as a choice variable. The two dimensions include the number of health signals and the variation in health signals. The number of health signals is determined by how long a patient chooses to participate in the experiment. The variation of health signals can be manipulated using dose as a choice variables. Similar to the concepts of linear regression, a patient's estimate of the treatment effect improve when there is more variation in dose. As her estimates improve, the level of uncertainty decreases.

Side effect sensitivity is smaller in the dynamic model than in the static model. The difference in magnitude between models is explained by the differences in behavioral assumptions of the two models. The direct impact of side effects is concentrated in the current period for patients in the static model, but patients in the dynamic model spread the disutility of side effects across current and future periods. Forward looking patients appear to tolerate side effect much better than myopic patients due to these behavior assumptions. For example, parathesia (side effect 2) decrease utility by a statically insignificant amount within the dynamic framework, but in the static model experiencing parathesia is equivalent to decreasing a patient income by 1.4%. On the other hand, fatigue is found to be a statistically

significant source of disutility in both models. In the dynamic model an instance of fatigue is equivalent to an increase of GGT from 300 to 302. In the myopic model, experiencing fatigue is equivalent to increasing GGT from 300 to 307 (or decreasing income by 2%). Both coefficients are statistically significant at the 1% level.

The model captures a patient's response to compensation by estimating a patient's income elasticity. In the myopic setting, the effect of compensation on attrition is strong. A 1% increase in income leads to a 27% increase in utility. In the dynamic setting, a patient's income elasticity is -.04% and is not statistically different from zero. The large difference in income elasticity between the two models displays the sensitivity of parameter estimates with respect to the assumption placed on the discount factor.[21]

Lastly, I include an additional utility parameter in the last period of the experiment 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 1.67 utils (equivalent of reducing GGT from 300 to 285) in the dynamic model and 0.09 utils (equivalent of reducing GGT from 300 to 299) 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.

Outside Option

The outside option 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 external options including detoxification clinics, support meetings, or drinking. The outside option also internalizes long term effects of the experimental drug even after the patient has stopped taking the drug. These long term effects are captured

[21] Using patient demographic characteristics to estimate a heterogenous discount factor for each patient may alleviate this problem.

by interacting treatment group beliefs with the value of the outside option. Long term effects can range from chronic side effect experiences or permanent improvements in health. Additionally, the inclusion of patient beliefs in the outside option can capture "placebo effects," which may improve (or hinder) unobserved components of health.

The study's data does not contain information on patient choices after exiting the trial. Therefore, the model takes a simplistic approach in handling external options. The outside option is comprise of a constant and a linear time trend. These two components are allowed to vary with patient beliefs on group assignment. The outside option constant represents the benefit measured in utils of exiting the trial in the first period. The value of the outside option constant is 1.59 in the dynamic model and -2.38 in the myopic model. The outside option improves quickly in subsequent periods as indicated by the time trend (dynamic model = 1.71 and myopic model = 2.45). These improvements can be the result of a decrease in alcohol prices or increase access to alternative treatment options. 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 value of the outside option decreases in both structural models. The dynamic model finds a patients utility for the outside option decreases by -6.06 utils as beliefs and time increase. The myopic model also finds a decreasing effect on the outside option (-4.49) as belief and time increase. Patients who *believe* they are receiving treatment are *less* likely to exit the trial, ceteris paribus. Therefore, the value of participating in the clinical trial increases as the patient's belief on receiving treatment increases even when health remains constant. This result may be interpreted as a type of "participation placebo effect."

Goodness of Fit

The ability of the dynamic model to explain patient behavior is measured by simulating dose and attrition decisions conditional on the estimated parameters, then comparing these outcomes with observed decisions in the experiment. First, the initial health state for each patient is their original GGT values in period 1. These values are used as a starting

point on health histories. To capture unobservable health shocks, 10 patient specific errors are drawn from the unobserved health shock and a four random health shocks from the idosyncratic health error. 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, the estimated treatment effects, and the health-side effect covariance matrix to complete a patient's health and side effect history.[22] 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 $(\widehat{\tau}\pi)^2/6$.

The simulated health histories and utility shocks are then used to solve each patient's value function. A patient's dose choice may take on one of d possible values in a given evaluation period. There are a total of d^{T-1} possible dose combinations for any given patient over the duration of the experiment lasting T periods.[23] The value function is solve for each dose sequence. The dose sequence that maximizes the value function is the optimal dose choice for the patient. If a patient's value function is > 0, then the patient remains in the trial; otherwise, the patient exits the trial.

Figure 4 exhibits the observed survival rate versus the predicted survival rate for each model. The dynamic model explains 91% of the variation in overall survival (attrition) rates between weeks 0 - 9, but precision is lost when evaluating predicted survival rates in the last period. Overall, the dynamic model can account for 45.8% of the variation in attrition decisions. On the other hand, the static model provides a poor fit to the data. The static model tends to over-estimate the rate of attrition throughout the experiment and explains only 38.4% of the attrition rate. Both models tend to underpredict survival rates in each

[22] For each patient in the orginial trial, fourty health histories are simulated using the different combinations of patient specific errors and random health shocks.

[23] I use both d=3 and d=4 for a total of 81 and 256 possible dose combinations, respectively.

treatment group. Again, the dynamic model out performs the static model in predicting attrition. The dynamic model captures 62% of the variation in attrition within the placebo group and 31% within the Topamax group. The static model captures 51.5% of the variation within the placebo group and only 27.4% within the Topamax group. By accounting for non-compliant behavior, the dynamic models is better able to predict attrition decisions between treatment groups by exploiting relative differences in dose choices between groups.

Goodness of Fit: Overall Survival Probabilities

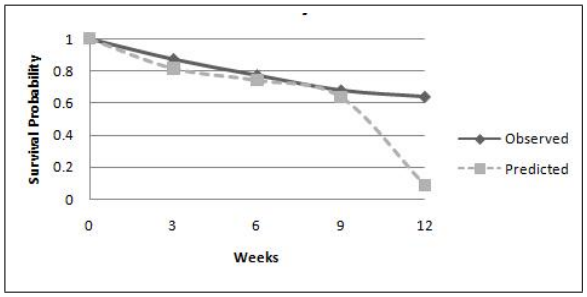

Dynamic Model

Goodness of Fit: Survival Probabilities by Group

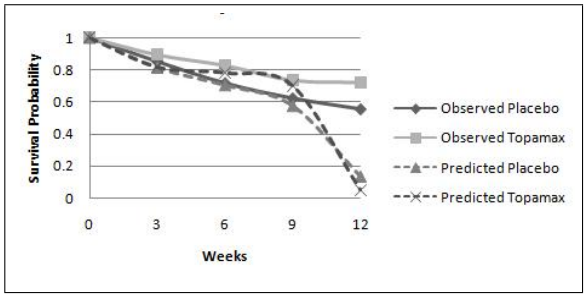

Dynamic Model

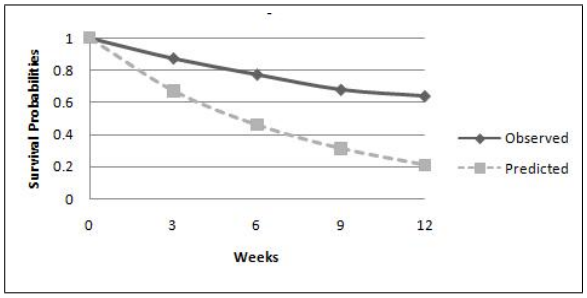

Static Model

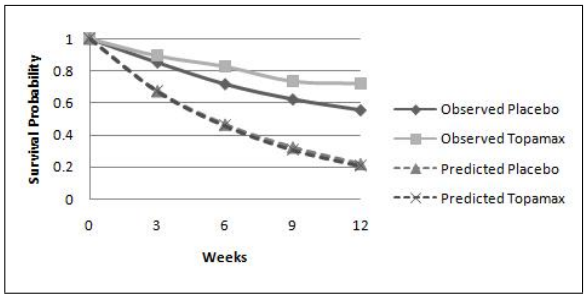

Static Model

Figure 4

Lastly, I compare the relative goodness of fit between the two structural models using a likelihood ratio test. The test setup includes a null hypothesis supporting patient behavior which is not "forward looking" (the static model) and an alternative hypothesis supporting forward looking patients who strategically experiment with dose to increase their rate of

learning. The F-statistic is the log difference in the likelihood functions of the two models, $F = -2(LLF_{static} - LLF_{dynamic}) = 254.2$. The likelihood ratio test rejects the null hypothesis at the 1% level of significance. Therefore, attrition behavior in clinical trials is *more likely* caused by forward looking patients *who are learning* treatment effects through dose experimentation than a static utility maximizers.

Dose choice

A primary objective of this work is to study patient dose choice. The model allows for patients to choose a level of dose, which maximizes both utility (maximize health/minimize side effects) and learning. In this section, I compare the dose choices made by patients in both treatment groups under two learning assumptions. The dynamic model with "forward looking" patients does correctly predict that patients in the Topamax group are more likely to choose higher consumption levels than patient in the placebo group, but the amount consumed in the Topamax group is overstated and the amount in the placebo group is understated. The static model with "short sighted" patients also predicts the experimental group will consume more drugs than the placebo group, but the magnitude of consumption is understated in both treatment groups. On average, dose consumption is 6.5% higher in the Topamax group and 9.2% higher in the placebo when comparing forwarding looking versus myopic patients. This result is consistent with theoretical models of learning with experimentation. Mirman, Grossman, and Kihlstrom (1977) discuss a theoretical model of consumer experimentation with a product of unreliable quality, eg. a drug. The authors find that if health increases as dose consumption increases, then the patient chooses a level of dose that *exceeds* the optimal single period utility maximizing dose level. By "over dosing" a patient jointly maximizes learning and utility. The inclusion of patient demographic characteristic as explanatory variables on dose may be used to explain the difference in magnitude

between predicted and observed dose levels.

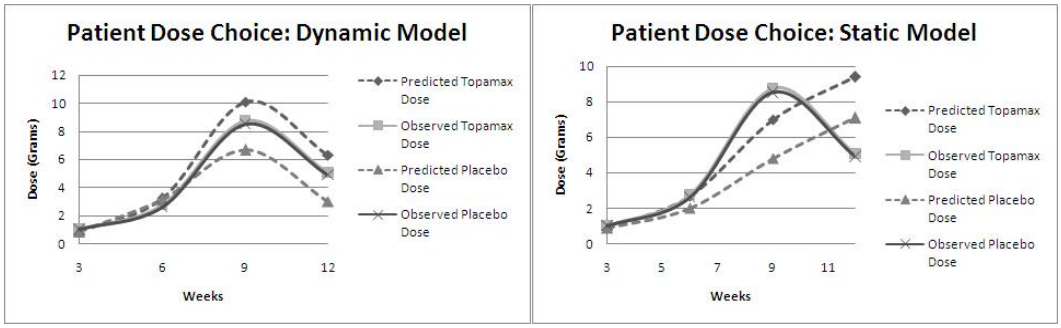

Figure 5: Dose Simulation

Treatment Group Beliefs: Simulation

If patients are proactive about learning, then they will tend to vary dose consumption in such a matter as to maximizes learning. To display this phenomenon, I use the simulated dataset of health histories and side effects along with the structural parameters to find the evolution of patient group beliefs. In the static model, patients choose dose consumption as to maximize single period utility conditional on treatment group/drug quality beliefs. In the dynamic model, patients choose to maximize lifetime utility which is subject to uncertainty. These forward looking patient take a proactive approach to reducing their uncertainty by systematically varying dose in each period as to maximize learning. For this reason, aprior one would expect learning rates to be higher in the dynamic model than the static model. Figure 6 contains the average belief probability of being in the experimental group for both treatment groups. As a patient's belief tends towards one, then the patient believes she is receiving the experimental. As a patient's belief tends towards zero, then she believes she is receiving the placebo pill. In the dynamic model, patients' group beliefs are on average greater than or equal to 0.5. Forward looking patients place more probability in believing they are receiving treatment than patients in the static model. Note as well that patients' beliefs experience larger changes (variation) in the static model than patients' beliefs. Experimental group patients within the dynamic model converge faster to their

actual group assignment (by week 6) than experimental patients in the static model (by week 9). The simulation exercise produces mixed results on learning in the placebo group of both models. While it appears that the static model has placebo patients' beliefs converging fast to zero than the dynamic model, it is difficult to determine if these beliefs will remain there or rise as they did between weeks 3 and 9. Further, it may become increasing difficult for a placebo patient to distinguish between a poor performing experimental drug (TE = 0) or a placebo pill. For example, one placebo patient may determine that she is in the experimental group, but the treatment effect is small. Therefore, it should be quite reasonable to expect clinical trial exit interviews to report that on average placebo patients could not identify their treatment group, but could identify that the drug they were consuming provided little to no benefit. This result of course ignores the possibility of placebo effects.

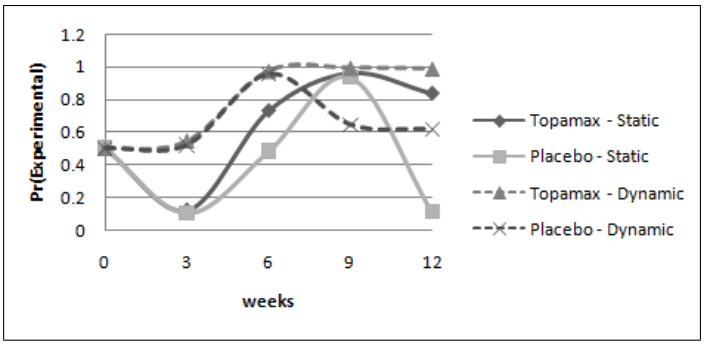

Figure 6: Patient Beliefs on Treatment Group Assignment

Caveats

Although there are many advantages to using structural modeling in the analysis of patient behavior there are some caveats of which the reader should be aware. The model is only as good as the underlying behavioral assumptions. In this regard, extensions may be made of this model where patients display alternative preferences such as constant relative risk aversion or quadratic utility. The model also assumes each patient can update her beliefs using Bayes' Law. Bayesian learning is currently the standard, but recent developments

in behavioral economics have identified belief "anchoring" as a potential source of bias. Anchoring bias occurs when patients are less willing to update their beliefs away from starting values when presented with new evidence (see Tversky and Kahneman, 1974). Lastly, the model remains agnostic in regards to "placebo effects." Presumably, placebo effects are caused by patients *believing* they are receiving the experimental treatment. A potential method to capture these effects within the current model structure is to allow treatment effects in the health equation to be interacted with patient group beliefs.

Conclusion

This paper presents an empirical examination of patient learning and experimentation in the context of clinical trials. A structural model of patient behavior and learning is proposed to capture the proactive nature of clinical trial participants when choosing dose consumption. The structural model is used to test two leading hypotheses, which explain the causes of patient non-compliance. The first hypothesis states that patient non-compliance is derived from side effect avoidance. The second hypothesis believes that patients are forward looking and deliberately vary their dose consumption to infer their treatment group assignment. This work demonstrates the first empirical study to capture learning and dose quantity experimentation in the market for pharmaceutical goods. The structural model is estimated using data from a clinical trial studying the effects of Topamax on alcoholism. A likelihood ratio test is used to reject the null hypothesis that patients choose to be noncompliant as a result of side-effect avoidance only. In addition to the learning results, the model does identify some evidences of attrition bias caused by learning. The estimated treatment effect found in the structural model is nearly twice the magnitude of the OLS estimate.

The results found in this study give rise to at least two future lines of research. First, work is still needed to determine an optimal incentive package to minimize the attrition bias on treatment effect measures associated with patient learning in clinical trials. For example, researchers could propose a progressive financial reward given to those patients who remain

compliant each period of the experiment. Second, the interpretation of the discount factor commonly found in dynamic discrete choice models is 1/(1+r) where r is the real interest rate. But do unhealthy individuals discount at the same rate as healthy individuals? Considering the dynamic model ($\beta = 0.98$) tends to overstate the effects of experimentation and the static model ($\beta = 0$) tends to understate these effects, it may be fruitful to estimate the discount factor.

Finally, the subject of learning under ambiguity is not limited to clinical trials, but may be applied in many other topics of interest to economist. In the area of education, the study of student attrition from college courses could be viewed as learning ones ability (treatment effect) in a class given the ambiguity of a teacher's quality (group assignment). In labor economics, the performance of a particular employee (treatment effect) is dependent upon the quality of her manager (group assignment). With respect to quantity experimentation, the industrial organization literature has many examples. The most notable example is Mirman, Samuelson, and Urbano (1993) where a monopolist varies quantity to learn demand even when these quantity choices are not profit maximizing with respect to static profits. These application can further our understanding of how economic agents use and manipulate signals to verify an unknown quality and accelerate learning.

Bibliography

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

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

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

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

Chan T. and Hamilton B. "Learning, Private Information and the Economic Evaluation of Randomized Experiments," *Journal of Political Economy*, 2006, vol. 114, no. 6

Coscelli, A. and Shum, M. "An empirical model of learning and patient spillovers in new drug entry," *Journal of Econometrics*, October 2004 vol. 122(2), pages 213-246.

Crawford, G. and Shum, M. "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, Vol. 73, No. 4 (Jul., 2005), pp. 1137-1173

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

Efron, B. and Feldman, D. "Compliance as an explanatory variable in clinical trials" *Journal of the American Statistical Association* 1991; 86:9-26.

Ellickson, P., Stern, S., and Trajtenberg, M. "Patient Welfare and Patient Compliance: an empirical framework for measuring the benefits from pharmaceutical innovation," NBER Working Paper 1999; 6890

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

Fernandez, J. "An Empirical Model of Learning under Ambiguity: The Case of Clinical Trials," mimeo, Department of Economics, University of Louisville. (2008)

Fischer, K. and Goetghebeur, E. "Structural mean effects of noncompliance: Estimating interaction with baseline prognosis and selection effects," *Journal of the American Statistical Association* 2004; 99: 918-928.

Goetghebeur, E. and Lapp, K. "The effect of treatment compliance in a placebo-controlled trial: regression with unpaired data," *Applied Statistics* 1997; 46: 351-364 .

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

Hrobjartsson, A. Gotzsche, P. C. "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 (2001).

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

Lamiraud, K., and P. Geoffard. Therapeutic non-adherence: a rational behavior revealing patient preferences?. *Health Economics* 16:11, 1185-1204 (2007)

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," *Journal of Econometrics* 144 (2008): 341-351

Mirman, L., Grossman, S., and Kihlstrom, R. "A Bayesian Approach to the Production of Information and Learning by Doing," *Review of Economic Studies*, Oct, 1977.

Mirman, L., Samuelson, L., and Urbano, A. "Monopoly Experimentation," *International Economic Review*, Vol. 34, No. 3 (Aug., 1993), pp. 549-563

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

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

Rust, J. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher" (1987) *Econometrica* 55–5, 9991033.

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

Tversky, A. and D. Kahneman, "Judgment under uncertainty: heuristics and biases," *Science*, 185(4157), 1124-1131 (1974).

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

Appendix: Additional Figures and Tables

Table 1: Patient Characterisitcs

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 2: Side Effects

Side Effect	Severity		Duration		N	
	Placebo	Topamax	Placebo	Topamax	Placebo	Topamax
Anorexia	2.1 (0.86)	2.29 (0.79)	5.02 (2.29)	5.48 (2.8)	41	65
Ataxia	1.58 (0.51)	2 (1)	2.56 (1.88)	4.86 (3.76)	12	7
Confusion	1.64 (0.92)	2.04 (0.79)	4.27 (2.9)	4.93 (2.72)	11	25
Depression	2.36 (0.87)	2.19 (0.83)	5.25 (2.7)	5.75 (2.03)	39	54
Dizziness	1.64 (0.84)	1.83 (0.7)	3.54 (2.37)	3.52 (2.38)	14	30
Fatigue	2.16 (0.76)	2.15 (0.78)	4.8 (2.87)	5.58 (2.69)	82	127
Nervousness/Anxiety	2.28 (0.87)	2.12 (0.75)	4.74 (2.51)	4.42 (2.76)	67	67
Nystagmus	1 (0)	1.75 (0.71)	20 (0)	4.75 (2.66)	1	8
Parasthesia	1.96 (0.91)	1.9 (0.78)	3.09 (2.33)	5.31 (2.41)	24	197
Psychomotor slowing	1.92 (0.8)	2.07 (0.64)	4.17 (3.35)	5.59 (2)	12	42
Somnolence	2.15 (0.72)	2.15 (0.74)	4.93 (2.82)	5.16 (2.32)	68	82
Speech disorder	1.43 (0.79)	1.61 (0.61)	3.71 (4.5)	3.89 (3.46)	7	18
Tremor	1.5 (0.52)	1.83 (0.58)	3.92 (2.64)	4.08 (2.91)	12	12
Weight loss	1.57 (0.74)	1.61(0.68)	6.55 (1.95)	6.55 (3.19)	35	93

Table 3: Tobit Model of Non-Compliance

Variable	Coefficient	(t-value)
Time	49.48[a]	10.76
Gender	42.35	0.87
AGE	-7.06[a]	-8.41
BMI	0.3014	0.23
Treatment	-120.13[a]	-2.76
White	-3.38	-0.09
Protestant	-250.78[a]	-5.08
Catholic	-198.85[a]	-4.48
Education	-19.02[a]	-7.20
Skilled	276.13[a]	3.53
Semi-skilled	207.04[a]	5.33
Income	0.0008	0.32
(Time)x(Treatment)	6.60	1.12
Constant	-140.26[a]	-3.99
σ_e	750.76[a]	25.56
LL	-4080	
Pseudo R^2	0.1112	
N	1395	

Note: Significance level: a =99%, b = 95%, and c = 90%

Table 4: Sample Selection Health Equation with Dose

Variable	(1)	(2)	(3)	(4)	(5)	(6)
(Constant)	−0.1149 (0.0600)	−0.1074 (0.0600)	−0.0922 (0.0617)			
Treatment	0.3510 (0.1863)	0.3914 (0.1872)	0.4524 ^{b} (0.1958)			
Time	−0.0118 (0.0073)	−0.0218 ^{b} (0.0091)	−0.0208 ^{b} (0.0092)	−0.0192 ^{a} (0.004)	−0.034 ^{a} (0.007)	−0.0288 ^{a} (0.008)
log(mg)	−0.0477 ^{c} (0.0215)	−0.0502 ^{b} (0.0215)	−0.0637 ^{b} (0.0249)			
$\widehat{\log(\text{mg})}$				−0.0213 (0.02)	−0.015 (0.02)	−0.0624 (0.035)
Mills		0.200 (0.1116)	0.1305 (0.1298)		0.297 (0.126)	0.177 (0.145)
(Treatment)x(Mills)			0.2042 (0.1926)			0.357 (0.214)
R^2	4.10%	4.9%	5.2%			
Patient FE ?	no	no	no	yes	yes	yes

Note: (1) Dependent Variable: LN(GGT); (2) Significance level: a =99%, b = 95%, and c = 90%

Table 5a: Learning with Experimentation Structural Model Parameters

Variables	Dynamic	Static
Prior Distribution of Treatment Effects (TE)		
Prior TE constant	-0.1131^a (0.0430)	-0.9410^a (0.0467)
Prior TE dose elasticity	-0.1024^a (0.0375)	-3.8643^a (0.1673)
Prior Variance (TE constant)	0.9266^a (0.0361)	0.7296^a (0.0567)
Prior Variance (TE dose elasticity)	0.9990^a (0.0368)	0.6666^a (0.0304)
Utility Function Parameters		
Utility: Unobserved Heterogeneity	0.8190^a (0.0307)	0.8156^a (0.0273)
Coef. of Absolute Risk Aversion	0.6773^a (0.0260)	0.6096^a (0.1798)
Income Elasticity	-0.0441 (0.0440)	27.2026^a (1.4389)
Side Effect 1 Sensitivity	-0.1406^a (0.0375)	-0.5879^a (0.0413)
Side Effect 2 Sensitivity	-0.0107 (0.0375)	-0.3781^a (0.0386)
Finishing reward	1.6696^a (0.0376)	0.0895 (0.1083)
Tau	170.82^a (13.9142)	170.82 (fixed)
Outside Option Parameters		
Outside Option Constant	1.5890^a (0.0379)	-2.3766^a (0.2007)
Outside Option Time	1.7093^a (0.1543)	2.4548^a (0.1221)
Outside Option E_{it}(G = Treat)	0.0031 (0.0762)	-1.917^a (0.0938)
Outside Option E_{it}(G = Treat)*Time	-6.064^a (0.0656)	-4.494^a (0.1195)
Discount Factor	0.98 (fixed)	0 (fixed)
Log Likelihood Value	-698.11	-825.1884
F	254.1568^a	
N	149	

Note: (1) standard errors are in parenthesis; (2) F stat: Likelihood ratio test (dynamic vs static);

(3) Significance level: a =99%, b = 95%, and c = 90%; (4) † evaluated at mean dose = log(3361 mg)

Table 5b: Learning with Experimentation Structural Model Parameters

Health Outcomes Parameters

Variables	Dynamic	Static
TE	0.4510^a (0.0073)	0.4495^c (0.2541)
TE log(dose (mg) elasticity	-0.0829^a (0.0020)	-0.0524^c (0.0298)
Placebo constant	-0.1399^a (0.0368)	-0.2513^a (0.0511)
Placebo Time	0.0017 (0.0167)	-0.0161 (0.0232)
Health Variance	0.0692^a (0.0029)	0.0689^a (0.0051)
Health: Unobserved Heterogeneity	0.2198^a (0.0083)	0.2106^a (0.0281)
Placebo Side Effect 1 Mean	0.1162^a (0.0041)	0.1122 (0.0184)
Placebo Side Effect 2 Mean	0.0625^a (0.0059)	0.0574^a (0.0141)
Topamax Side Effect 1 Mean†	0.1423^a (0.0204)	0.1459^a (0.0304)
Topamax Side Effect 2 Mean†	0.1400^a (0.0213)	0.1427^a (0.0343)
Variance Side Effect 1	0.1115^a (0.0219)	0.1103^a (0.0204)
Variance Side Effect 2	0.0861^a (0.0211)	0.0864^a (0.0207)
Cov(Health,Side Effect 1)	0.0069 (0.0455)	0.0063 (0.0761)
Cov(Health, Side Effect 2)	-0.0020 (0.0267)	-0.0026 (0.0248)
Cov(Side Effect 1, Side Effect 2)	0.0530^c (0.0299)	0.0529^c (0.0293)

Note: (1) standard errors are in parenthesis; (2) F stat: Likelihood ratio test (dynamic vs static);

(3) Significance level: a =99%, b = 95%, and c = 90%; (4) † evaluated at mean dose = log(3361 mg)

Appendix I: Discretization Method

The learning model present in this paper requires that patient's beliefs on health outcomes follow a normal mixture model[24]. The model is a linear

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

[24]This distribution is commonly used by computer scientist in "machine learning" applications. See "unsupervised learning" for more informaiton.

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.

$$\text{(A2)} \qquad \mu_{ig0} = \begin{pmatrix} \lambda 1_{i0} \\ \lambda 2_{ig0} \end{pmatrix} \tilde{} N \left(\underline{\mu}_{g0} = \begin{bmatrix} \underline{\lambda 1}_0 \\ \underline{\lambda 2}_{g0} \end{bmatrix}, \Sigma = \begin{bmatrix} \sigma^2_{\lambda 1} & 0 \\ 0 & \sigma^2_{\lambda 2} \end{bmatrix} \right)$$

Initially, the priors on the dose effect, $\lambda 1_{i0}$, 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 $\lambda 1$ and $\lambda 2$, $\Lambda = [\lambda_{-2}, .\lambda_0, ..\lambda_2]$, 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

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

where $\Phi(\cdot)$ is the normal cumulative density function. An analogous equation defines the probability mass function for $\lambda 2$. 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 GGT levels, reports of fatigue, and report of parathesia. Initially, the side effect state space takes on a simple definition: 1 if a side effect is experienced and 0 otherwise[25]. GGT levels are dis-

[25] This simple definition is necessary to increase computation efficiency and decrease memory requirements.

cretize into 10 equispace values over the sample range. Therefore, the health outcomes state space for a given value of λ, side effect means $(s1_g, s2_g)$, 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 (# of health states) x (# of experimental effect states) x (# of Groups) $=$ 2000, which must be carried in memory for each patient and evaluated at each level of dose in each evaluation period. If there are d discrete values for drug dose in each t periods, then there are d^t potential dose combination over the length of the experiment. Therefore, the total state space is $2000d^t$ for each patient.

Appendix II: Clinical Trial Documentation

University of Texas Health Science Center at San Antonio

We are asking you to take part in a research study of Topamax for the treatment of alcohol dependence. Currently, there is no standard treatment for people who are alcohol dependent. While there are several treatment programs available, it is not clear how well these programs work. We want to find out if a drug called Topamax (Topiramate), a drug approved by the FDA in December 1996 as adjunctive therapy for all adults with partial onset seizures (maximum dose 400 mg/day) and Brief Behavioral Compliance Enhancement Treatment (BBCET) is more effective than BBCET without Topamax (Placebo medication) for the treatment of alcohol dependence. We are asking you to take part in this study because you are alcohol dependent and wish to stop drinking alcohol. It is anticipated that approximately 80 patients will participate in this study.

If you decide to take part in this study, we will conduct an initial screening visit, in which you will have a complete medical and psychiatric evaluation. If you qualify (you do not have any medical and/or psychiatric conditions that jeopardize your physical or mental health) for the study, you will be asked to participate in a 12-week double-blind

randomized research treatment study. You will be randomized (like flipping a coin) to one of the following treatments: a) Topamax up to 400 mg/day + BBCET, or b) Placebo + BBCET. There is a 50% chance that you will receive Topamax and a 50% chance that you will receive placebo. This is a double-blind study, which means neither you nor your doctor will know which treatment you are receiving. However, this information is available if needed to treat you in an emergency. You will be required to attend the clinic weekly for medication and BBCET. BBCET sessions are 15-20 minutes in length, and involve you meeting with a member of the research staff to discuss medication issues and your status in the treatment program. During the 12-week study you will be required to complete self-report measures (pencil and paper rating forms) of substance use, craving, withdrawal symptoms, mood, and other psychological states (e.g., feelings), and provide both urine and blood samples for drug screening, pregnancy screening, and laboratory evaluations. Your study medication will be gradually increased during the first 10 weeks of the study, to the maximum tolerated dose (400mg/day or placebo) allowed by the study. Dr. Johnson may adjust your study medication dose as is necessary.

You will be expected to attend the clinic for an initial screening visit (approximately 6 hours) and a weekly visit (approximately 20-30 minutes) for each of the 12 study weeks.

Drinking levels required for participation in this study are set at > (greater than or equal to) 14 alcohol units/week for women and > (greater than or equal to) 21 units/week for men in the last 30 days to avoid enrolling individuals who may experience severe withdrawal or medical complications. Alcohol units are measured by the amount and type of beverages you consume (e.g., 12 oz. of beer = 1 alcohol unit, 1.5 oz. of liquor = 1). Individuals who are experiencing physical signs of withdrawal will be referred for medical care.

Risks

The most commonly observed side-effects associated with the use of Topamax at dosages between 200 to 400 mg/day in previous studies are: somnolence (sleepiness), dizziness, ataxia (problems with coordination), speech disorders and related speech problems, psychomotor slowing (slowed thinking and movement), nystagmus (eyes jumping), and paresthesia (tingling of the skin). These side effects are not related to the dose of Topamax.

In clinical trials, 11% of patients discontinued due to adverse events. Compared to placebo, 1% of patients treated with Topiramate (Topamax) (200-400 mg) had a greater incidence of adverse events. Most patients who experienced adverse events during the first 8-weeks of treatment, no longer experienced these effects at their last visit.

The most common side effects which may appear as the dose is increased from 200 to 1,000 mg/day are: fatigue, nervousness, difficulty with concentration or attention, confusion, depression, anorexia (i.e., loss of appetite), language problems, anxiety, mood problems, cognitive problems not otherwise specified, weight decrease, and tremor. In addition, a total of 1.5% of patients exposed to Topamax can develop kidney stones. The use of Topamax at the same time as other carbonic anhydrase inhibitors, such as acetazolamide and zonisamide, may increase the risk of kidney stone formation, and should therefore be avoided. Drinking plenty of liquids is recommended to reduce new stone formation. You will be monitored weekly for the occurrence of side effects and other potential difficulties associated with this treatment study.

The use of antiepileptic (anti-seizure) drugs (i.e., Topamax) during pregnancy has the potential for causing serous birth defects in offspring (children). If you become pregnant during this study, you will be discontinued from the study and the medication will be stopped after a tapering (lowering of the medication dose) period, if necessary. If you are pregnant, you cannot take part in this study. We will perform a pregnancy test at screening (beginning of study), and weeks 3, 5, 7, 9, 11 & 13 to make sure that you are not pregnant. You should use an effective method of birth control (i.e., spermicide + barrier, pill, or hormonal implants)

while you are taking part in the study. The interaction of Topamax and oral contraceptives (the pill) may render them less effective. Thus, if you are taking the oral contraceptive pill, you should consider utilizing an additional acceptable method of contraception. If you think you might be pregnant at any time during the study, you must tell us, and a pregnancy test will be performed. Your physician will explain the potential hazards (dangers) which may affect the fetus and possible alternatives, which may include ending the pregnancy.

There is the possibility of bruising, swelling and/or infection at the site of the venipuncture which is needed for the collection of blood samples.

Due to the possibility of dizziness or drowsiness, you will be cautioned regarding operating a vehicle or operating heavy machinery while in this study. The interaction between Topamax and central nervous system depressants such as alcohol may lead to excessive sedation. If you experience excessive somnolence (sleepiness), sedation, dizziness, or ataxia (problems with coordination), you should contact your study doctor immediately.

There is the possibility of feeling distressed or uncomfortable which can occur as part of the clinical interview, or completing self-report ratings and questionnaires. While the possibility of such events is low, our research staff is trained to look for these situations and provide a needed level of support.

Not all of the effects of Topamax are known at this time. There may be risks involved which are currently unforeseeable. You will be notified if significant new findings become known that may affect your willingness to continue participating in the study.

Benefits

At screening, you will receive a medical and psychiatric evaluation at no cost. You will also receive study medication and BBCET for 12-weeks at no cost. Your participation in this study may lead to the development of new treatments for alcohol dependence. We do not guarantee that you will benefit from taking part in this study.

Alternatives

You do not have to participate in this research study. Your alternatives would be

to be referred to another alcohol treatment program in the San Antonio area. In the San Antonio area, there are numerous medically (e.g., inpatient, day hospital) and psychologically (e.g., Alcoholics and Narcotics Anonymous) focused alternative treatment programs for the treatment of alcohol dependence and abuse. If you do not wish to enroll in the alcohol treatment research program described in this document, then we will refer you to these alternative programs.

You will be paid for your participation in the 12-week study at a rate of $20/week for a total of $240, if you complete all of the study visits. You will only be paid for the visits that you complete. You will also be paid for parking and/or transportation to the clinic for study visits.

Everything we learn about you in the study will be confidential. If we publish the results of the study in a scientific magazine or book, we will not identify you in any way. This study is being supported by Ortho-McNeil Pharmaceutical. The results of the study will be given to Ortho-McNeil Pharmaceutical, the company that makes Topamax. The Food and Drug Administration of the U.S. Government and Ortho-McNeil Pharmaceutical may need to see your records which identify you as a subject in this study, usually to confirm study participation.

Your decision to take part in the study is voluntary. You are free to choose not to take part in the study or to stop taking part at any time. If you choose not to take part or decide to stop at any time, it will not affect your future medical care at the University of Texas Health Science Center at San Antonio. We will tell you about any significant new findings which develop during the course of this research which may relate to your willingness to continue taking part.

Your participation in this study may be discontinued without your consent if: you do not follow the study procedures, you develop any serious medical or psychiatric complications, or you break any of the clinic guidelines for appropriate conduct. If this occurs, you must understand that it is important to notify your doctor so that he may plan for your

continuing medical care. This may involve gradually decreasing the dosage of Topamax, and visits to the clinic until you are no longer taking Topamax.

If you are injured as a result of the research procedures, medical care will be provided. You will be responsible for all charges. We are not able to give you money if you are injured. You or your insurance company will have to pay for all costs of care related to any injuries resulting from your participation in this research study. Insurance companies and Medicare may not pay for the costs of some research studies like this one. If your insurance company does not cover the costs of care, then you will have to pay these costs. You have the right to ask what it will cost you to take part in this study or to have other treatments.

If you have questions now, please do not hesitate to ask us. If you have additional questions later or you wish to report a medical problem which may be related to this study, Dr. Bankole Johnson can be reached at (210) 567-5475. If he is not available, Dr. Bordnick may be reached at (210) 567-5475. If you encounter any problems or need assistance after clinic hours, please page our Clinic Director (Dr. Alison Jones) at (210) 235-2435 or Rene' Beauregard, M.Ed. at (210) 230-4263 who will assist you. To use the pager, you need a touch tone (push button) telephone. Dial the pager number as you would any phone number. When you hear 3 short high-pitched beeps, dial the number where you want the doctor to call you back. Push the # button, hang up and wait for the doctor to return your call. The University of Texas Health Science Center at San Antonio committee that reviews research on human subjects (Institutional Review Board) will answer any questions about your rights as a research subject (210-567-2351).

We will give you a signed copy of this form to keep.

YOUR SIGNATURE INDICATES THAT YOU HAVE DECIDED TO TAKE PART IN THIS RESEARCH STUDY AND THAT YOU HAVE READ AND UNDERSTAND THE INFORMATION GIVEN ABOVE AND EXPLAINED TO YOU.

_____ _____

Signature of Subject Signature of Witness

_____/_____ _____

Date/Time Signature of Investigator

_____

Signature of Person Enrolling Subject

(If other than Investigator)