Estimating Models of Learning in Individual Decision making with an Application to Youth Smoking

Brett Matsumoto

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics.

Chapel Hill
2015

Approved by:

Donna Gilleskie

Peter Arcidiacono

Clement Joubert

Brian McManus

Helen Tauchen

ABSTRACT

BRETT MATSUMOTO: Estimating Models of Learning in Individual Decision making with an Application to Youth Smoking.
(Under the direction of Donna Gilleskie)

In the first chapter of my dissertation, I examine the dynamics of youth smoking behavior using a model of rational addiction with learning. Individuals in the model face uncertainty regarding the parameters that determine their utility from smoking. Through experimentation, individuals learn about how much they enjoy smoking cigarettes as well as the effects of reinforcement, tolerance, and withdrawal. The addition of learning to the dynamic optimization problem of adolescents provides an explanation for the experimentation of the non-smoker. I estimate the parameters of the model using data from the National Longitudinal Survey of Youth 1997 and compare the overall fit of the model to the model without learning. The estimated model is also used to analyze the effect of cigarette taxes and anti-smoking policies. I find that the model with learning is better able to fit the observed data and that cigarette taxes are not only effective in reducing the level of youth smoking, but can even increase welfare for some individuals.

In the second chapter (with Jonathan James), we show how the conditional choice probability (CCP) estimation procedure of Arcidiacono and Miller (2011) can be extended to feasibly estimate structural learning models. Although the focus of the paper is the specific application to learning models, the procedure could be used to estimate any model with continuous unobserved heterogeneity. Monte-Carlo simulations show that the CCP method can provide significant computational savings relative to Simulated Maximum Likelihood.

In the third chapter (with Forrest Spence), we investigate whether an individual's subjective price beliefs reflect the empirical distribution of prices and whether an individual learns about features of the price distribution through experience in the market. We use data on subjective price beliefs from a survey of 1,224 college students, and find that inexperienced individuals tend to expect online

prices to be higher than what is observed empirically. However, consumers with more experience in the marketplace generally have more accurate beliefs about the price distribution, which is consistent with learning.

ACKNOWLEDGMENTS

I have benefitted greatly from the advice and support of many people. I would like to thank my advisor Donna Gilleskie and committee members Peter Arcidiacono, Clement Joubert, Brian McManus, and Helen Tauchen. I would also like to thank Wayne-Roy Gayle, Juan Pantano, Tiago Pires, Steven Stern, and the UNC Applied Microeconomics Workshop for helpful comments and suggestions. Additionally, conversations and discussions with Forrest Spence, Matt Harris, and many other graduate students in the Economics Department were instrumental in shaping my research agenda.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

EXPLAINING YOUTH SMOKING INITIATION IN THE CONTEXT OF A RATIONAL ADDITION MODEL WITH LEARNING

1.1 Introduction

Despite its historically low level in the U.S., cigarette smoking remains a major public health concern. The Surgeon General estimates that tobacco use causes approximately 480,000 deaths per year in the United States and is estimated to cause between $289-332.5 billion in economic costs (USDHHS 2013).[1] Tobacco use is the leading preventable cause of death, yet people continue to smoke despite the high level of public awareness of its adverse health effects. Because cigarettes are addictive, it may be easier to discourage smoking initiation than to encourage smoking cessation. Also, cigarette manufacturers have historically targeted their advertisements to young people in the hopes of cultivating lifelong customers. Among adults who become daily smokers, approximately 90 percent smoke for the first time before age 18 (USDHHS 2012). For these reasons, policy interventions aimed at reducing the level of smoking in the population often target young people.

The decision to engage in a harmful addictive behavior, such as smoking, seemingly presents a problem for standard economic models. Consuming a harmful addictive substance would be an irrational act for a forward-looking utility-maximizing agent. The Rational Addiction (RA) model of Becker and Murphy (1988) shows that consumption of an addictive substance can be explained using the standard economic framework. Their explanation of addictive behavior centers around the concept that past utilization of addictive goods impacts current utility from consumption of these goods. A major criticism of the Becker and Murphy model is the implication that individuals are always acting optimally, so addicts do not regret their decision to consume the addictive good. In

[1] Economic costs include direct medical costs in addition to the lost productivity attributable to smoking related illnesses. Estimates are for the years 2009-2012.

their model addiction is not a problem or even an undesirable outcome, so there is no place for policy intervention to treat or prevent addiction. Empirical evidence suggests that many individuals regret their decision to smoke. Approximately 70% of adult smokers wish to quit smoking entirely and over half have attempted to quit smoking in the past year (NHIS, 2010).

Another limitation of the RA model as a model of youth smoking behavior is that it treats smoking initiation as exogenous. In this paper, I extend the RA model so that it is better able to explain the individual's smoking initiation decision. Specifically, I relax the assumption of perfect information in the RA model by incorporating learning about one's preferences. The parameters that determine the utility one receives from smoking are initially unknown, but the individual has beliefs about their true value. As an individual experiments with smoking, he receives utility signals and updates his beliefs. The addition of uncertainty and learning to the optimization problem of adolescents provides an explanation for the experimentation of the non-smoker and allows for the possibility that an individual who starts to smoke may later regret that decision. Therefore, policies that prevent an individual from experimenting with cigarettes may be welfare improving as the individual would be prevented from making a decision he may later regret.

The main purpose of this paper is to quantify the effectiveness of anti-smoking policies and to evaluate the resulting impact on individual welfare. In order to do this, I recover the policy-invariant utility function parameters of a rational addiction model with learning by fitting a dynamic discrete choice model of optimal smoking decision making to the observed data. As the first attempt to estimate the structural parameters of a rational addiction model with learning about preferences, this research allows for empirical testing of the perfect information assumption in the RA model (i.e., the assumption that individuals know their utility function parameters). I estimate the model parameters using the National Longitudinal Survey of Youth 1997 (NLSY97).

Estimation of the parameters of a dynamic discrete choice model is generally computationally intensive as each iteration over the parameter space requires re-solving the dynamic optimization problem. The inclusion of uncertainty and learning over multiple parameters further complicates estimation of the model. To circumvent these computational issues, I use the Expectation Maximization (EM) algorithm in conjunction with Conditional Choice Probability (CCP) estimation and

Monte Carlo simulation to estimate the model parameters. The estimation procedure provides a significant computational advantage, which allows for the estimation of a more complex model than is feasible using full-solution techniques.

Preliminary estimation results demonstrate that allowing for uncertainty and learning in a dynamic model of youth smoking significantly improves the overall fit of the model. Results from counterfactual policy simulations suggest that policies that impact individuals' initial beliefs about their utility function parameters are effective in reducing youth smoking. Taxes are also shown to be effective in reducing the level of smoking. The estimated model predicts that a doubling of the price of cigarettes would reduce the prevalence of youth smoking by 12.3% and adult smoking by 12.6%. An increase in the legal purchasing age from 18 to 19 years old would decrease youth smoking by 21.7%. However, there would be no effect on adult smokers as the higher legal purchasing age would only cause a delay in smoking initiation. The results of the welfare analysis show that increasing cigarette taxes would only lead to a relatively small loss in total welfare as the welfare gains to keeping those who would later regret the decision to smoke from starting to smoke offset the loss of welfare from smokers having to pay a higher price for cigarettes.

The remainder of the paper proceeds as follows: Section 2 reviews the related literature. Section 3 presents the model. Section 4 discusses the data. Section 5 develops the estimation routine. The estimation results are presented in section 6, and section 7 concludes.

1.2 Related Literature

Becker and Murphy (1988) developed the RA model to show that seemingly irrational behavior could be explained using a standard economic framework of a forward-looking utility-maximizing agent. The model's welfare implications have caused many to abandon the general framework of the RA model and to develop "irrational" models to explain the time inconsistency of addictive behavior. These alternative theoretical models generally feature dual-states of the world or individuals with dual-selves.[2] Addiction results when an individual is in an addictive state of the world or if the behavior of the individual is being controlled by the self that is more prone to addiction.

[2]Papers that use the dual-state approach include Winston (1980) and Bernheim and Rangel (2004). Papers that use the dual-self approach include Thaler and Shefrin (1981) and Benabou and Tirole (2004).

Other models of the consumption of addictive goods generate time-inconsistent behavior by deviating from the standard assumptions regarding how future utility is discounted. The simplest deviation is the myopic model. A myopic individual completely discounts future utility and only considers the current period's utility when making decisions. Other deviations from the standard assumptions regarding time preferences include an endogenous discount factor (Orphanides and Zervos 1998) or hyperbolic discounting (Gruber and Koszegi 2001). Finally, Orphanides and Zervos (1995) argue that the problem with the RA model is not the assumption of a rational, forward-looking agent but the assumption of perfect information. An individual in their model can be one of two types (addict or not an addict). The individual learns which type he is if he consumes the addictive good. The model estimated in this paper is an extension of the theoretical model proposed by Orphanides and Zervos (1995).

The RA model assumes that individuals are forward-looking, and there have been many studies that attempt to test the validity of this assumption empirically in the context of consumer demand for an addictive good. The evidence is generally consistent with forward looking behavior (Becker, Grossman, and Murphy 1994; Chaloupka 1991).[3] One of the limitations of the empirical addiction literature is that papers primarily attempt to compare the rational addiction model to the myopic model. No work (of which the author is aware) has been done to estimate alternative models or to empirically test the other assumptions of the RA model. Most of the literature involves reduced form estimation, but a few papers have estimated the structural parameters of an addiction model (Arcidiacono, Sieg, and Sloan 2007; Choo 2000; Gordon and Sun 2009; Darden 2011).[4]

Much of the analysis in the economics literature of policy interventions on youth smoking has focused on cigarette taxes. The rational addiction framework implies that individuals who are not currently consuming the addictive good should be more responsive to changes in the price of that good than current users. Many studies have found a significant effect of taxes on smoking initiation. Some studies, however, have found that cigarette taxes have little to no significant effect on youth

[3] See Chaloupka and Warner (2000) for a thorough summary of the empirical literature.

[4] There is a learning component to the life-cycle model of Darden (2011), but the learning is over the health effects of smoking. Individuals are assumed to know their preferences (i.e., utility function parameters).

smoking initiation (DeCicca, Kenkel, and Mathios 2002; DeCicca, Kenkel, Mathios, Shin, and Lim 2008; Emery, White, and Pierce 2001). Importantly, some of the studies in this literature find that nonsmokers are more price sensitive than smokers while also controlling for unobserved heterogeneity (Fletcher, Deb, and Sindelar 2009; Gilleskie and Strumpf 2005). Finally, some studies have found that taxes merely delay smoking initiation rather than prevent people from becoming smokers (Glied 2002). There have been fewer papers that examine the effect of other anti-smoking policies on youth smoking and the results have been mixed (Tworek, Yamaguchi, Dloska, Emery, Barker, Giovino, O'Malley, and Chaloupka 2010).[5]

One of the main applications of learning models in economics is in the area of consumer learning from experience goods (Erdem and Keane 1996; Ackerberg 2003).[6] These models estimate the learning process involved when consumers purchase unfamiliar goods. The consumer learns about the utility he receives from consuming these goods and updates his beliefs each time the good is consumed. This paper fits into the structural learning literature because the utility that the individual receives from consuming an addictive good is initially unknown and is learned over time if the individual consumes the addictive good. This paper extends the standard models used by incorporating the unique features of consuming an addictive good.

1.3 Model

This section sets up the individual's decision problem regarding optimal smoking behavior. An individual receives utility from consuming cigarettes as well as the consumption of other goods. In order to incorporate the features of consuming an addictive good, the individual's utility in the current period also depends on past levels of smoking in a manner consistent with the scientific literature on addiction (Laviolette and van der Kooy 2004; Nestler and Aghajanain 1997). Past consumption of the addictive good affects current utility through *reinforcement*, which occurs when the marginal utility of smoking is increasing in the level of past smoking. As the body becomes accustomed to consuming an addictive substance, larger quantities of the substance must be consumed to achieve

[5]For an overview of the effectiveness of anti-smoking legislation in general, see Goel and Nelson (2006).

[6]See Ching, Erdem, and Keane (2011) for an overview of the empirical economic applications of learning models.

a similar effect. This physical transition is referred to as developing *tolerance*. Habitual use of an addictive good also generates physical dependence. As a result, the individual experiences adverse effects from attempting to lower the level of consumption of the addictive good. This transition may result in a *withdrawal* effect. Withdrawal is modeled as an asymmetric adjustment cost, i.e. a cost associated with decreasing the amount consumed.[7] These effects are parameterized in the model (ρ, τ, and ω for reinforcement, tolerance, and withdrawal respectively), the magnitude of these effects depends on the level of past smoking, and these parameters vary across individuals. For certain combinations of these individual specific parameter values, the combined effect of reinforcement, tolerance, and withdrawal generates *adjacent complementarity* in the consumption of cigarettes. Adjacent complementarity, which Becker and Murphy (1988) use as the defining characteristic of addiction, occurs when current consumption of a good is increasing in past consumption.

1.3.1 Utility

Each year, individual n makes an annual smoking decision and chooses his level of smoking from a discrete set of alternatives, $a_j \in \{a_1, a_2, \ldots, a_J\}$, which reflect the average daily cigarette consumption during the year. The decision not to smoke is represented by the level of smoking a_1. The price of a single cigarette in period t is denoted p_t. The addictive stock is denoted as S_{nt} and is defined as the level of smoking in the prior year.[8] The contemporaneous utility associated with alternative $j > 1$ for individual n at time t if the individual did not smoke in the previous period ($S_{nt} = 0$) is:

$$u_{nt}^{j} = \big(\alpha_n + \xi_j X_{nt}\big) z(a_j) - \gamma_n p_t a_j + \epsilon_n^j \tag{1.1}$$

where ϵ is a vector of independent and identically distributed alternative-specific preference shocks that follow a Generalized Extreme Value (GEV) distribution. The utility from smoking depends

[7]This approach of explicitly modeling withdrawal effects as asymmetric adjustment costs to achieve adjacent complementarity in a rational addiction model was developed by Suranovic, Goldfarb, and Leonard (1999).

[8]This definition of the addictive stock implies full depreciation which is justified by the frequency of the smoking decision. Future versions of this paper will test whether the parameter estimates of the model change if this assumption is relaxed.

upon the individual-specific match parameter α_n, demographic variables (X_{nt}), the level of smoking through the function $z(a)$ (explained below), and the expenditure on smoking (which depends on both the price of cigarettes and level of smoking). The parameter γ_n measures the individual's sensitivity to the price of cigarettes and is a function of age, work status, and income.[9] Additionally, the individual's demographic variables affect utility by imposing additional costs or benefits on different levels of smoking. If a variable only affects the utility of smoking versus not smoking and does not affect the decision of how much to smoke conditional on smoking, then the coefficient ξ_j will be constant for $j > 1$. Variables in this category include the individual's race or religion. These variables may affect the social acceptance of smoking within the individual's culture. Variables that potentially affect utility differently for different levels of smoking could include whether the individual is under 18 years of age or whether the individual has older siblings. These variables were shown to be significant in the smoking decision of young people in Gilleskie and Strumpf (2005).

If the individual has a positive level of smoking stock (i.e., $S_{nt} > 0$) then the utility for alternative $j > 1$ is:

$$u_{nt}^j = \big(\alpha_n + \rho_n g(S_{nt}) + \xi_j X_{nt}\big) z(a_j) - \tau_n S_{nt} - \omega_n q(a_j, S_{nt}) \mathbf{1}[a_j < S_{nt}] - \gamma_n p_t a_j + \epsilon_{nt}^j \quad (1.2)$$

The addictive stock affects the marginal utility of smoking through the reinforcement, tolerance, and withdrawal terms. The reinforcement effect, $\rho g(S)$, increases the marginal utility of smoking for every positive level of smoking. The tolerance effect, τS, enters current period utility for positive levels of past and current consumption and decreases the utility associated with each positive level of smoking. The adjustment cost or withdrawal cost, $\omega q(a, S)$, only enters the current period's utility when the individual reduces his consumption from one period to the next. The utility of not smoking ($j = 1$) is normalized to only include the withdrawal term (if $S_{nt} > 0$) and the preference shock. The functions $z, g,$ and q have the following properties:

1. $z'(a) > 0$, $z''(a) < 0$, $\lim_{a \to 0} z'(a) < \infty$

[9]The utility from consuming one's entire income in other goods is normalized to zero.

2. $g(0) = 0$, $g'(S_{nt}) > 0$, $g''(S_{nt}) < 0$

3. $q(a_j, S_{nt}) \geq 0$ for all $a_j \leq S_{nt}$ and $q(a_j, S_{nt}) = 0$ if $a_j = S_{nt}$

The assumptions on the function z allow for a corner solution since the marginal utility from smoking is finite when the individual chooses not to smoke. The function q, which is a component of the withdrawal effect, is also assumed to be increasing in the size of the decrease in smoking from one period to the next. The functions g and q allow the reinforcement and withdrawal effects to be nonlinear.[10] The Estimation section discusses the specific functional forms used. The individual's smoking preference parameters are $\theta_n = (\alpha_n \ \rho_n \ \tau_n \ \omega_n)'$. The parameter α_n determines the individual's match quality for smoking. The parameters ρ_n, τ_n, and ω_n correspond to the effects of reinforcement, tolerance and withdrawal, respectively. The parameters in θ_n vary across individuals and are jointly normally distributed in the population: $\theta_n \sim N(\bar{\theta}, \Sigma)$.

1.3.2 Timing

The individual does not initially know the value of his smoking preference parameters (θ_n). He makes an annual smoking decision based on his beliefs about the parameters. At the start of the period, the individual observes prices, government tobacco policies, demographic variables (X), and the alternative specific preference shock. Then, the individual chooses a level of smoking and receives a utility signal. The individual uses this signal to update his beliefs at the end of the period.

An individual who has never smoked before the current period faces a sequential smoking decision within the period, where he first decides whether to experiment with smoking before making a smoking consumption decision for the year. The consumer learning literature generally finds that learning about match quality occurs relatively quickly. Since it would not take a full year to learn the match quality parameter α, an individual who has never smoked must first decide whether to experiment with smoking. Let a^E denote the level of consumption associated with experimentation. If he chooses to experiment, he learns his true value of α and proceeds to make a smoking decision for the rest of the period. If he chooses not to experiment, his smoking consumption for the period is zero and he will face the experimentation decision again in the next period. In periods after the individual

[10]The tolerance term could also be allowed to be nonlinear in S.

experiments, the only decision is about annual smoking consumption. The utility of experimenting is:

$$u^E_{nt} = (\alpha_n + \xi^E X_{nt}) z(a^E) - \gamma p_t a^E + \epsilon^E_{nt} \tag{1.3}$$

The utility shock for experimenting is assumed to be from a Type I Extreme Value distribution. For the sequential decision, individuals observe the preference shock for experimenting at the start of the period but do not observe the preference shock for the smoking decision until after they experiment.

1.3.3 Beliefs and Learning over the Utility Function Parameters

The individual's initial prior beliefs are denoted as $\theta_{n,0} \sim N(m_{n,0}, \Sigma_{n,0})$. Assuming Rational Expectations, the mean and variance of the individual's initial prior beliefs equal the population mean and variance of θ.[11]

The individual updates his beliefs according to a Bayesian learning process based on the signals received. After experimenting, the individual learns his true value of α. Without loss of generality, assume that the individual first experiments with the addictive good in period 0. The initial prior for the period 0 consumption decision is the initial prior distribution conditional on the realized value of α. Let $m_{n,0|\alpha_n}$ and $\Sigma_{n,0|\alpha_n}$ denote the mean and covariance matrix of the initial prior distribution conditional on $\alpha = \alpha_n$. This conditional distribution becomes the initial prior distribution for the subsequent learning over the parameters ρ, τ, and ω.

In every period that an individual chooses to smoke, he receives utility signals about the value of the reinforcement and tolerance parameters. If the individual reduces his level of smoking in period t from the level in period $t-1$, he receives a signal for the withdrawal parameter. For the level of smoking a_j and past smoking $\{S_{nl}\}_{l=0}^{t}$, the signals are as follows:

$$\delta_{nt} = \begin{cases} (\rho_n + \lambda_{nt})\mathbf{1}[a_j > 0] & \lambda_{nt} \sim i.i.d.\ N(0, \frac{\sigma^2_\lambda}{a_j(1+g(S_{nt}))}) \\ (\tau_n + \psi_{nt})\mathbf{1}[a_j > 0] & \psi_{nt} \sim i.i.d.\ N(0, \frac{\sigma^2_\psi}{1+S_{nt}}) \\ (\omega_n + \eta_{nt})\mathbf{1}[a_j < S_{nt}] & \eta_{nt} \sim i.i.d.\ N(0, \frac{\sigma^2_\eta}{S_{nt}-a_j}) \end{cases} \tag{1.4}$$

[11] Some restriction on the initial prior beliefs is required for identification. It may be possible to introduce heterogeneity in the initial priors by allowing the parameters of the initial prior beliefs to vary by observable characteristics.

The variation in the observed signal around its true value is assumed to be uncorrelated with the other parameters. The accuracy of the reinforcement signal is proportional to the quantity consumed as well as the level of past consumption, which implies that individuals face a trade-off between the speed of learning and the risk of becoming addicted. The accuracy of the tolerance signal is greater for higher levels of past consumption, and the accuracy of the withdrawal signal increases with larger decreases in consumption. The individual uses this utility signal to update his beliefs about his true parameters. I assume that the individual is able to distinguish between the signals if multiple signals are received in a given period and that the signal noises are uncorrelated (conditional on a_j and S_{nt}).

The individual's posterior beliefs at the end of period t after choosing a level of smoking equal to a_j (i.e., the individual's beliefs after receiving the signals associated with the smoking decision) are:

$$\theta_{n,t+1|\alpha} \sim N(m_{n,t+1|\alpha}, \Sigma_{n,t+1|\alpha}) \tag{1.5}$$

where

$$m_{n,t+1|\alpha} = \Sigma^{-1}_{n,t+1|\alpha}(\Sigma^{-1}_{n,t+1|\alpha} m_{nt|\alpha} + \Phi^{-1}_{nt}\delta_{nt}) \tag{1.6}$$

$$\Sigma_{n,t+1|\alpha} = (\Sigma^{-1}_{nt|\alpha} + \Phi^{-1}_{nt} B_{nt})^{-1} \tag{1.7}$$

$$\Phi^{-1}_{nt} = \begin{pmatrix} \frac{a_j(1+g(S_{nt}))}{\sigma^2_\lambda} & 0 & 0 \\ 0 & \frac{1+S_{nt}}{\sigma^2_\psi} & 0 \\ 0 & 0 & \frac{S_{nt}-a_j}{\sigma^2_\eta} \end{pmatrix} \tag{1.8}$$

$$B_{nt} = \begin{pmatrix} \mathbf{1}[a_j > 0] & 0 & 0 \\ 0 & \mathbf{1}[a_j > 0] & 0 \\ 0 & 0 & \mathbf{1}[a_j < S_{nt}] \end{pmatrix} \tag{1.9}$$

Equations (1.6) and (1.7) are the updating equations for the mean and variance of the individual's beliefs. The updated mean is a weighted average of the prior mean and the signal, where the weights are the precision (inverse of the variance) of the prior and the signal. Φ is a diagonal matrix of the signal precision, and B is a diagonal matrix with indicators for a given signal being received. As the individual receives more signals, the precision of his beliefs increases. Since the signals are unbiased, the individual's beliefs converge to the true parameter values.

Note that even though the signal noises are uncorrelated, the learning process for each parameter is not independent of the learning process for the other parameters. Since the parameters are correlated in the population and the population covariance matrix is the variance of the individual's initial prior beliefs, there is correlation in the learning process among the parameters. Even if the individual never receives a withdrawal signal, his beliefs about the value of his withdrawal parameter will change as he receives more information about the value of his other parameters.

1.3.4 Expectation of Future Prices, Policies, and State Variables

There are two components of the retail price of cigarettes: the manufacturer's price of the product and state and federal excise taxes. Determinants of the price of the product include the price of tobacco, production technology, labor costs, and other costs of production and distribution. Since surveyed individuals are not typically asked about their subjective expectations for future prices, some assumption must be made for how individuals forecast prices. One possible specification is to assume that the base component of the price follows a simple stochastic process (e.g., time trend with an AR(1) error). The justification for this specification is that individuals likely have some idea as to any time trend in the price as well as some realization that price shocks are persistent over time.

The other component of price, the excise tax, is much more difficult for the individual to forecast because it is determined by the political system. Specifying how individuals form expectations over other future tobacco policies presents a similar challenge. Estimates of the model presented in this work will impose the likely unrealistic assumption of perfect foresight.[12]

The endogenous state variables include the individual's beliefs and the addictive stock. The addictive stock is defined as the prior period's level of smoking, so the addictive stock evolves deterministically conditional on a particular smoking choice. The individual uses his current beliefs about smoking preferences to evaluate the different smoking alternatives, while taking into account the potential information that he will receive from each possible choice. The individual also has

[12]Other possibilities include assuming that the individual expects current tobacco taxes and policies to continue indefinitely or that individuals form expectations regarding the frequency and magnitude of excise tax changes based upon recent experience (i.e., a form of adaptive expectations).

perfect foresight regarding the observed exogenous state variables in X.[13]

1.3.5 The Individual's Problem

Each period, the individual chooses a level of smoking that maximizes his expected discounted lifetime utility given his beliefs and the value of the other state variables. The individual evaluates his expected discounted lifetime utility using backwards recursion. Let T denote the final period the individual is observed in the data, and let d_{nt}^j be an indicator variable that equals one if the individual selects alternative j in period t. Then the value function in period T is:

$$V_{nT}(S_{nT}, m_{nT}, \Sigma_{nT}, X_{nT}) = \mathbf{E}\Big[\max_j \ d_{nT}^j \Big(u^j(\theta_{nT}, S_{nT}, X_{nT}) + \beta\mathbf{E}[V_{n,T+1}(S_{n,T+1}, m_{n,T+1}, \Sigma_{n,T+1}, X_{n,T+1}, H_{n,T+1})|S_{nT}, m_{nT}, \Sigma_{nT}, X_{nT}, d_{nT}^j = 1]\Big)\Big] \quad (1.10)$$

The continuation value function V_{T+1} contains an additional state variable H that contains the individual's cumulative smoking history (i.e., total number of years smoked at each level of smoking).[14] The cumulative smoking history affects the individual's utility later in life through potential adverse health effects of smoking. The expectation over the discounted future value term is taken with respect to the future state variables. Current period utility is the expected utility given the current period's prior beliefs. Since the parameters in θ enter the utility function linearly, the expected utility for the current period is just the utility evaluated using the mean of the individual's current prior. The value function for earlier periods can be defined recursively starting from the terminal period value function:

$$V_{nt}(S_{nt}, m_{nt}, \Sigma_{nt}, X_{nt}) = \mathbf{E}\Big[\max_j \ d_{nt}^j \Big(u^j(\theta_{nt}, S_{nt}, X_{nt}) + \beta\mathbf{E}[V_{n,t+1}(S_{n,t+1}, m_{n,t+1}, \Sigma_{n,t+1}, X_{n,t+1}) \mid S_{nt}, m_{nt}, \Sigma_{nt}, X_{nt}, d_{nt}^j = 1]\Big)\Big] \quad (1.11)$$

[13] For some variables, such as the individual's age, this assumption is not unrealistic.

[14] The state variable H is suppressed in the value functions of earlier periods to simplify notation. Although the cumulative history does not affect utility in earlier periods, it still impacts the individual's behavior by changing the discounted expected future lifetime utility in period T.

If the individual has never smoked prior to period t, the value function for the experimentation decision is defined as:

$$V^E_{nt}(m_{n,0},\Sigma_{n,0},X_{nt}) = \max\left\{u^E_{nt} + \mathbf{E}_\alpha[V_{nt}(m_{n,0|\alpha},\Sigma_{n,0|\alpha},0,X_{nt})]\ ,\ \beta\mathbf{E}[V^E_{n,t+1}(m_{n,0},\Sigma_{n,0},X_{n,t+1})]\right\} \quad (1.12)$$

The first term inside the max operator is the value from experimenting in the current period. This term includes the utility from experimenting plus the value of the consumption decision for the current period. The value of the consumption decision depends upon a particular realization of α, which is unknown at the time of the experimentation decision, so the expected value of the consumption decision is calculated by integrating over potential realizations of α. The second term inside the max operator is the value associated with not experimenting, which is the discounted expected future value of the next period's experimentation decision.

The individual's problem is to choose the optimal sequence of experimentation and consumption in order to maximize his discounted lifetime expected utility. In the first period, the individual's beliefs are the initial prior beliefs and the individual has no experience with smoking.

1.4 Data

The data used to estimate the structural parameters of the model are from the NLSY97. The first wave of the survey was conducted in 1997 and included 8,984 individuals who were born between 1980 and 1984 (age at first interview ranged from 12 to 18). Subsequent waves have been conducted annually and are ongoing. This paper uses the first 13 waves of the data (through the 2009 wave). There are several advantages of using this data set for the study of youth smoking initiation. First, the individuals in the data set are surveyed at a young age during which the decision to begin smoking is made. Second, the survey is conducted annually, which is generally the shortest interval between observations in large nationally-representative panel data sets. The learning process is better identified with annual observations as opposed to less frequent observations.[15] Finally, the questions related

[15]If the individuals are only observed infrequently, then it is likely that much of the uncertainty would be resolved after a relatively small number of observations. It would be difficult to identify the dynamic learning process if the

to smoking are asked every wave. I supplement the geocoded restricted use version of the NLSY97 data set with tobacco policy data by matching individuals with the tobacco policies in their state. Relevant policies for this study include the cigarette excise tax, restrictions on tobacco advertising, spending on anti-smoking policies, and indoor smoking bans.

1.4.1 Sample Selection and Attrition

In a dynamic structural model, missing choice data add additional complexity in estimation. If an individual is in the sample, leaves, and later re-enters the sample, then the estimation routine has to integrate over all possible sequences of choices in the missing periods to calculate the value of the state variable when the individual re-enters the sample. One alternative is to only estimate the model on individuals who are observed in each time period. Restricting the sample to individuals observed in every time period avoids the difficulties in estimation, but the resulting sample may no longer be representative of the population if attrition is non-random. Table 1.1 reports the proportion of individuals with a given number of missing waves. Only about 60% of the original sample (5,385 of the original 8,984 individuals) is observed in every wave. Approximately 11% of this sample has one missing observation, and an additional 10% have either two or three missing observations. The preliminary estimation sample only includes the individuals who are observed in every wave. An additional 598 individuals are excluded due to missing smoking, demographic, or geographic data. The preliminary estimation sample contains the 4,787 individuals observed in every wave with nonmissing data for the key variables.

Table 1.1: Individual Level Survey Participation

Total years missing	0	1	2	3	4	5	6	7+	Total
Frequency	5,385	1,011	582	378	330	254	219	825	8,984
Percent	59.94	11.25	6.48	4.21	3.67	2.83	2.44	9.18	100

econometrician only had a few observations per individual where uncertainty and learning mattered.

1.4.2 Data Summary and Construction of Key Variables

In the NLSY97, individuals are asked whether they have smoked since the previous interview. If the answer is yes, the individuals are asked about their smoking behavior over the month prior to the interview. Specifically, the question asks, "during the past 30 days, on how many days did you smoke a cigarette?" If the answer is greater than zero, the next question asks, "when you smoked a cigarette during the past 30 days, how many cigarettes did you usually smoke each day?" I construct a categorical smoking variable from the answers to these two questions. The total number of cigarettes smoked in the past month is simply the product of the answer to these two questions and is divided by 30 to give the average number of cigarettes smoked per day. The range of possible values for the average number of cigarettes smoked per day is divided into four intervals to create the discrete choice variable a_j. These intervals correspond to not smoking, light smoking (0-5 cigarettes per day), moderate smoking (5-15 cigarettes per day), and heavy smoking (more than 15 cigarettes per day).

Table 1.2: Categorical Smoking Statistics

Smoking Level	Range (cigarettes per day)	Frequency (in person years)	Percent	$\mathbf{E}[a\|a_j]$
None	$a_1 = 0$	$44,186$	69.65	0
Light	$0 < a_2 \leq 5$	$9,714$	16.50	1.63
Moderate	$5 < a_3 \leq 15$	$5,469$	9.21	10.50
Heavy	$15 < a_4$	$2,862$	4.64	23.51

Table 1.2 reports the range of each of the intervals, the number of observations (in person years) in each interval, and the mean of average cigarettes smoked per day conditional on being in the range of the interval. The distribution of the average cigarettes smoked per day is skewed to the right with the majority of the observations concentrated at the mass point of zero.

Table 1.3 reports the transition probabilities for the smoking categories. The transition probabilities illustrate several key features of the data. First, individuals increase their level of smoking gradually. Individuals are more likely to increase to the next highest level than they are to jump

Table 1.3: Cumulative Smoking Transition Probabilities

	Smoking level at t			
Smoking level at $t-1$	None	Light	Moderate	Heavy
None	0.900	0.081	0.014	0.005
	(36,748)	(3,321)	(562)	(193)
Light	0.297	0.537	0.141	0.026
	(2,679)	(4,848)	(1,269)	(233)
Moderate	0.097	0.175	0.575	0.154
	(482)	(871)	(2,862)	(767)
Heavy	0.065	0.059	0.249	0.627
	(169)	(155)	(650)	(1,635)

Note: frequencies in parentheses

Table 1.4: Under 18 Smoking Transition Probabilities

	Smoking level at t			
Smoking level at $t-1$	None	Light	Moderate	Heavy
None	0.882	0.099	0.015	0.004
Light	0.360	0.448	0.151	0.040
Moderate	0.116	0.146	0.517	0.221
Heavy	0.076	0.093	0.271	0.559

several levels. Also, for any given level of smoking, there is a high probability that individuals will transition to a lower level of smoking. For light and moderate levels of smoking, the probability that individuals decrease the amount they smoke is approximately 30%. For the heaviest smokers, this probability is almost 40%. The amount of decreases in the level of smoking observed in the data is difficult to reconcile with the standard RA model, but is consistent with the model of behavior that incorporates uncertainty and learning. Table 1.4 reports the transition probabilities for individuals under 18 years old. Relative to the full sample there is less persistence in smoking choices, with more movement (both upward and downward) between smoking categories. This is consistent with the learning model since it will take some experience before individuals are able to determine what level of smoking is optimal for their specific utility function parameters.

Table 1.5 presents summary statistics for smoking behavior and demographic variables in three of the early waves.[16] Over these waves, the proportion of individuals who currently smoke increases, however, it does fall in later waves. The other variables included in the table enter the individual's decision to smoke, either directly through the utility received from smoking or through the cost of smoking. The NLSY does not ask about parent's smoking behavior. Parental smoking behavior potentially enters the individual's smoking decision through the individual's beliefs as well as through the cost of smoking. Parental education and other parental characteristics could serve as a proxy for parent smoking behavior.

Figure 1.1 presents the proportion of individuals in each smoking category by age. The proportion of individuals choosing to smoke increases steadily during the teenage years, reaches a peak for individuals in their early 20s, and declines slightly as individuals progress through their 20s. The decline in smoking rates for individuals in their 20s is primarily due to a lower proportion of light smokers. The proportion of moderate and heavy smokers remains relatively constant after reaching a peak around the age of 20. Figure 1.2 presents the proportion of current smokers by gender and race. Blacks have a substantially lower rate of smoking compared to other ethnic groups, and females have a lower smoking rate than males.

[16]See the Data Appendix for summary statistics for all waves.

Table 1.5: Summary Statistics of Smoking and Demographic Variables in Select Years

	Year					
	1997		1999		2001	
Time-Varying Variables	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Ever smoked	0.363	0.481	0.517	0.500	0.598	0.490
Current smoker	0.156	0.363	0.265	0.441	0.315	0.464
Number of cigarettes per day	0.541	2.559	1.618	5.064	2.368	6.135
Age	14.23	1.474	16.82	1.432	18.88	1.430
Employed	0.447	0.497	0.530	0.499	0.708	0.455
Real annual income*	247.5	762.6	1,208	3,178	3,889	6,295
Income $>$ $20,000	0.000	0.014	0.010	0.098	0.028	0.164
Married	0.000	0.014	0.013	0.115	0.053	0.224
Any children in household	0.007	0.085	0.047	0.212	0.107	0.310
High School student	0.982	0.133	0.691	0.462	0.292	0.455
College student	0.000	0.020	0.128	0.334	0.303	0.460
High School graduate	0.001	0.035	0.240	0.427	0.605	0.489
Time-Invariant Variables						
Female	0.536	0.499				
Black	0.255	0.436				
Father's educ (years)	10.32	5.752				
Mother's educ (years)	11.79	4.242				

* In year 2000 dollars.

Figure 1.1: Smoking Choice Probabilities by Age

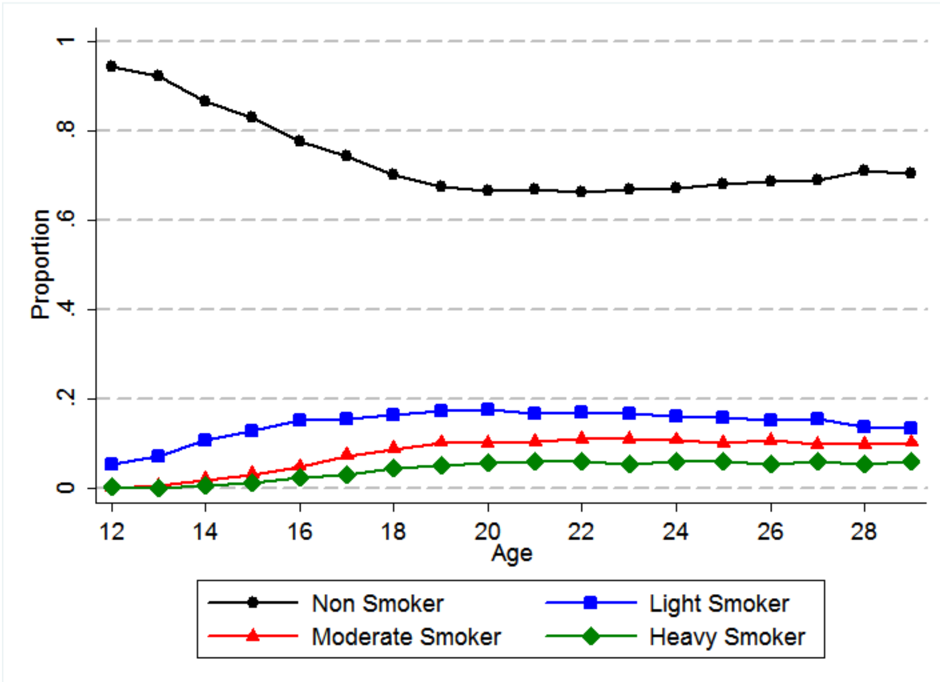

Figure 1.2: Gender and Racial Differences in Smoking Rates by Age

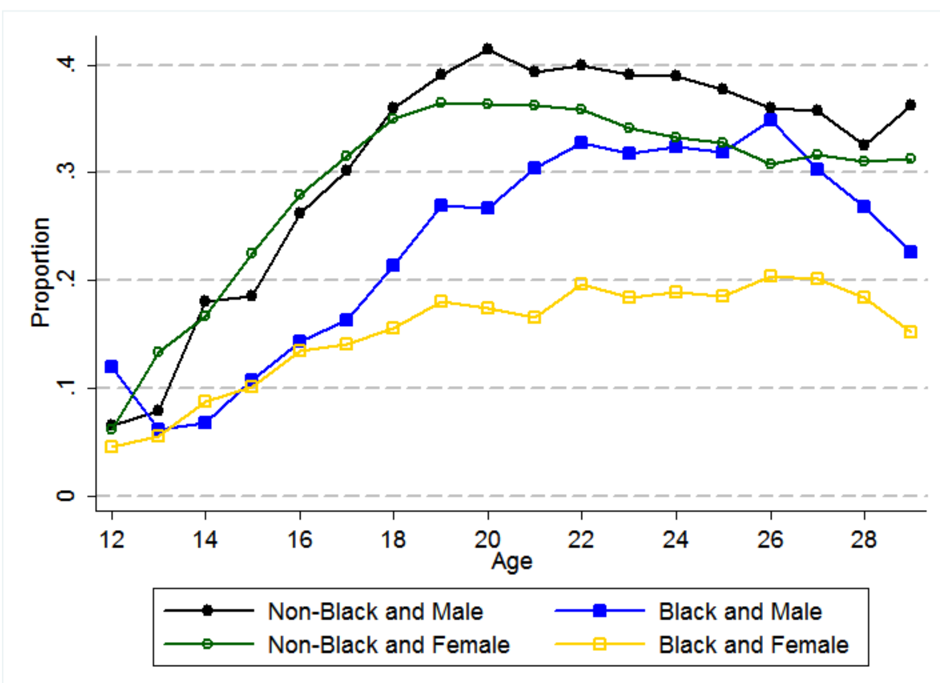

1.4.3 Cigarette Prices and State Excise Tax Data

The cigarette tax and price data used in this paper are from Orzechowski and Walker's *Tax Burden on Tobacco*. The price used is a sales weighted average of the premium brand cigarettes sold in a given year. Cigarettes are taxed at the federal and state level. In some instances they are also taxed at the county and municipal level. The federal cigarette tax in 2011 was $1.01 per pack. The tax rates vary considerably across states. In 2011, state cigarette taxes ranged from a low of $0.17 per pack in Missouri to a high of $4.24 in New York. At the start of the sample period in 1997, state cigarette taxes ranged from a low of $0.025 in Virginia to a high of $0.825 in Washington. Historically, the states with the lowest tax rates on tobacco are the tobacco-producing states of the southeast. From 1997-2011, only two states have had a constant tax rate, and most states have had multiple tax increases over the period. The variation in tax rates is largely responsible for the variation in the retail price of cigarettes across states. In 2011, the average retail price of cigarettes per pack ranged from $4.70 in Missouri to $10.29 in New York.

Table 1.6: Summary Statistics of State Tobacco Price and Taxes

Year	Real Price				Real Tax (State + Federal)			
	Mean	SD	Min	Max	Mean	SD	Min	Max
1997	2.265	0.327	1.796	3.305	0.633	0.217	0.284	1.143
1998	2.477	0.353	2.013	3.576	0.661	0.256	0.280	1.310
1999	3.200	0.361	2.698	4.304	0.670	0.271	0.274	1.282
2000	3.318	0.393	2.777	4.512	0.760	0.279	0.365	1.450
2001	3.500	0.370	3.035	4.458	0.752	0.291	0.355	1.410
2002	3.787	0.550	3.107	5.671	0.927	0.432	0.397	1.819
2003	3.843	0.567	3.157	5.452	1.041	0.453	0.388	2.284
2004	3.815	0.615	3.088	5.343	1.064	0.515	0.378	2.598
2005	3.847	0.643	3.095	5.292	1.157	0.528	0.406	2.513
2006	3.772	0.673	2.899	5.365	1.135	0.527	0.393	2.533
2007	3.883	0.652	2.906	5.520	1.191	0.518	0.382	2.462
2008	3.896	0.737	2.893	5.687	1.244	0.579	0.368	2.511
2009	4.711	0.846	3.406	6.458	1.856	0.628	0.867	3.588

Table 1.6 present the summary statistics across the 50 states and the District of Columbia of the real price of cigarettes as well as the real total tax. The average real price approximately doubles

over the sample period, and the amount of the average real tax increases by about three times. Over the time frame, the variability in both the prices and taxes across states increases. Most years the average real price increases due to increases in taxes. In years when there are no tax changes in a state, the real price of cigarettes falls as the nominal price increases less than inflation.

Figure 1.3 shows how real retail cigarette prices and taxes have changed over time in New York and North Carolina. Much of the price difference between these two states can be attributed to the difference in their cigarette taxes. Also, the increase in the price of cigarettes over time is driven by the increase in the tax rates. Other factors behind the increase in cigarette prices over this time period are the Tobacco Master Settlement Agreement in 1998, and the increase in the federal cigarette tax rate in 2009.[17] Figure 1.4 shows the distribution of state cigarette tax rates over time. At the beginning of the sample period, state cigarette taxes were relatively low. Over time, both the mean and variance of the state cigarette tax distribution increased.

1.4.4 State Level Tobacco Policy Data

In addition to tobacco excise taxes, there are many other policies that states can pursue to influence the level of youth smoking. Some of these policies enter into the individual's problem through the budget constraint by imposing non-monetary costs on obtaining tobacco. Some examples of policies that enter the individual's problem in this way are restrictions on the sale of tobacco to minors, bans on the sale of tobacco in vending machines, and restrictions on free samples of tobacco products. Another way for tobacco policies to influence behavior is through restrictions on tobacco consumption. The overall utility one receives from smoking will be less if there are restrictions on where and when one can smoke. Examples of restrictions on tobacco consumption are indoor smoking bans and smoke-free schools. Finally, some tobacco policies influence the individual's beliefs and expectations. In the context of this paper, these policies influence the individual's initial prior beliefs. Examples include restrictions on cigarette advertisements, funding of tobacco prevention

[17] In 1998, 46 states came to an agreement with the four largest cigarette manufacturers. The states agreed to drop their lawsuits against the tobacco companies, which sought compensation for the treatment of tobacco-related illnesses in the Medicaid system. In exchange, the tobacco companies agreed to a monetary settlement, restrictions on the marketing of tobacco products to young people, and the funding of a national anti-smoking organization. The tobacco companies raised the price of cigarettes by 45 cents per pack in response to the settlement to cover the payments to the states.

Figure 1.3: Real Cigarette Taxes and Prices in NY and NC (in year 2000 dollars)

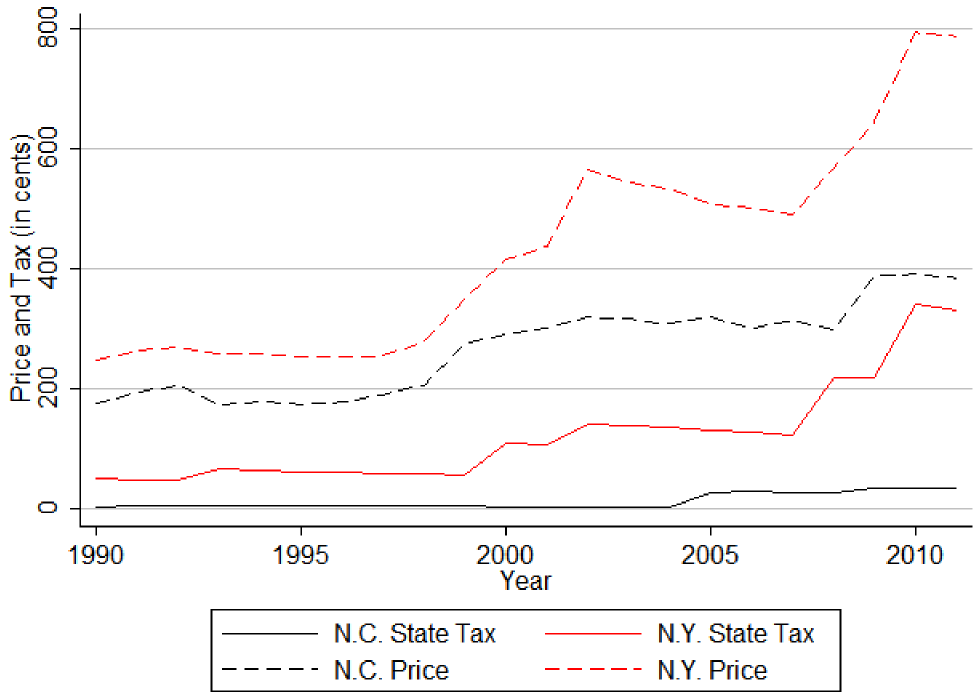

Figure 1.4: Distribution of Real State Cigarette Taxes by Year

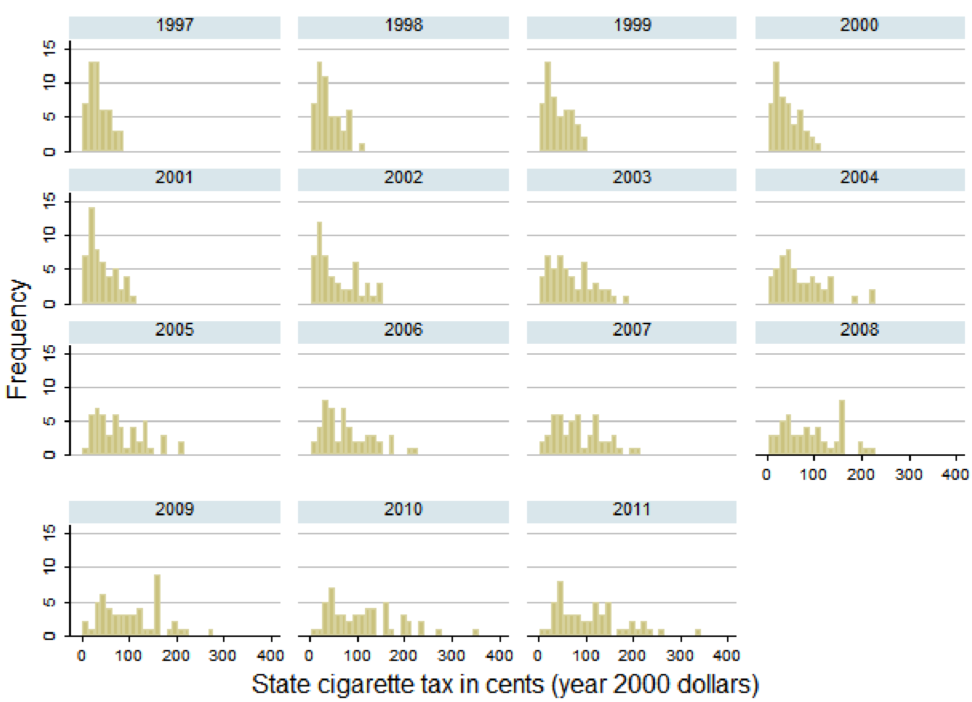

and education programs, and requiring tobacco education in schools. The data on state tobacco policies are from the Centers for Disease Control (CDC), the National Cancer Institute (NCI), and the Substance Abuse and Mental Health Services Administration (SAMHSA).

1.5 Estimation

1.5.1 Likelihood Function

Define the conditional value function for alternative j as the deterministic portion of flow utility from that alternative (i.e., utility minus the preference shock) plus the discounted expected future value of lifetime utility conditional on alternative j being chosen. Then, the conditional value function associated with alternative j in period t is given by:

$$\begin{aligned} v_{nt}^{j}(S_{nt}, \Gamma_{nt}, X_{nt}) &= \big(\alpha_n + \mathbf{E}_t[\rho_n|\Gamma_{nt}]g(S_{nt}) + \xi_j X_{nt}\big)z(a_j) - \mathbf{E}_t[\tau_n|\Gamma_{nt}]S_{nt}\mathbf{1}[a_j > 0] \\ &- \mathbf{E}_t[\omega_n|\Gamma_{nt}]q(a_j, S_{nt})\mathbf{1}[a_j < S_{nt}] - \gamma_n p_t a_j + \beta\mathbf{E}_t[V_{n,t+1}(S_{n,t+1}, \Gamma_{n,t+1}, X_{n,t+1})|d_{nt}^{j} = 1] \quad (1.13) \end{aligned}$$

where

$$V_{n,t+1}(S_{n,t+1}, \Gamma_{n,t+1}, X_{n,t+1}) = \mathbf{E}[\max_j v_{n,t+1}^{j}(S_{n,t+1}, \Gamma_{n,t+1}, X_{n,t+1}) + \epsilon_{n,t+1}^{j}] \quad (1.14)$$

The expectation over the future value term is taken with respect to the distribution of future beliefs, future demographic state variables, and future prices. The evaluation of current period utility depends upon the mean of the prior beliefs only. The variance of the prior does affect the expectation over future beliefs. The utility from not smoking is normalized to include the cost of withdrawal only, so $\xi_1 = 0$. The state variables are the level of smoking stock (i.e., last period's smoking decision) and the individual's beliefs, denoted by Γ, which include beliefs about parameter values and future prices.[18] I assume an i.i.d. type I extreme value (EV) preference shock.[19] The choice probabilities

[18]The price process has yet to be formally incorporated into the model, so the following estimation routine assumes perfect knowledge of future prices. The proposed estimation routine can be extended to estimate the parameters of a random price process.

[19]One of the major limitations of the multinomial logit model is the assumption that the shocks are uncorrelated over alternatives (i.e., the Independence of Irrelevant Alternatives (IIA) assumption). The use of random parameter, or mixed, logit can overcome the limitations of this assumption. In fact, mixed multinomial logit can approximate any discrete

after experimentation are given by:

$$P_{nt}^{j} = \frac{e^{v_{nt}^{j}}}{\sum_{k=1}^{J} e^{v_{nt}^{k}}} \quad \text{for } j = 1, \ldots, J \tag{1.15}$$

For individuals who have never smoked, they first choose whether or not to experiment, and then, conditional on experimenting, they decide the level of smoking. Let d_{nt}^{E} be an indicator variable that equals one if the individual experiments in period t. The conditional value of experimenting is:

$$v_{nt}^{E} = (\mathbf{E}[\alpha_n|\Gamma_{nt}] + \xi^{E} X_{nt}) z(a^{E}) - \gamma p_t a^{E} + \mathbf{E}_t[V_{nt}|d_{nt}^{E} = 1] \tag{1.16}$$

The conditional value function of not experimenting is simply the discounted expected maximum of the next period's value function conditional on not experimenting and not consuming any of the addictive good. The probability for experimenting, P_{nt}^{E}, is given by the Logistic cumulative distribution function. For an individual who has never smoked prior to period t, the behavior in period t is captured by the joint probability of experimenting and level of smoking ($P_{nt}^{E} P_{nt}^{j}$). The decision to experiment is made based on the individual's belief about his level of α, so P_{nt}^{E} is calculated based on an individual's beliefs. If he decides to experiment, he learns his true level of α, so P_{nt}^{j} is calculated using the individual's true value of α.

There are a total of N individuals, and each individual is observed for a total of $T+1$ periods. The likelihood of individual n making the sequence of choices $\{\cup_j \{d_{nt}^{j}\}, d_{nt}^{E}\}_{t=1}^{T}$ is:

$$L_n\Big(\gamma, \xi \mid \theta_n, \Gamma_{n,0}, \Lambda_n\Big) = \prod_{t=0}^{T} \Big((\prod_{j=1}^{J} P_{nt}^{j\, d_{nt}^{j}})^{A_{nt}} * \big[(1 - P_{nt}^{E})^{1-d_{nt}^{E}} \big(P_{nt}^{E} \prod_{j=1}^{J} P_{nt}^{j\, d_{nt}^{j}}\big)^{d_{nt}^{E}}\big]^{1-A_{nt}} \Big) \tag{1.17}$$

where A_{nt} is an indicator for the individual having ever smoked prior to period t. If the individual has smoked prior to period t (i.e., $A_{nt} = 1$), the individual makes a consumption decision. If the individual has not smoked prior to period t (i.e., $A_{nt} = 0$), then the individual makes a sequential

choice model derived from a random utility model to within any arbitrary degree of precision (McFadden and Train 2000).

experimentation and consumption decision. This individual likelihood is conditional on the individual's true addictive parameters (θ_n), the distribution of individual's initial prior beliefs ($\Gamma_{n,0}$), and a given sequence of signal noise draws ($\Lambda_n = \{\psi_{nt}, \lambda_{nt}, \eta_{nt}\}_{t=0}^{T}$). This formulation is equivalent to conditioning on the individual's beliefs at time t since the beliefs in time t are completely determined by the individual's initial prior, the sequence of signal noise, and the sequence of choices. Since the individual's true parameters and signal noise sequences are not observed by the researcher, the unconditional likelihood is calculated by integrating the conditional likelihood over the distribution of these unobserved variables:

$$L_n(\gamma, \xi, \sigma_\psi^2, \sigma_\lambda^2, \sigma_\eta^2, \bar{\theta}, \Sigma) = \int_\theta \int_\Lambda L_n\Big(\gamma, \xi \mid \theta_n, m_{n,0}, \Sigma_{n,0}, \Lambda_n\Big)\, dF(\Lambda | \sigma_\psi^2, \sigma_\lambda^2, \sigma_\eta^2)\, dF(\theta | \bar{\theta}, \Sigma) \quad (1.18)$$

and the full log-likelihood function is given by:

$$\mathscr{L}(\gamma, \xi, \sigma_\psi^2, \sigma_\lambda^2, \sigma_\eta^2, \bar{\theta}, \Sigma) = \sum_n log\Big(L_n(\gamma, \xi, \sigma_\psi^2, \sigma_\lambda^2, \sigma_\eta^2, \bar{\theta}, \Sigma)\Big) \quad (1.19)$$

The total dimensions of unobserved variables is $3 * T + 4$.[20] The integrals do not have a closed form solution, so they must be approximated numerically. The parameters to be estimated include the utility function parameters (γ, ξ), the mean and covariance matrix of the population distribution of the rational addiction parameters $(\bar{\theta}, \Sigma)$, and the variances of the signal noise distributions $(\sigma_\psi^2, \sigma_\lambda^2, \sigma_\eta^2)$.

1.5.2 Identification

The model parameters are identified through the observed sequences of smoking decisions. The parameters ξ and γ are identified through differences in smoking decisions between individuals with different observable characteristics. The price sensitivity parameter γ is identified by both cross-sectional variation and variation over time in the price of cigarettes. The utility from not smoking when the smoking stock is zero is normalized to zero. The parameter α affects the utility for each

[20]The dimension of the unobserved signals is likely to be less than 3*T since some of the signals are observed by the researcher. Based upon the sequence of actions, the researcher knows whether or not a signal is received in a given period.

level of smoking regardless of past smoking. The reinforcement parameter captures the effect of the interaction between the current level of smoking and the smoking stock. The tolerance parameter only depends on the smoking stock, so, for a given level of smoking stock, a change in the tolerance parameter only affects the probability of smoking versus not smoking. The reinforcement parameter affects the probability of smoking versus not smoking, but it also affects the probability of each level of smoking. The withdrawal parameter only affects the utility of a reduction in the level of smoking from one period to the next, so this parameter is identified by smokers who reduce their level of smoking or quit smoking entirely. The match, tolerance, reinforcement, withdrawal, and price sensitivity parameters do not vary across alternatives. Differences in utility for the different levels of smoking for these parameters are ultimately a result of the functional form assumptions.

The individual-specific parameters are not point identified for each individual. There is no way to estimate a specific value of these parameters for each individual. Also, since these parameters are continuous, a distributional assumption is required for the population distribution of parameters. Then, given that the conditional value function is defined over the support of the distribution of the unobserved continuous variables, the parameters of the population distribution (mean and covariance) are identified. The identification behind the learning process is driven by the fact that the valuation an individual attributes to each alternative depends upon the individual's current beliefs only and not the individual's true parameters. The individual's beliefs converge to the true parameters as the individual receives additional signals. Therefore, individuals with a lot of experience will behave according to their true parameter values. Also, if an individual knows his true parameter values, he can use the model to calculate an optimal consumption sequence. Differences between the optimal consumption sequence if the individual knows his true parameter values and the decisions of the individual when he is inexperienced are driven by the difference between the individual's beliefs and his true parameter values. The speed at which the individual's consumption sequence converges to the optimal consumption sequence with full knowledge identifies the speed of learning (i.e., the variance of the signals). Additional restrictions on the learning process are necessary for identification. These include restrictions on the initial prior beliefs (Rational Expectations), distributional assumptions for the beliefs and signals (both Normal), and Bayesian updating.

1.5.3 Estimation Procedure

There are several computational requirements that make estimation of the parameters of the model by Full Information Maximum Likelihood difficult. The main issue is that the evaluation of the log likelihood function requires integrating over the continuous distribution of population parameters and over all possible sequences of signal noise. Simulated maximum likelihood is one method that is used to overcome this problem. The unconditional likelihood function is approximated numerically by taking random draws from the distribution of the unobserved variable, evaluating the conditional likelihood, and taking the average of the conditional likelihoods over the draws. Evaluating the conditional likelihood, however, for a single draw still involves significant computation. The solution to the individual's problem requires integrating over future beliefs, which are multidimensional continuous variables. One way to reduce the computational burden of evaluating the value function is to use the Conditional Choice Probability (CCP) method of Hotz and Miller (1993).

Hotz and Miller (1993) show that when the preference shock has a GEV distribution, the future value term in the conditional value function can be expressed as a function of future flow utilities and conditional choice probabilities (CCPs). For certain classes of problems (e.g., optimal stopping problems), taking the difference in conditional value functions leads to the future value term only containing one period ahead flow utilities and CCPs. In other problems, the future value term associated with the difference in conditional value functions contains flow utilities and CCPs for a finite number of future periods. This property is called finite dependence, and it is a feature of the problem in this paper.[21] Standard CCP estimation involves estimating the CCPs in a first stage using the data and using the estimated CCPs to calculate the individual's value function. One limitation of the standard method is that it do not allow for unobserved heterogeneity. Arcidiacono and Miller (2011) develop a method of CCP estimation that allows for a finite distribution of unobserved heterogeneity by using the Expectation Maximization (EM) algorithm. The unobserved heterogeneity in this paper are the individual's beliefs and the individual's true parameter values, which are both continuous. James and Matsumoto (2013) extend the work of Arcidiacono and Miller (2011) to allow for a

[21] See the Estimation Appendix for the derivation of the CCP representation of the future value term.

continuous distribution of unobserved heterogeneity.

It can be shown that the values of the parameters that maximize the likelihood function (1.19) also maximize the following transformed likelihood function:[22]

$$\mathscr{L}(\gamma,\xi,\sigma_\psi^2,\sigma_\lambda^2,\sigma_\eta^2,\bar{\theta},\Sigma)=\sum_n\int_\theta\int_\Lambda \pi_n(\theta_n,\Lambda_n)\bigg(\sum_t (1-A_{nt})\big[(1-d_{nt}^E)log(1-P_{nt}^E)+d_{nt}^E log(P_{nt}^E)\big]+\sum_j d_{nt}^j log(P_{nt}^j)\bigg)\,d\Lambda\,d\theta \tag{1.20}$$

where π is the conditional probability that the parameter values are θ, θ_0, and Λ given the observed choices. This conditional probability is given by:

$$\pi_n(\theta_n,\Lambda_n)=\frac{f(\theta_n|\bar{\theta},\Sigma)f(\Lambda_n|\sigma_\psi^2,\sigma_\lambda^2,\sigma_\eta^2)\prod_t L_{nt}(\theta_n,\Lambda_n)}{\int_\theta\int_\Lambda\prod_t L_{nt}(\theta_n,\Lambda_n)f(\Lambda_n|\sigma_\psi^2,\sigma_\lambda^2,\sigma_\eta^2)f(\theta_n|\bar{\theta},\Sigma)\,d\Lambda\,d\theta} \tag{1.21}$$

The estimation routine in this paper used the likelihood function in equation 1.20. The procedure starts by taking M draws from the distribution of the unobserved variables for each individual as well as initial guesses for the values of the parameters and the CCPs. The estimation proceeds by using the EM algorithm, specifically a simulated EM algorithm (SEM). The EM algorithm is an iterative procedure that alternates between an expectation step (or E-step) and a maximization step (or M-step). The E-step updates the CCPs and π using the prior iteration values of the parameters and CCPs. The M-step updates the value of the parameters by maximizing the likelihood function using the updated CCPs and π. The estimation continues to iterate over these two steps until the parameter estimates converge. The use of the EM algorithm to incorporate unobserved heterogeneity has several advantages.[23] The most significant advantage is that the EM algorithm, or the SEM algorithm in the current context, reintroduces additive separability of the likelihood function. This property allows for sequential estimation of the likelihood function. In the current context, additive separability of the likelihood function allows for the parameters of the experimentation and consumption

[22]This is the expected conditional (on the unobserved variables) likelihood, where the expectation is taken with respect to the distribution of the unobserved variables conditional on the observed variables and the choices.

[23]See Arcidiacono and Jones (2003) for a full discussion.

decisions to be estimated separately. The estimation procedure is presented in greater detail in the Estimation Appendix.

1.5.4 Initial Conditions

The first period that individuals are observed in the NLSY97 is not the same as the initial period of the individual's optimization problem. That is, individuals may enter the estimation sample having already smoked. The values of the state variables in the initial wave of data depend on prior decisions and state variables that are not observed by the researcher. Some individuals have never smoked by the first wave. Others have smoked at some point prior to the first wave but are not observed to smoke in the first wave. Finally, some individuals are regular smokers at the first wave. The latter two groups present an initial conditions problem both in that the prior year's smoking is not observed in the first period and it is not observed how much they have learned. Individual's initial prior beliefs also present an initial conditions problem. I assume that individual initial priors are identical to the population distribution of the parameters (i.e., Rational Expectations).[24]

For individuals who have smoked prior to the first wave, the amount smoked in the period prior to the first wave is treated as discrete unobserved heterogeneity. The individual likelihood is calculated for each possible alternative in period $t = 0$. The probability that the individual selected alternative j in period $t = 0$ is:

$$P_{n,0}^{j} = \frac{1}{1 + exp(\xi_{IC}^{j} X_{n,0}^{IC})} \tag{1.22}$$

The individual likelihood is calculated by multiplying the likelihood conditional on selecting alternative j in period $t = 0$ by the probability $P_{n,0}^{j}$ and summing over the alternatives.

1.5.5 Functional Forms

The utility for the smoking level associated with alternative j, contains several modifying functions. The purpose of these functions is to allow for utility to be nonlinear in both the level of smoking and the level of past smoking. In order to estimate the parameters of the model, these generic functions must be replaced with specific functional forms. The function $z(a)$ incorporates

[24]In future work, I will attempt to parameterize the initial priors by allowing the mean of the initial priors (and perhaps the variance as well) to be functions of individual characteristics and state tobacco policies.

the standard utility function assumptions except that the marginal utility of smoking is positive for a level of smoking equal to zero. Also, the utility from not smoking is normalized to zero. The function $z(a)$ is assumed to take the following form: $z(a) = log(1 + a)$. The function that modifies the effect of reinforcement takes the following form: $g(S_t) = \sqrt{S_t}$. Finally, the function that modifies the withdrawal effect has the following form:

$$q(a_j, S_{nt}) = S_{nt} * \Big(1 - exp\Big(c * (a_j - S_{nt})\Big)\Big) \tag{1.23}$$

When the individual smokes the same amount as the prior period, $q = 0$. If the individual smokes less than the prior period, the withdrawal cost is positive. For a given level of last period smoking, the withdrawal cost decreases as the individual smokes more in the current period. This decrease occurs at an increasing rate. The parameter c affects the curvature of the function q as well as the maximum possible withdrawal cost. This parameter is initially fixed at a value of 0.15. Finally, the discount parameter β is set to 0.95.

1.6 Results

1.6.1 Parameter Estimates

This section presents the parameter estimates for the model. The estimation sample includes white males who are observed in every time period. The version of the model that is estimated differs from the model presented earlier in that the tolerance parameter τ is not estimated and set to zero. Table 1.7 presents the parameter estimates for the model with learning as well as the model without learning. The match parameter is negative for a large majority of the population. Even individuals with a negative match parameter could receive positive utility from smoking due to the effect of reinforcement. Individuals below the age of 18 experience a utility cost from smoking, which is likely due to their inability to purchase cigarettes legally. This cost is increasing in the level of smoking. The variance of the signals is significantly different from zero, which suggests that the learning component of the model is significant.

In order to test the importance of learning, I estimate a version of the model without learning. In the model without learning individuals are assumed to know the value of their parameters, but the

Table 1.7: Estimation Results, Population Distribution Parameter Estimates

Parameter	Description	Model with Learning	Model without Learning
$\bar{\alpha}$	Mean of match parameter	-3.567 (0.291)	-1.351 (0.042)
$\bar{\rho}$	Mean of reinforcement parameter	1.573 (0.105)	0.363 (0.011)
$\bar{\omega}$	Mean of withdrawal parameter	0.595 (0.021)	0.150 (0.014)
$Var(\alpha)$	Variance of match parameter	0.988 (0.100)	4.723 (0.221)
$Var(\rho)$	Variance of reinforcement parameter	0.290 (0.019)	0.167 (0.009)
$Var(\omega)$	Variance of withdrawal parameter	0.444 (0.019)	0.038 (0.004)
$Cov(\alpha, \rho)$	Covariance of match and reinforcement	-0.301 (0.245)	0.867 (0.031)
$Cov(\rho, \omega)$	Covariance of reinforcement and withdrawal	0.221 (0.012)	0.061 (0.006)
$Cov(\alpha, \omega)$	Covariance of match and withdrawal	-0.364 (0.023)	0.274 (0.032)
σ_λ	Standard deviation of reinforcement signal	0.912 (0.051)	-
σ_η	Standard deviation of withdrawal signal	1.010 (0.071)	-

Note: Standard Errors in parentheses

parameters vary across individuals.[25] Estimates of the parameters from the model without learning differ in important ways from the parameter estimates from the model with learning. The mean of the population distribution of the match parameter is larger in magnitude in the model without learning and has much higher variability in the population. The mean value of the reinforcement parameter is nearly the same in both models, and the mean of the withdrawal parameter is smaller in the model without learning. The population variance of the reinforcement and withdrawal parameters is smaller in the model without learning.

Table 1.8 presents the estimates of the coefficients on the observable variables. The first panel includes the estimates for the variables that enter the utility function as preference shifters. The next panel includes the variables that affect price sensitivity, and is followed by the parameters that affect the probability of different levels of prior unobserved consumption. The last panel includes the variables that enter the utility of experimentation. Note that there is no experimentation decision in the model without learning since individuals already know they value of the match parameter. Other than the age variables, the coefficients on observable characteristics tend to be relatively small

[25]The model without learning corresponds to a restricted version of the model with learning. Specifically, the model without learning is equivalent to the model with learning where the mean of the initial prior is set to the individual's true parameter value and the variance of the initial prior is set to zero.

Table 1.8: Estimation Results, Coefficients on Observable Variables

Variable	Model with Learning	Model without Learning
Preference Shifting Variables		
Years until age 18, light smoking	-0.075 (0.088)	-0.435 (0.052)
Years until age 18, moderate smoking	-0.229 (0.186)	-0.903 (0.176)
Years until age 18, heavy smoking	-0.053 (0.357)	-1.016 (0.337)
Years until age 18 squared, light smoking	-0.055 (0.018)	0.005 (0.012)
Years until age 18 squared, moderate smoking	-0.199 (0.044)	0.001 (0.039)
Years until age 18 squared, heavy smoking	-0.305 (0.090)	-0.013 (0.074)
Married, light smoking	-0.146 (0.134)	-0.052 (0.039)
Married, moderate smoking	-0.064 (0.166)	-0.127 (0.051)
Married, heavy smoking	-0.273 (0.191)	-0.041 (0.048)
Has children in household, light smoking	0.090 (0.142)	0.167 (0.039)
Has children in household, moderate smoking	-0.057 (0.187)	0.107 (0.051)
Has children in household, heavy smoking	0.027 (0.215)	-0.032 (0.049)
Price Sensitivity Variables		
$\bar{\gamma}$, Mean Price sensitivity	0.143 (0.111)	0.442 (0.025)
Under age 18	0.042 (0.031)	-0.053 (0.112)
Employed	0.076 (0.028)	-0.009 (0.018)
Income greater than $20k	-0.022 (0.107)	-0.024 (0.009)
Unobserved Prior Consumption Variables		
Constant, light smoking	0.050 (0.741)	5.913 (4.978)
Constant, moderate smoking	0.277 (1.359)	1.047 (1.189)
Constant, heavy smoking	-0.182 (1.654)	-0.543 (1.623)
Years since first smoked, light smoking	-0.068 (0.476)	-3.581 (5.137)
Years since first smoked, moderate smoking	1.040 (0.492)	-0.135 (0.213)
Years since first smoked, heavy smoking	-0.377 (0.919)	-0.153 (0.330)
Experimentation Variables		
Years until age 18	-0.475 (0.096)	-
Years until age 18 squared	0.128 (0.017)	-
Married	-0.155 (0.437)	-
Has children in household	0.066 (0.317)	-
a^E	0.949 (0.290)	-

in magnitude

The price sensitivity parameter is larger for the model without learning. The model without learning is limited in terms of explaining quitting (or any reduction is smoking). Since prices are increasing throughout the sample, the model without learning attributes any reduction in smoking to the increase in prices. Although an increase in the price of cigarettes is one reason why an individual would reduce his level of smoking, the model with learning allows for other potential reasons. As an individual experiments with smoking, he may discover that his true utility from smoking is less than he initially believed it to be. An individual is also able to learn about the withdrawal cost through reductions in smoking. So in the model with learning, reduction in smoking could be due to the increase in price, new information about the utility from smoking, and strategic reductions in smoking in order to learn about the withdrawal cost. By ignoring the mechanisms through which the learning process generates endogenous quitting (or reduction), the model without learning overstates the importance of price in explaining the observed level of quitting.

1.6.2 Model Fit

Table 1.9 presents the observed transition probabilities from the data as well as the transition probabilities from simulated outcomes generated using the model with learning and the estimated parameters. The model is able to fit the observed transition probabilities well. For smoking transitions for individuals under 18, the simulated data tends to overstate the persistence in smoking behavior, particularly for remaining a nonsmoker and a heavy smoker. The model is better able to fit the transition probabilities for individuals over 18 years old. Also, the simulated data tends to underestimate the probability of quitting, particularly for heavy smokers.

Figure 1.5 shows the proportion of individuals in each smoking category by age for both the observed and simulated data using the estimated model with learning. The simulated data closely match the observed age profile of smoking behavior. Figure 1.6 compares the proportion of individuals in each smoking category by age for the observed data and for simulated data using the model without learning. The model without learning does a relatively poor job in matching the observed data.

Table 1.9: Transition Probabilities, Observed and Simulated Data

Under 18 years old								
	Smoking level at t							
	Observed Data				Simulated Data			
Smoking level at $t-1$	None	Light	Moderate	Heavy	None	Light	Moderate	Heavy
None	86.81	10.47	1.83	0.88	90.26	8.23	0.90	0.62
Light	35.21	44.92	14.25	5.62	35.06	52.28	9.08	3.58
Moderate	15.70	10.74	49.59	23.97	9.37	16.97	50.32	23.34
Heavy	9.84	8.20	22.95	59.02	4.24	4.75	18.79	72.21
Over 18 years old								
	Smoking level at t							
	Observed Data				Simulated Data			
Smoking level at $t-1$	None	Light	Moderate	Heavy	None	Light	Moderate	Heavy
None	89.60	8.09	1.52	0.79	85.95	11.41	1.76	0.88
Light	27.12	57.88	12.24	2.76	32.64	56.23	9.31	1.82
Moderate	9.90	16.30	55.82	17.99	7.63	13.13	59.98	19.26
Heavy	6.30	5.09	21.26	67.35	3.66	3.22	18.88	74.23

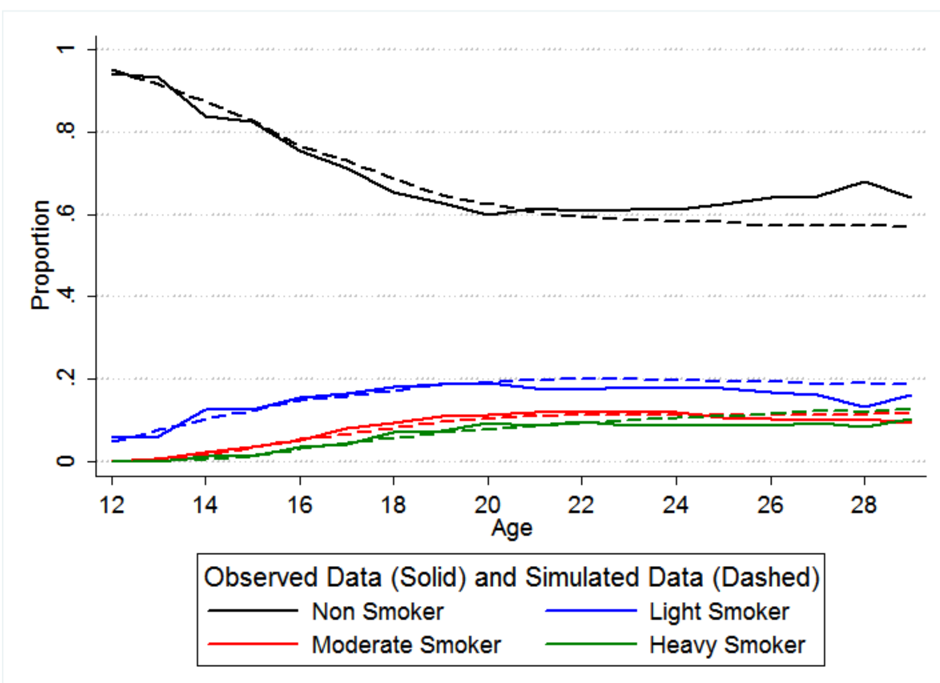

Figure 1.5: Smoking Rates by Age, Observed and Simulated Data from the Model with Learning

Figure 1.7 illustrates the model fit for the experimentation decision. This figure shows the probability by age that an individual who has not previously experimented with smoking will experiment as well as the cumulative probability that an individual has experimented at some point in the past. Up until age 19, the model is able to closely fit the observed experimentation probability. After age 20, the probability that an individual experiments drop quickly to between 1 and 2 percent, however the model predicts a much more gradual decline.

1.6.3 Policy Simulations

In this section, I use the parameter estimates from the model with learning to conduct policy counterfactual experiments. I consider policies that alter cigarette prices, beliefs about withdrawal, and the legal smoking age.[26]

Prices

The tobacco excise tax is a popular policy tool among policymakers and anti-smoking advocates to reduce the level of smoking. The specific policy experiment is doubling the price of cigarettes. Under the counterfactual policy, individuals are faced with a price of cigarettes that is two times what is observed in the data. This counterfactual measures the long run impact of a change in the price of tobacco. Figure 1.8 depicts the smoking rates by age for the baseline simulation and the simulated data under the counterfactual prices.

Doubling prices has a dramatic effect on the proportion of smokers. For individuals over 18 years old, the proportion of nonsmokers increases by about 5 percentage points as a result of the higher prices. The proportion of light smokers decreases by around 15 percent, and the proportion of moderate and heavy smokers decreases by around 10 percent.[27]

Beliefs

Prior to experimentation with cigarettes, a young individual has beliefs about his smoking preference parameters. The model imposes rational expectations for the initial beliefs. However, the

[26] The legal smoking age is the minimum age at which an individual can legally purchase tobacco products.

[27] One concern with this counterfactual is that the CCPs are estimated using the data, so they are only identified for values of the state variables that are observed in the data.

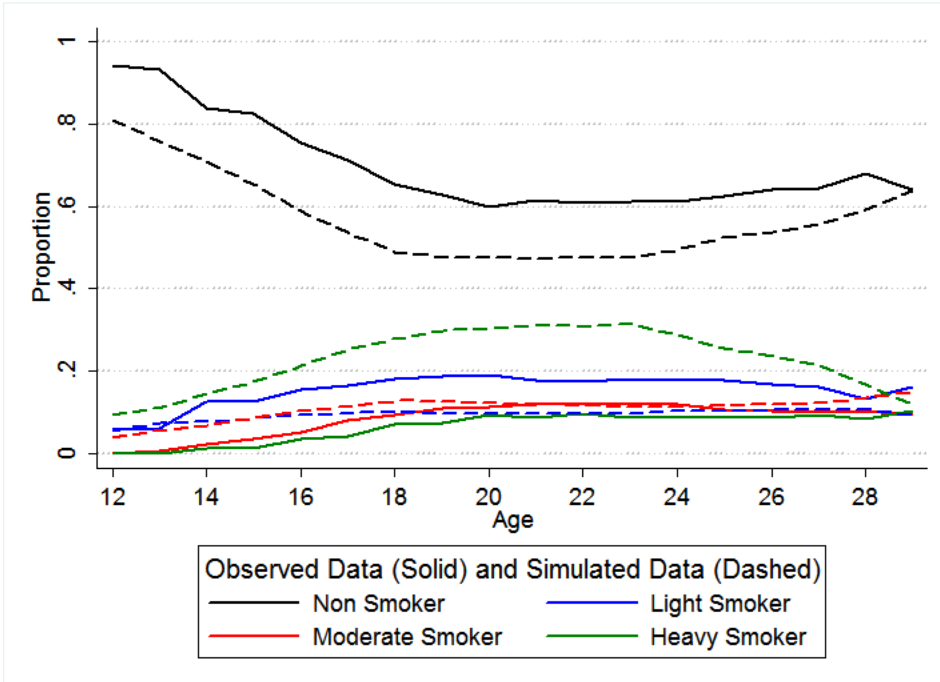

Figure 1.6: Smoking Rates by Age, Observed and Simulated Data from Model without Learning

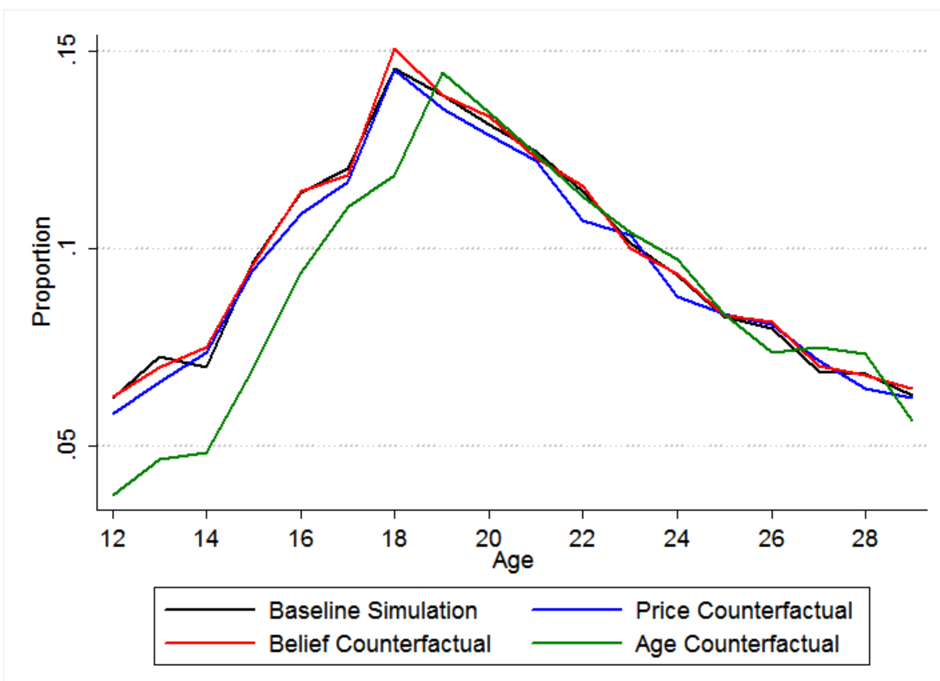

Figure 1.7: Experimentation Rates by Age, Observed and Simulated Data from Model with Learning

individual's initial beliefs are likely influenced by a number of factors and could potentially be influenced by anti-smoking policies. For example, advertisements that highlight the addictive nature of cigarettes and the difficulty of quitting smoking may affect the individual's beliefs about the value of the parameters that govern the effects of reinforcement, tolerance, and withdrawal.

In this counterfactual experiment, the mean initial prior for the withdrawal parameter is increased by one standard deviation of the population distribution. Only the initial belief about the withdrawal parameter changes. The actual distribution of the withdrawal parameter in the population is the same.

Increasing the mean value of the withdrawal parameter in one's initial prior beliefs causes a reduction in the overall level of smoking. Individuals are now less likely to experiment with smoking given the higher anticipated cost of quitting. The higher expected withdrawal cost results in a large reduction in the proportion of moderate and heavy smokers. The proportion of light smokers is slightly less relative to the baseline simulation. Some individuals who were light smokers under the baseline simulation decide not to smoke under the counterfactual. These individuals are offset by those who were moderate and heavy smokers under the baseline who remain light smokers for a longer period before transitioning to higher levels of smoking under the counterfactual. The increase in the expected withdrawal cost has the effect of extending the experimentation period.

Legal Smoking Age

Another possible policy tool that targets youth smoking is the minimum legal age to purchase tobacco. In this counterfactual experiment, the effect of increasing the minimum legal purchase age to 19 years old. Relative to the baseline simulation, increasing the purchase age is effective in reducing smoking among teenagers. However, increasing the legal purchase age only delays smoking rather than preventing it. The smoking rates converge to the baseline simulation for all smoking categories once individuals are able to legally purchase tobacco.

These counterfactual simulations confirm that increasing the price of cigarettes is an effective policy tool to reduce the prevalence of smoking. Changing the legal smoking age would have the effect of reducing youth smoking, but would likely have only a minimal impact in reducing smoking in the broader population. Policies that target an individual's initial prior beliefs about the utility of smoking could be very effective in reducing smoking. Specifically, increasing an individual's beliefs

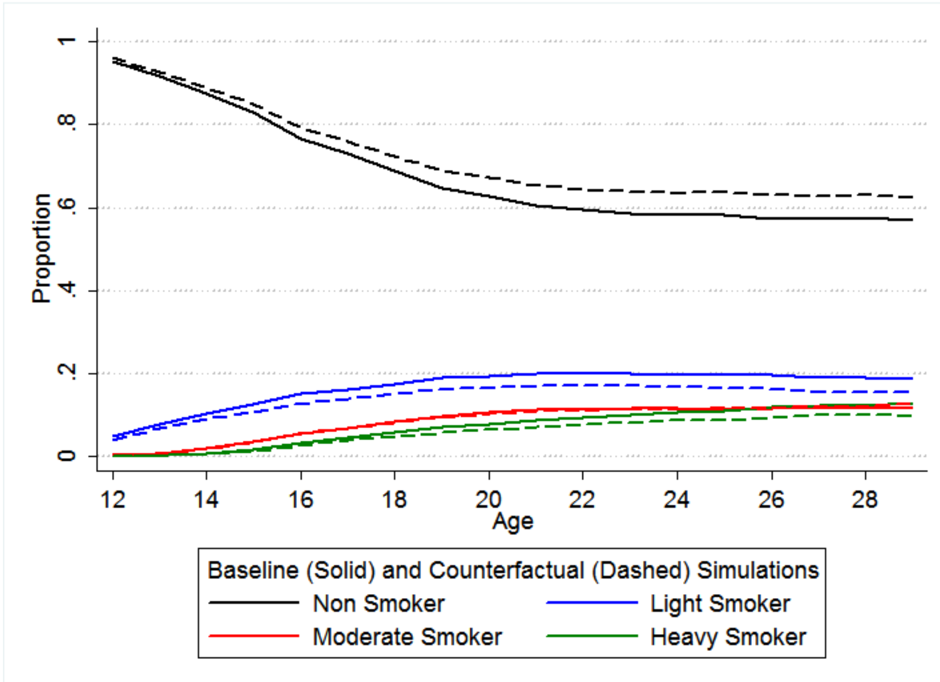

Figure 1.8: Smoking Rates by Age, Baseline Simulation and Price Counterfactual Data

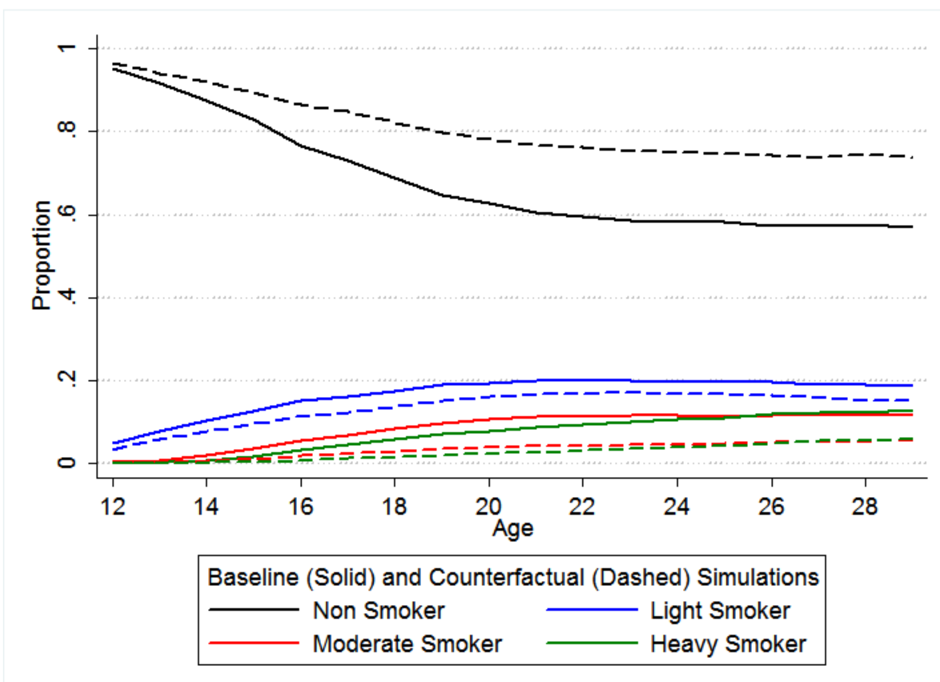

Figure 1.9: Smoking Rates by Age, Baseline Simulation and Beliefs Counterfactual Data

about the withdrawal cost would lead to a relatively large reduction in the probability that individual would become a heavy smoker.

Figure 1.11 shows the experimentation probabilities by age for the different counterfactual simulations. The belief counterfactual has no effect on the probability that an individual experiments. This is because an individual who experiments who does not continue to smoke will not experience withdrawal. Increasing the price only has a minor effect on the probability that an individual experiments. Even with the higher prices, the monetary cost of experimenting is minor. However, the individual would be less likely to continue smoking as a higher value of the match parameter would be required to offset the higher monetary cost of smoking. Increasing the smoking age simply delays experimentation.

1.6.4 Welfare Analysis

Thus far, the policy analysis has evaluated the effectiveness of alternative policies in reducing the level of smoking without taking into account the effect of the policies on an individual's welfare. One of the key advantages of the model with learning is the ability of the model to explain regret. In the standard RA model, policies that increase the cost of smoking will lower the welfare of every individual. In the model with learning, a policy that increases the cost of smoking may lower the welfare of some individuals, but it may increase the welfare of others. If an individual who would later regret the decision to become a smoker decided not to smoke because of the policy, that policy would increase his welfare (assuming the policy did not affect the utility from not smoking).

To evaluate the effect of the alternative policies on welfare, I calculate the expected lifetime utility (ELU) for each individual's simulated sequence of choices. This measure includes the deterministic portion of utility evaluated using the individual's true parameters.[28] Since different sequences of choices lead to a different state space in the terminal period, the individual's value function in the final period is added to the sequence of flow utilities.[29] The objective in constructing this measure of

[28]Note that expected lifetime utility refers to the ex-post deterministic flow utility up to the end of the sample period.

[29]Including the terminal period value function adds the value of information to the measure of welfare. An individual would not regret smoking just because the ex-post flow of utility from smoking was less than not smoking, as long as the value of information more than offsets the difference.

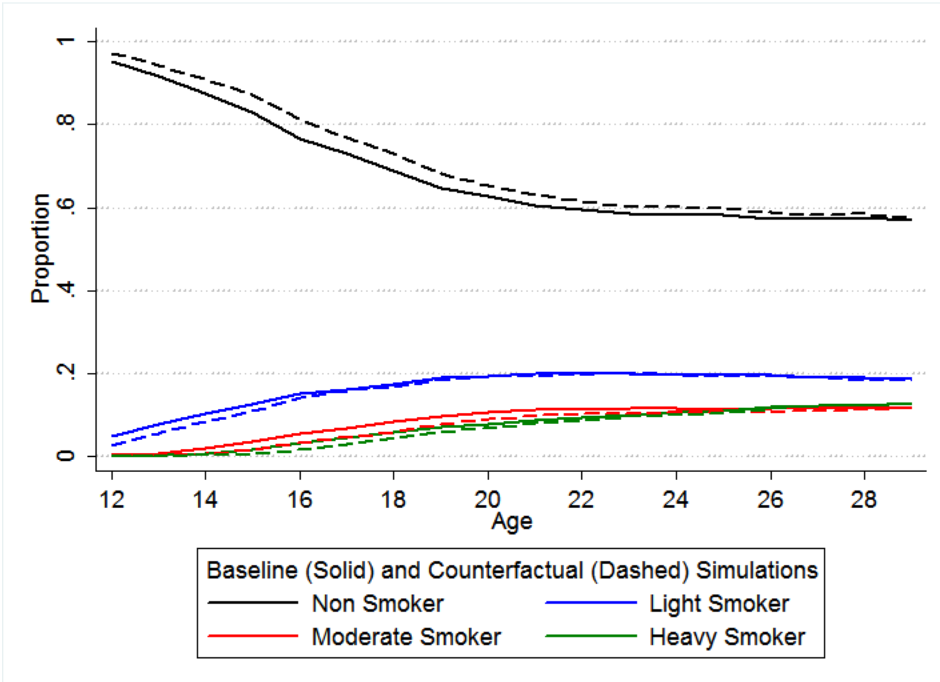

Figure 1.10: Smoking Rates by Age, Baseline Simulation and Smoking Age Counterfactual Data

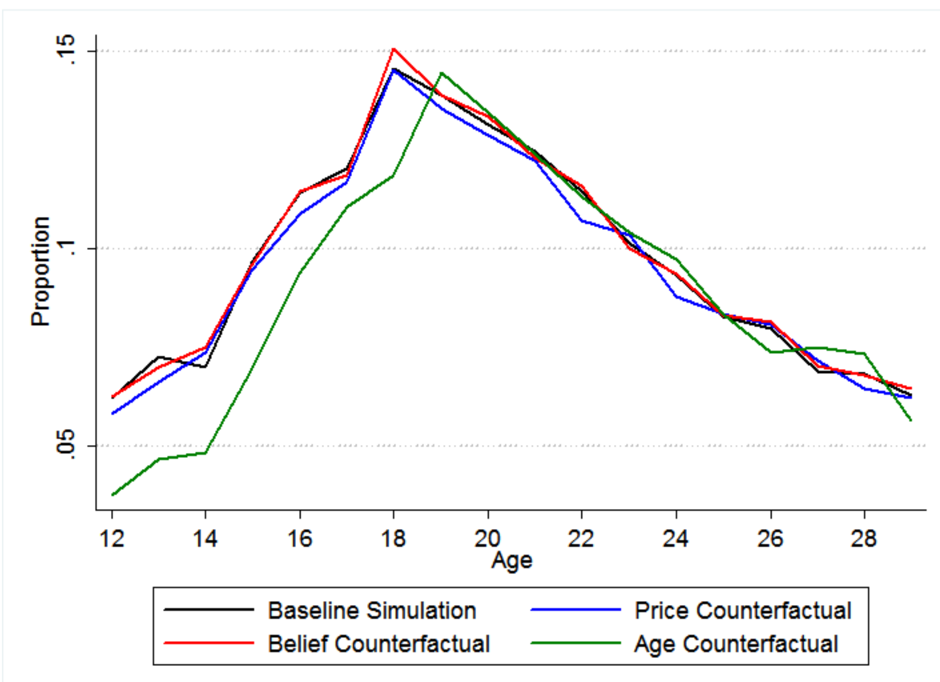

Figure 1.11: Experimentation Rates by Age, Baseline Simulation and Counterfactual Data

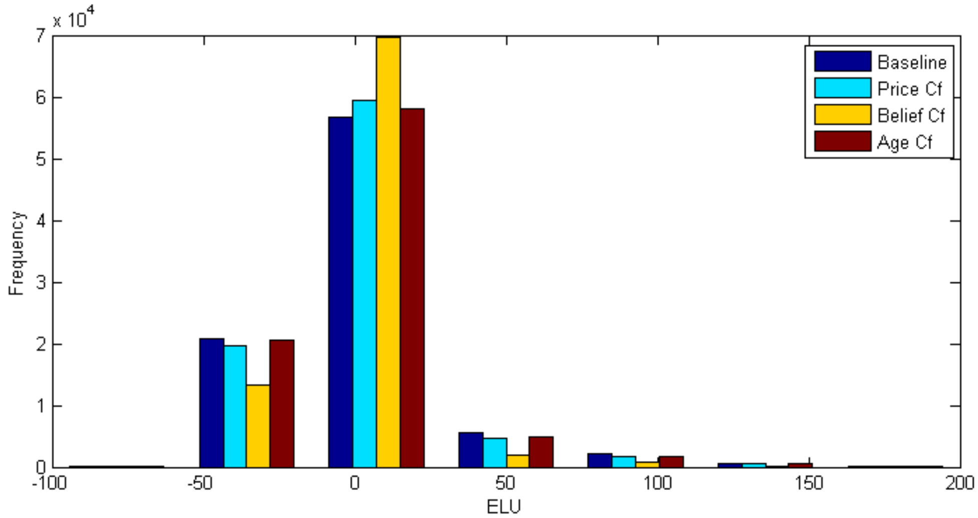

Figure 1.12: Distribution of Expected Lifetime Utility

welfare is to determine the effect of the alternative policies on the individual's ex-post welfare. The welfare measure is normalized with respect to the sequence of never smoking. The expected flow utility portion of the welfare measure is already normalized from the normalization of the utility function. The terminal period value function is normalized by subtracting the terminal period value function of the individual if he had never smoked.

Table 1.10 shows the summary statistics for the ELU measure of individual welfare for the baseline simulation as well as under the different counterfactuals. The ELU of an individual is zero for an individual who never smokes. The mean ELU for the baseline simulation is -0.660 and the median is -5.870. A vast majority of individuals have an ELU that is less than zero, indicating that they are ex-post worse off than if they had never smoked. In the price counterfactual, the proportion of individuals who never smoke increases. By increasing the proportion of individuals who never smoke, these policies decrease the proportion of individuals with a negative ELU, but the policies also decrease the proportion of individuals with a positive ELU. Individuals who enjoy smoking are negatively affected by the increase in the price of cigarettes. Increasing initial beliefs about the difficulty of quitting smoking decreases welfare on average as it decreases the proportion of individuals with a positive ELU and increases the proportion of individuals with a negative ELU.

Figure 1.12 is a histogram that shows the distributional effects of the different smoking policies.

Table 1.10: Welfare Analysis

	Mean ELU	Std ELU	Med ELU	Pr(ELU<0)	Pr(ELU=0)	Pr(ELU>0)
Baseline	-0.660	27.388	-5.870	0.722	0.055	0.223
Price Cf	-2.063	24.352	-5.686	0.739	0.062	0.199
Belief Cf	-4.191	16.828	-4.981	0.791	0.075	0.135
Age Cf	-1.981	25.042	-5.951	0.738	0.058	0.204

The first bar in each set represents the outcomes under the baseline simulation. The second, third, and fourth bars are the results for the price, belief, and age counterfactuals respectively. All of the policies reduce the welfare of individuals who enjoyed smoking under the baseline simulation as there are fewer individuals in each range of positive ELU. The price and belief counterfactuals also reduce the number of individuals in each range of negative ELU and increase the number of individuals with an ELU close to zero. The individuals with a large negative ELU include those who are "trapped" in their addiction. These are individuals who initially overestimated their true tolerance and underestimated their true withdrawal. They receive a negative utility from smoking but are not able to quit because they face a large withdrawal cost.

1.7 Conclusion

This research develops a model of rational addiction with learning in order to explain the smoking initiation decision of young people. Estimation of the structural parameters of the model requires significant computational resources, and is not computationally feasible using a full solution estimation routine unless significant restrictions are placed on the model. Therefore, this paper proposes the use of an alternative estimation routine. This estimation routine uses the EM algorithm and CCP estimation, which reduces the computational burden of estimating the structural parameters of the model.

Overall, the model is able to fit the data well. In particular, the model with learning fits the data significantly better than the model without learning. The estimated parameters of the model are used to conduct counterfactual policy experiments. Since an individual's decision to smoke depends upon his beliefs about his smoking preference parameters, policies that affect one's beliefs can have a significant impact on smoking behavior. Increasing individuals' beliefs about the difficulty of quitting smoking is effective at reducing the number of heavy smokers. Increasing cigarette prices is shown to be an effective policy tool to reduce youth smoking, although the model without learning overstates the importance of the price of cigarettes. An increase in the legal age to purchase cigarettes would lead to a decrease in the number of youth smokers, but it would only delay smoking initiation so adult smoking behavior would not be affected. The analysis of individual welfare supports the use of taxes as an anti-smoking policy tool. An increase in the price of cigarettes improves the ex-post level of utility for some individuals by discouraging those who would later regret the decision to smoke from ever experimenting with cigarettes, but hurts those who do enjoy smoking.

The results of this paper suggest several potential avenues of future research. First, the analysis performed considers the demand side of the market. Although the analysis in this paper demonstrates the importance of learning in explaining cigarette demand, the model would need to be extended to incorporate optimal firm behavior in order to better capture the general equilibrium effects of policy changes. An individual's initial beliefs are an important determinant of early smoking behavior, as changing the initial beliefs was shown to have a large effect on behavior in the counterfactual simulation. Additional exploration of how these initial beliefs are formed would be useful. In particular, to what degree are the individual's initial beliefs influenced by the smoking behavior of others (e.g., parents, siblings, peers) or by advertising (either pro- or anti-smoking). Finally, the importance of learning in explaining youth smoking behavior begs the question of how learning about cigarette smoking preferences may impact learning about preferences for consuming other addictive goods such as alcohol or illegal drugs. There are potential knowledge spillovers about the dynamic effects of consuming an addictive good (i.e., tolerance, reinforcement, and withdrawal), which may be correlated across different addictive goods for an individual.

CHAPTER 2

USING CONDITIONAL CHOICE PROBABILITIES TO ESTIMATE DYNAMIC DISCRETE CHOICE MODELS WITH CONTINUOUS UNOBSERVED HETEROGENEITY WITH AN APPLICATION TO LEARNING MODELS (CO-AUTHOR JONATHAN JAMES)

2.1 Introduction

This paper develops a method for estimating dynamic discrete choice models with continuous unobserved state variables using the conditional choice probability (CCP) method developed by Hotz and Miller (1993). The use of CCP estimation results in substantial computational savings for a range of dynamic discrete choice models. CCP estimation techniques were extended by Hotz, Miller, Sanders, and Smith (1994) to cover a larger range of discrete choice models and by Arcidiacono and Miller (2011) to incorporate a finite number of unobserved states. Allowing for unobserved heterogeneity overcomes one of the primary obstacles to the practical application of CCP methods. In this paper, we extend prior CCP estimation techniques to allow for a continuous distribution of unobserved heterogeneity. One area where this technique could potentially yield the greatest benefit is in the estimation of structural learning models. Even with recent computational advances, estimation of learning models is often infeasible in many applications unless strong restrictions are placed on the model.

In dynamic models, the current period value function not only depends on the current period utility, but also the discounted expected future utility. In recursively defined models with a continuous state variable, calculating the expected future value term requires integrating over potential realizations of the state variable for each future time period. Even when numerical methods are used to simulate these integrals, estimation of these dynamic models can become infeasible as the number of time periods increases or as the state space becomes richer. One way to deal with continuous variables is through discretization. Continuous state variables can also be incorporated in full solution estimation techniques using the method developed by Keane and Wolpin (1994). They use simulation techniques to calculate the value function at particular points (often randomly chosen) within the state space. When solving the individual's optimization problem, the individual's value function at other points in the state space is calculated through interpolation. When the continuous state variable is unobserved, an additional computational burden is added. In estimation, the likelihood function is formed by integrating over the unobserved variable. With time varying unobserved heterogeneity, the individual's problem, which requires integrating over future values of the unobserved variable, must be solved for all possible current period values of the continuous variable.

The use of CCPs can reduce the computational burden of estimating models with continuous unobserved heterogeneity by reducing the burden of solving the individual's problem. In discrete choice models, the choice probabilities are functions of the conditional value functions. Hotz and Miller (1993) showed that for certain discrete choice models, there is an inverse mapping, so the conditional value functions could be written as functions of the conditional choice probabilities. In particular cases, the conditional value functions can be written as functions of flow utilities and conditional choice probabilities for only a few periods into the future. With a continuous state variable, the conditional choice probability representation of the value function still requires integrating over future conditional choice probabilities. However, the conditional choice probability representation of the value function can reduce the number of integrals by reducing the number of future value terms that are required to evaluate the conditional value functions.

2.2 Simple Learning Framework

2.2.1 Model

Each period the individual chooses among J alternatives plus an outside option. The utility from option $j \in \{1, \ldots, J\}$ is:

$$u_{jt} = \alpha_j X_t + \mu_j + \epsilon_{jt} \tag{2.1}$$

The utility from the outside option, $j = 0$, is normalized to zero. The vector ϵ includes alternative- and time-specific preference shocks assumed to be i.i.d. Type I extreme value. The vector $\mu = (\mu_1, \ldots, \mu_J)$ consists of individual-specific tastes for each of the alternatives, and individuals do not know their true value of μ. The population distribution of μ is normal with mean $\bar{\mu}$ and variance Σ. Although individuals do not know their true value of μ, they have a belief about its value. Individuals learn about their true value of μ through experience.

The individual's prior belief in period t is normally distributed with mean m_t and variance Φ_t ($\mu_t \sim N(m_t, \Phi_t)$). After choosing alternative j, the individual receives the signal:[1]

$$\delta_{jt} = \mu_j + \xi_{jt} \tag{2.2}$$

$$\begin{pmatrix} \xi_1 \\ \vdots \\ \xi_J \end{pmatrix} \sim N(0, B) \tag{2.3}$$

$$B = \begin{pmatrix} \sigma^2_{\xi_1} & 0 & \cdots & 0 \\ 0 & \sigma^2_{\xi_2} & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma^2_{\xi_J} \end{pmatrix} \tag{2.4}$$

Let d_{jt} be an indicator variable that equals 1 when the individual chooses alternative j. Then the individual updates the mean and variance of his beliefs according to the following Bayesian updating equations:

$$m_{t+1} = \Phi_{t+1}(\Phi_t^{-1} m_t + B^{-1} D \delta_{jt}) \tag{2.5}$$

$$\Phi_{t+1} = [\Phi_t^{-1} + B^{-1} D D']^{-1} \tag{2.6}$$

$$D_t = \begin{pmatrix} d_{1,t} \\ \vdots \\ d_{J,t} \end{pmatrix} \tag{2.7}$$

The information that the individual uses when evaluating all of the possible alternatives is contained in a vector of state variables. These state variables are the mean and variance of the individual's prior

[1]There is no signal associated with the outside option.

beliefs, m and Φ, and the individual's observable characteristics, X. The conditional value function for alternative j is:

$$v_{jt}(m_t, \Phi_t, X_t) = \alpha_j X_t + \mathbf{E}[\mu_j] + \sum_{X'} \beta \mathbf{E}[V(X', m_{t+1}, \Phi_{t+1} | m_t, \Phi_t, D_t)] * Pr(X_{t+1} = X' | X_t, D_t) \quad (2.8)$$

Evaluating the future value term requires integrating over all possible values of the signal since different realizations of the signal will generate different beliefs in the future period. Also, the variable μ_j enters the utility function linearly, so the current period utility can simply be evaluated using the mean of the prior distribution ($\mathbf{E}[u_j(\mu_j)] = u_j(\mathbf{E}[\mu_j])$).[2] The variance of the prior distribution affects the likelihood of different realizations of the mean of the posterior distribution.

2.3 CCPs and Finite Dependence

In this section we show that the learning framework naturally generates a relatively simple expression for the future value term in the conditional value function. Calculation of the conditional value function for each alternative requires solving the full dynamic learning model. The conditional value function for alternative j is period t is given by 2.8. The expected future value term is calculated by integrating over possible future state variables, including future beliefs. The individual's mean prior in period $t+1$ depends upon his beliefs in period t as well as the realization of the signal. So integrating over possible future beliefs requires integrating over possible realizations of the signal. The variance of the future beliefs is a function of the variance of the current beliefs and the variance of the signal. Since the individual knows the variance of the signal, the individual can calculate the variance of future beliefs conditional on the choice. The variance of the prior beliefs transitions deterministically conditional on the choice and does not depend on the actual realization of the signals.

Consider the future value term in the conditional value function for alternative j. This term can be expressed as a function of a conditional choice probability and the conditional value function of

[2] If μ_j enters the utility function non-linearly, then calculating expected utility requires integrating over the prior distribution.

any of the alternatives:

$$\begin{aligned}
\mathbf{E}[V(m_{t+1},\Phi_{t+1}|m_t,\Phi_t,D_t)] &= \mathbf{E}_m[log(\sum_{j=0}^{J} exp(v_{j,t+1}(m_{t+1},\Phi_{t+1}))] \\
&= \mathbf{E}_m[v_{0,t+1} + log(1+\sum_{j=1}^{J} exp(v_{j,t+1} - v_{0,t+1}))] \\
&= \mathbf{E}_m[v_{0,t+1} - log(P_{0,t+1}(m_{t+1},\Phi_{t+1}))] \\
&= \mathbf{E}_m[u_{0,t+1} + \beta * \mathbf{E}[V(m_{t+2},\Phi_{t+2}|m_{t+1},\Phi_{t+1},D_{t+1})] \\
&\quad -log(P_{0,t+1})]
\end{aligned} \tag{2.9}$$

The CCP representation still contains a future value term for period $t+2$. The future value term associated with the conditional value function for the outside option has the following CCP representation:

$$\begin{aligned}
\mathbf{E}[V(m_{t+1},\Phi_{t+1}|m_t,\Phi_t,D_t)] &= \mathbf{E}_m[log(\sum_{j=0}^{J} exp(v_{j,t+1}(m_{t+1},\Phi_{t+1}))] \\
&= \mathbf{E}_m[v_{j,t+1} + log(1+\sum_{k\neq j} exp(v_{k,t+1} - v_{j,t+1}))] \\
&= \mathbf{E}_m[u_{j,t+1} + \beta * \mathbf{E}[V(m_{t+2},\Phi_{t+2}|m_{t+1},\Phi_{t+1},D_{t+1})] \\
&\quad -log(P_{j,t+1})]
\end{aligned} \tag{2.10}$$

When forming the choice probability in the likelihood function, what matters is the difference in conditional value functions. Arcidiacono and Miller (2011) show that it may be possible in a given problem to choose a particular form of the CCP representation such that the future value term k periods in the future cancels when taking the difference in conditional value functions. Finite dependence is the term they use to define the property of expressing the difference in conditional value functions in such a way to generate the cancellation of the future value terms. Given how beliefs transition in learning models, it is possible to generate finite dependence in beliefs.[3]

[3]Finite dependence must also hold for the other state variables along the choice sequence used to generate finite dependence in beliefs in order for the estimation procedure to be feasible.

The period $t+2$ value functions in equations (2.10) and (2.11) are equal to each other if the period $t+2$ state variables are the same after the sequence $\{j, 0\}$ and the sequence $\{0, j\}$ in expectation. For each of the sequences, the individual will receive one signal about the utility of alternative j. The variance of the priors will be the same after both sequences. The individual's expectation in period t of the mean of the priors in period $t+2$ will also be the same for each sequence. Prior to the realization of the signals, all that matters to the individual when forming expectations of future beliefs is which signals will be received, not the timing of the signals.

2.3.1 Likelihood Function

The probability of individual n selecting alternative j in period t conditional on the state variables is:

$$P^j_{nt}(m_{nt}, \Phi_{nt}, X_{nt}) = \frac{exp(v^j_{nt}(m_{nt}, \Phi_{nt}, X_{nt}))}{\sum_{k=0}^{J} exp(v^k_{nt}(m_{nt}, \Phi_{nt}, X_{nt}))} \tag{2.11}$$

The individual's contribution to the likelihood function conditional on his true parameter values and sequence of signal noises is:

$$L_n(\mu_n, \Delta_n) = \prod_t \prod_j P^j_{nt}(m_{nt}, \Phi_{nt}, X_{nt} | \mu_n, \Delta_n) \tag{2.12}$$

where $\Delta_n = \{\delta_{n,1}, \ldots, \delta_{nt}\}$. The unconditional likelihood function is obtained by integrating over the distribution of parameter values and signals.

$$L_n = \int_{\mu,\Delta} L_n(\mu_n, \Delta) f(\Delta, \mu | \bar{\mu}, \Sigma, \sigma_\xi) \; d\Delta \; d\mu \tag{2.13}$$

The full sample likelihood is calculated as the product over the individuals' likelihoods. The parameters that need to be estimated are the parameters of the population distribution of μ, the variance of the signal, and the coefficients of the observed state variables in the utility function.

2.4 The Estimation Algorithm

The model parameters can be estimated using Simulated Maximum Likelihood. The estimation procedure we propose is based on the EM Algorithm. First, we will present the estimation procedure, and then we will discuss the advantages of using the EM Algorithm over Simulated Maximum

Likelihood.

Denote the parameters to be estimated as θ. Then the parameters that maximize the log likelihood function will also maximize the following equation:

$$\mathscr{L}(\theta) = \sum_n \int_{\mu,\Delta} q(\mu_n, \Delta_n | D, X, \theta) log(L_n(\mu_n, \Delta_n) f(\mu_n, \Delta_n|\theta)) \; d\mu \; d\Delta \tag{2.14}$$

where

$$q(\mu_n, \Delta_n | D, X, \theta) = \frac{L_n(\mu_n, \Delta_n) f(\mu_n, \Delta_n|\theta)}{\int_{\mu,\Delta} L_n(\mu, \Delta) f(\mu, \Delta|\theta) \; d\mu \; d\Delta} \tag{2.15}$$

Equation 2.14 is the expected conditional likelihood of the joint probability of the individual's choices and the unobserved heterogeneity, where the expectation is over the distribution q, which is the conditional probability distribution of the unobserved heterogeneity. This expected likelihood function serves as the basis of the estimation procedure.

The estimation method in this section extends the method of Arcidiacono and Miller (2011). Their method incorporates unobserved heterogeneity as a finite mixture distribution over unobserved types using the Expectation Maximization (EM) algorithm. We extend this approach to allow for any continuous distribution of unobserved heterogeneity by using simulation methods. The estimation procedure uses a CCP representation of the individual's value function in the maximization step of a Simulated EM (SEM) algorithm. An individual's "type" corresponds to a draw from the distribution of unobserved heterogeneity.

Each step of the estimation procedure will be covered in detail after an overview of the entire process. The algorithm begins with initial guesses for the parameters and iterates over the following steps:

1. E-step, part 1: Use the current parameter values to update q.

2. E-step, part 2: Update the Conditional Choice Probabilities (CCPs) using the current parameter values.

3. M-step: Update the parameter estimates by maximizing the simulated likelihood function using the updated CCPs and values of q.

The process terminates when the parameter estimates converge. Since the EM algorithm is an iterative procedure, the maximization step (or M-step) must be performed on each iteration. The SML estimator only requires maximizing the likelihood function a single time. Despite the iterative nature of the estimation, the EM algorithm can still yield considerable computational savings. The use of CCPs in the maximization step can require far fewer computations than full solution methods. Also, the EM algorithm reintroduces additive separability in the maximization step. When the likelihood function is additively separable, it may be possible to estimate the parameters sequentially.

2.4.1 E step

The Expectation step uses the prior iteration parameter estimates, $\hat{\theta} = (\alpha, \bar{\mu}, \Sigma, \sigma_\xi)$, to update the conditional probability distribution of the unobserved variable, q, and the CCPs, $\hat{P}$. The probability distribution of the unobserved variables and the CCPs are functions of μ, which can take an infinite number of values. Each individual, however, can only take on a finite number of values corresponding to the draws needed to simulate the likelihood function. Therefore, these functions only need to be evaluated at a finite number of points. Denote the iteration number with a superscript.

The first step is to use the prior iteration estimates of the population distribution parameters to update the individual parameter values:

$$\mu_n^{m,s+1} = \bar{\mu}^s + chol(\Sigma^s)\eta_n^m, \quad \text{for } m = 1, ..., M \tag{2.16}$$

where $\{\eta_n^m\}_{m=1}^M$ are size J vector draws from the standard normal distribution, and $chol(\Sigma^s)$ is the lower triangular Cholesky decomposition of the population variance matrix. The value of the sequences of signal draws are updated similarly using the current iteration estimate of the standard deviation of the signal. The following equation updates q:

$$q^{m,s+1}(\mu_n, \Delta_n) = \frac{L_n(\mu_n^m, \Delta_n^m)}{\frac{1}{M}\sum_{m=1}^M L_n(\mu_n^m, \Delta_n^m)} \tag{2.17}$$

The CCPs are updated as a weighted multinomial logit of the outcome on a flexible polynomial of the state variables where the weights are q. Alternatively, the CCPs can be updated using the structure of the model. The model can be used to calculate the probability of a given choice at different points in the state space and interpolation methods can be used to generate estimates of the CCPs at other points in the state space.

2.4.2 The M Step

The maximization step uses the updated CCPs and the updated q's to maximize the simulated version of equation (2.14). The updated parameter estimates are:

$$\theta^{s+1} = \max_{\theta} \sum_{n=1}^{N} \frac{1}{M} \sum_{m=1}^{M} q_n^{m,s+1} log(L_n(\theta, \hat{P}^{s+1})) \tag{2.18}$$

It is important to note that equation (2.14) is additively separable in the likelihood of the choice (L_n) and the likelihood of the unobserved heterogeneity ($f(\mu, \Delta|\theta)$). The latter term can be used to update the parameters of the distribution of unobserved variables. The updated distribution parameters are the ML estimate of the mean and variance of the multinomial normal distribution, which simply becomes the mean and covariance matrix of the sample ($\{\mu_n, \Delta_n\}_{n=1}^{N}$) with weights q. The closed form solution for the updated distribution parameters follows Train (2007). Additionally, the EM Algorithm introduces additive separability into the choice likelihood, L_n, which could allow for sequential estimation of the other model parameters (Arcidiacono and Jones 2003). Finally, it is possible to use an alternative version of the EM algorithm that replaces the full maximization in the M step with a single iteration of an optimization procedure. This variant of the EM algorithm is known as the Generalized EM (GEM) algorithm. Using the GEM variant requires more iterations, but can substantially reduce the computation required for each iteration. Full maximization in the M step can be computationally intensive particularly if the optimization procedure uses numerical gradients.[4]

[4]See James (2015) for an application of the GEM algorithm to estimate random parameter Logit models.

2.5 Monte Carlo Results

2.5.1 Simple Learning Model

In this section we present Monte Carlo estimation results for a simplified version of the model where individuals are myopic. In the next section we present results for the estimation of the full model. The motivation for estimating the simplified model is to compare the basic estimation procedure with Simulated Maximum Likelihood. This comparison is not feasible in the full model as performing a large enough number of iterations of the estimation using SML is computationally prohibitive.

The values used in simulation are: $T = 20$ and $J = 2$. For each of 20 simulated data sets, estimation is performed using SML and the EM Algorithm for a given number of draws ($M = \{100, 500, 1000\}$). The simulations are repeated for different numbers of X variables ($K = \{0, 1, 5)\}$), and different sample sizes ($N = \{100, 500\}$). For the EM Algorithm, we use a sequential M step. The updated distribution parameters have a closed form solution. The coefficients of the X variables are updated as a single NR iteration of the simulated likelihood function. Optimization is performed using numerical gradients.[5] All computation is performed in MATLAB on a 12 processor computer.[6] The convergence criteria for the EM algorithm is that every parameter change by less than 0.5% across iterations, which is the criteria suggested by Train (2007), or a cumulative change in the parameters by less than $1e-4$. For SML, the MATLAB defaults for the fminunc optimization command were used. The true parameter values were used as the initial values for both procedures and the same set of random draws were used.[7]

Table 2.1 presents the results for $N = 100$ and $K = 0$. The table reports the mean, standard

---

[5]In the naive learning model, it would be possible to use analytic gradients, which would significantly improve the speed of the estimation. However, in the full dynamic learning model, only numerical gradients would be feasible.

[6]In the full dynamic learning model, certain computationally burdensome parts of the estimation algorithm use FORTRAN Mex files in MATLAB. FORTRAN Mex files are MATLAB executable subroutines that are written and compiled using FORTRAN.

[7]Starting at the true parameters biases the comparison in favor of SML. The EM algorithm performs better than gradient based optimization methods at points far from the optimum, but converges slower near the optimum.

deviation, and median for the set of parameter estimates across simulations for both the EM Algorithm and SML. These summary statistics are also provided for the time it took for the estimation procedure to converge as well as the number of iterations.[8] Finally, we report the number of times the estimation procedure successfully converged of the 20 total simulations.[9] The first panel in table 2.1 reports the results for 100 draws used in simulation. The second and third panel report the results for 500 and 1000 draws respectively. The variables reported include the mean of the population distribution (μ^1 and μ^2), the Cholesky decomposition of the covariance matrix, and the standard deviation of the signal distribution.

Particularly for the higher number of simulation draws, both estimation methods are able to recover the true parameters. The EM Algorithm has difficulty identifying the covariance term when few draws are used. The variability of the parameter estimates is higher for the EM Algorithm, but the EM Algorithm is significantly faster.[10] The M step in the EM Algorithm only uses part of the likelihood function, so we would expect a loss in efficiency relative to SML. For each number of simulation draws, SML takes about the same number of iterations on average to converge. Since the number of computations needed to calculate the simulated likelihood function is proportional to the number of draws, calculating the simulated likelihood function for 500 draws takes about 5 times as long as for 100 draws. Since the number of iterations does not change when 500 draws are used, SML takes about 5 times as long when going from 100 to 500 draws. On the other hand, the EM Algorithm converges in fewer iterations as more draws are used. Each iteration takes longer with more draws, but since fewer iterations are needed, the total estimation time increases by less than the increase in the number of draws. The number of iterations for the EM Algorithm contains some very large outliers (the mean time and mean iterations are much larger than the medians). As the number

[8]If computation were performed on a single processor, the estimation time would be approximately 12 times the numbers reported.

[9]The most common reason why the EM Algorithm failed to converge was due to the variance terms either going to zero or becoming very large.

[10]There is a tradeoff in setting the convergence criteria. A more difficult criteria could lead to less variability in the estimates but would add to the estimation time. However, it is not straight forward to set a consistent convergence criteria for both estimation methods.

of draws increases, these outliers occur less frequently.

Table 2.2 presents the results for $N = 100$ and $K = 1$ and table 2.3 presents the results for $N = 100$ and $K = 5$. The results are qualitatively similar to those without the observed X variables. There is more variation in the parameter estimates from the EM algorithm, but the estimation is significantly faster in most cases. When only 100 draws are used, the average estimation time is less for SML, but the median time is still higher. When more draws are used, the EM algorithm becomes faster on average.

The next set of results repeats the analysis using a larger sample size ($N = 500$). Table 2.4 presents the results for $K = 0$, table 2.5 presents the results for $K = 1$, and table 2.6 presents the results for $K = 5$. The larger sample size results in much less variation in the parameter estimates. Also, the increase in sample size improves the identification of the parameters when few draws are used. Even when 100 draws are used, the EM algorithm is now better able to estimate the covariance parameter. The EM algorithm converges with far fewer iterations and increasing the number of draws no longer results in a reduction in the number of iterations. The time differential between the EM algorithm and SML becomes much more dramatic with the larger sample size. For the case with no X variables ($K = 0$), the EM algorithm performs nearly as well as SML in recovering the true parameters in approximately 10% of the time.

2.5.2 Full Learning Model

For the full dynamic model, we consider the estimation procedure in a stationary infinite horizon framework. The additional complexity of solving the dynamic problem makes Simulated Maximum Likelihood infeasible for the purposes of this exercise. The values used in the simulations are the same as the previous section except the estimation is not done for the largest number of draws ($M = 1000$) or the largest number of variables in X ($K = 5$). The discount factor is fixed at $\beta = 0.9$. In integrating over future continuous state variables in the conditional likelihood function, 500 draws are used to approximate the integrals.

Table 2.7 presents the results for $K = 0$. For a sample size of $N = 100$, there is a lot of variability in the estimated parameters. Increasing the number of draws slightly improves the results. Going from 100 to 500 draws approximately doubles the time it takes to estimate the parameters. Even though each iteration takes longer with more draws, fewer iterations are required for convergence. Going to the large sample size greatly improves the parameter estimates. For the sample size of 500, increasing the number of draws does not reduce the median number of iterations required for convergence, but the average number of iterations falls due to fewer large outliers.

Table 2.8 presents the results for $K = 1$. The results are similar to those with $K = 0$. The larger sample size allows for better identification of the parameters. Increasing the number of draws reduces the number of iterations required for the smaller sample size, but has less of an effect for the larger sample size.

2.6 Conclusion

In this paper we extend CCP estimation techniques to allow for a continuous distribution of unobserved heterogeneity. In many cases, the use of a finite non-parametric distribution of unobserved heterogeneity (i.e., Arcidiacono and Miller, 2011) is likely to perform as well or better than a parametric continuous distribution. One significant exception is the when the parametric assumption is made in the model, as is the case with learning models. The estimation procedure developed in this paper provides a computationally feasible method for the estimation of learning models.

Table 2.1: Naive learning model Monte Carlo results for $N = 100$ and $K = 0$

$M = 100$

Variable	True Value	EM Mean	EM SD	EM Median	SML Mean	SML SD	SML Median
μ^1	1.0	1.430	1.095	1.109	0.923	0.100	0.920
μ^2	1.0	1.457	1.102	1.082	0.968	0.128	0.986
Ch_{11}	2.0	2.396	1.401	1.923	2.024	0.236	2.048
Ch_{22}	2.0	2.060	0.253	2.072	2.081	0.236	2.126
Ch_{21}	-0.5	0.372	2.010	-0.386	-0.438	0.287	-0.357
σ	0.25	0.198	0.103	0.239	0.420	0.245	0.385
Time (min)		8.4	11.3	2.2	9.8	1.7	9.6
Iterations		108.0	145.4	28.0	15.4	2.7	15.0
Successes			20			20	

$M = 500$

Variable	True Value	EM Mean	EM SD	EM Median	SML Mean	SML SD	SML Median
μ^1	1.0	1.190	0.524	1.051	0.947	0.095	0.922
μ^2	1.0	1.205	0.549	1.150	0.988	0.134	0.946
Ch_{11}	2.0	2.192	0.504	2.065	2.028	0.238	2.062
Ch_{22}	2.0	2.109	0.238	2.089	2.050	0.231	2.065
Ch_{21}	-0.5	-0.057	0.834	-0.278	-0.455	0.272	-0.403
σ	0.25	0.238	0.057	0.251	0.404	0.288	0.333
Time (min)		18.0	36.2	7.5	53.2	8.5	52.9
Iterations		47.9	96.1	20.0	16.6	1.8	16.5
Successes			19			20	

$M = 1000$

Variable	True Value	EM Mean	EM SD	EM Median	SML Mean	SML SD	SML Median
μ^1	1.0	1.186	0.596	1.055	0.957	0.095	0.936
μ^2	1.0	1.237	0.637	1.044	0.992	0.132	0.931
Ch_{11}	2.0	2.230	0.675	2.045	1.991	0.234	2.025
Ch_{22}	2.0	2.120	0.264	2.121	2.058	0.239	2.030
Ch_{21}	-0.5	0.058	1.210	-0.334	-0.439	0.231	-0.395
σ	0.25	0.253	0.007	0.252	0.293	0.232	0.236
Time (min)		15.1	8.6	12.7	110.5	21.8	103.7
Iterations		20.2	11.4	17.0	17.1	2.7	17.0
Successes			19			20	

Table 2.2: Naive learning model Monte Carlo results for $N = 100$ and $K = 1$

		EM			SML		
Variable	True Value	Mean	SD	Median	Mean	SD	Median
				$M = 100$			
α^1	1.0	1.035	0.437	0.960	0.953	0.282	0.933
α^2	-1.0	-0.876	0.563	-0.933	-0.958	0.310	-0.878
μ^1	1.0	1.081	0.685	0.914	1.020	0.153	1.005
μ^2	1.0	1.159	0.556	1.114	1.011	0.178	1.029
Ch_{11}	2.0	2.071	0.432	1.986	2.046	0.210	1.969
Ch_{22}	2.0	1.883	0.422	1.766	1.987	0.277	2.008
Ch_{21}	-0.5	-0.347	1.111	-0.745	-0.507	0.353	-0.433
σ	0.25	0.226	0.074	0.250	0.453	0.265	0.425
Time (min)		22.7	50.7	6.0	17.9	2.6	18.6
Iterations		77.6	165.1	23	21.2	3.1	21.0
Successes			20			20	
				$M = 500$			
α^1	1.0	0.968	0.259	0.962	0.934	0.217	0.947
α^2	-1.0	-0.975	0.312	-0.917	-0.943	0.267	-0.945
μ^1	1.0	1.080	0.648	0.864	1.006	0.132	1.004
μ^2	1.0	1.131	0.464	1.123	1.023	0.154	1.019
Ch_{11}	2.0	2.148	0.499	2.073	2.044	0.189	1.988
Ch_{22}	2.0	1.948	0.371	1.859	1.994	0.246	1.929
Ch_{21}	-0.5	-0.309	1.079	-0.671	-0.557	0.256	-0.547
σ	0.25	0.231	0.074	0.253	0.417	0.272	0.369
Time (min)		89.0	192.5	19.7	75.3	14.0	75.2
Iterations		72.9	147.7	19.0	19.4	3.28	19.5
Successes			20			20	
				$M = 1000$			
α^1	1.0	0.985	0.317	1.018	0.958	0.241	1.010
α^2	-1.0	-1.000	0.313	-0.944	-0.950	0.270	-0.911
μ^1	1.0	1.029	0.607	0.848	1.007	0.127	1.000
μ^2	1.0	1.073	0.404	1.088	1.021	0.150	1.008
Ch_{11}	2.0	2.164	0.529	2.092	2.042	0.190	2.041
Ch_{22}	2.0	1.909	0.337	1.864	1.959	0.226	1.956
Ch_{21}	-0.5	-0.420	0.896	-0.631	-0.561	0.282	-0.514
σ	0.25	0.250	0.014	0.253	0.360	0.275	0.270
Time (min)		54.3	63.9	38.9	147.7	31.8	143.6
Iterations		26.4	28.1	20.0	18.9	3.3	19.0
Successes			19			20	

Table 2.3: Naive learning model Monte Carlo results for $N = 100$ and $K = 5$

		EM			SML		
Variable	True Value	Mean	SD	Median	Mean	SD	Median
			$M = 100$				
$\bar{\alpha}^1$	1.0	1.063	0.385	1.046	1.067	0.335	1.029
$\bar{\alpha}^2$	-1.0	-1.018	0.328	-0.978	-0.971	0.274	-0.954
μ^1	1.0	1.002	0.407	1.010	1.029	0.173	1.041
μ^2	1.0	1.157	0.453	1.262	0.985	0.116	0.977
Ch_{11}	2.0	1.968	0.520	2.024	2.044	0.274	1.988
Ch_{22}	2.0	1.865	0.527	1.895	2.103	0.296	2.099
Ch_{21}	-0.5	-0.384	0.998	-0.207	-0.460	0.302	-0.499
σ	0.25	0.340	0.441	0.259	0.450	0.333	0.147
Time (min)		123.9	219.3	35.2	39.5	6.9	38.7
Iterations		141.1	240.0	44.0	27.2	4.9	27.0
Successes			18			20	
			$M = 500$				
$\bar{\alpha}^1$	1.0	1.058	0.321	1.034	1.094	0.325	1.041
$\bar{\alpha}^2$	-1.0	-0.968	0.318	-0.948	-0.996	0.282	-0.969
μ^1	1.0	1.035	0.447	1.035	1.011	0.176	1.030
μ^2	1.0	1.182	0.424	1.187	0.991	0.122	0.979
Ch_{11}	2.0	2.092	0.231	2.059	1.999	0.182	2.015
Ch_{22}	2.0	1.994	0.246	1.988	1.997	0.279	1.950
Ch_{21}	-0.5	-0.177	0.798	-0.365	-0.565	0.244	-0.552
σ	0.25	0.252	0.016	0.253	0.417	0.317	0.321
Time (min)		117.3	65.3	91.2	168.9	23.8	163.8
Iterations		38.8	20.4	29.5	23.4	3.22	23.5
Successes			20			20	
			$M = 1000$				
$\bar{\alpha}^1$	1.0	1.076	0.313	1.067	1.106	0.305	1.045
$\bar{\alpha}^2$	-1.0	-0.966	0.297	-0.938	-0.991	0.270	-0.968
μ^1	1.0	1.020	0.420	1.082	0.999	0.169	1.030
μ^2	1.0	1.166	0.407	1.208	0.995	0.120	0.993
Ch_{11}	2.0	2.076	0.189	2.063	1.992	0.181	1.958
Ch_{22}	2.0	2.009	0.254	1.973	1.989	0.301	2.041
Ch_{21}	-0.5	-0.225	0.703	-0.290	-0.546	0.226	-0.552
σ	0.25	0.240	0.022	0.247	0.409	0.350	0.295
Time (min)		268.7	177.5	196.6	321.9	44.2	332.4
Iterations		44.6	27.3	34.5	22.5	3.3	22.0
Successes			20			20	

Table 2.4: Naive learning model Monte Carlo results for $N = 500$ and $K = 0$

$M = 100$

		EM			SML		
Variable	True Value	Mean	SD	Median	Mean	SD	Median
μ^1	1.0	1.078	0.114	1.094	0.994	0.066	0.983
μ^2	1.0	1.044	0.166	1.026	0.964	0.053	0.966
Ch_{11}	2.0	1.951	0.115	1.946	1.980	0.106	1.984
Ch_{22}	2.0	1.973	0.125	1.991	2.060	0.166	2.025
Ch_{21}	-0.5	-0.475	0.247	-0.476	-0.450	0.115	-0.447
σ	0.25	0.248	0.008	0.250	0.229	0.145	0.224
Time (min)		6.6	3.6	5.7	57.0	9.4	57.1
Iterations		16.8	9.0	14.5	15.7	2.1	16
Successes			18			20	

$M = 500$

μ^1	1.0	1.069	0.128	1.057	1.017	0.056	1.023
μ^2	1.0	1.036	0.169	0.991	0.979	0.049	0.985
Ch_{11}	2.0	2.008	0.112	2.021	2.010	0.125	2.026
Ch_{22}	2.0	2.023	0.120	2.011	2.040	0.106	2.054
Ch_{21}	-0.5	-0.450	0.218	-0.437	-0.486	0.112	-0.472
σ	0.25	0.250	0.005	0.251	0.340	0.152	0.307
Time (min)		31.0	18.8	26.9	307.2	42.7	315.1
Iterations		16.1	9.7	14.0	17.3	2.2	17.5
Successes			20			20	

$M = 1000$

μ^1	1.0	1.072	0.127	1.071	1.017	0.054	1.020
μ^2	1.0	1.038	0.167	0.998	0.984	0.049	0.982
Ch_{11}	2.0	2.017	0.111	2.027	2.015	0.122	2.015
Ch_{22}	2.0	2.041	0.119	2.036	2.024	0.107	2.032
Ch_{21}	-0.5	-0.429	0.222	-0.394	-0.515	0.092	-0.498
σ	0.25	0.251	0.003	0.251	0.350	0.158	0.322
Time (min)		62.9	37.5	55.3	624.9	122.2	593.9
Iterations		16.5	9.9	14.5	18.4	2.7	19.0
Successes			20			20	

Table 2.5: Naive learning model Monte Carlo results for $N = 100$ and $K = 1$

		EM			SML		
Variable	True Value	Mean	SD	Median	Mean	SD	Median
				$M = 100$			
α^1	1.0	0.952	0.105	0.956	0.961	0.112	0.949
α^2	-1.0	-1.009	0.129	-1.008	-1.006	0.140	-1.010
μ^1	1.0	0.950	0.219	0.965	0.976	0.073	0.980
μ^2	1.0	1.073	0.136	1.046	1.016	0.064	1.011
Ch_{11}	2.0	1.991	0.157	1.987	2.066	0.110	2.070
Ch_{22}	2.0	1.940	0.159	1.923	2.012	0.120	2.029
Ch_{21}	-0.5	-0.513	0.277	-0.581	-0.478	0.104	-0.462
σ	0.25	0.250	0.011	0.250	0.276	0.171	0.269
Time (min)		21.1	14.0	16.1	73.4	12.5	70.9
Iterations		20.4	12.7	16.0	17.7	2.9	18.0
Successes			20			20	
				$M = 500$			
α^1	1.0	0.971	0.102	0.999	0.973	0.105	1.000
α^2	-1.0	-1.021	0.114	-1.024	-1.015	0.129	-1.008
μ^1	1.0	0.933	0.191	0.979	0.988	0.066	1.004
μ^2	1.0	1.061	0.120	1.068	1.032	0.068	1.044
Ch_{11}	2.0	2.047	0.116	2.038	2.048	0.101	2.037
Ch_{22}	2.0	1.982	0.115	1.984	2.038	0.113	2.058
Ch_{21}	-0.5	-0.511	0.214	-0.564	-0.494	0.092	-0.489
σ	0.25	0.251	0.005	0.249	0.383	0.153	0.411
Time (min)		82.8	40.6	73.9	394.6	77.1	376.7
Iterations		16.9	7.6	15.5	18.9	3.06	18.5
Successes			20			20	
				$M = 1000$			
α^1	1.0	0.980	0.090	0.980	0.980	0.094	0.998
α^2	-1.0	-1.012	0.121	-1.022	-1.009	0.128	-1.005
μ^1	1.0	0.933	0.187	0.933	0.989	0.066	1.015
μ^2	1.0	1.059	0.133	1.087	1.032	0.066	1.046
Ch_{11}	2.0	2.049	0.124	2.034	2.051	0.102	2.045
Ch_{22}	2.0	1.991	0.106	1.994	2.033	0.111	2.057
Ch_{21}	-0.5	-0.508	0.189	-0.508	0.513	0.102	0.505
σ	0.25	0.250	0.004	0.250	0.347	0.185	0.353
Time (min)		159.3	83.6	139.7	800.3	142.8	765.8
Iterations		16.6	8.0	15.0	19.2	2.3	19.5
Successes			20			20	

Table 2.6: Naive learning model Monte Carlo results for $N = 500$ and $K = 5$

		EM			SML		
Variable	True Value	Mean	SD	Median	Mean	SD	Median
			$M = 100$				
$\bar{\alpha}^1$	1.0	0.956	0.126	0.952	0.982	0.121	0.982
$\bar{\alpha}^2$	-1.0	-0.951	0.141	-0.960	-0.975	0.144	-0.979
μ^1	1.0	1.108	0.205	1.099	0.953	0.071	0.945
μ^2	1.0	1.105	0.219	1.137	0.979	0.087	0.967
Ch_{11}	2.0	1.964	0.122	1.947	2.032	0.132	2.037
Ch_{22}	2.0	1.977	0.201	1.996	2.017	0.130	2.009
Ch_{21}	-0.5	-0.349	0.314	-0.348	-0.511	0.152	-0.505
σ	0.25	0.253	0.023	0.248	0.231	0.147	0.243
Time (min)		103.8	75.0	65.9	163.6	27.1	157.8
Iterations		33.7	21.8	23	22.7	3.5	21
Successes			20			20	
			$M = 500$				
$\bar{\alpha}^1$	1.0	1.010	0.113	1.013	1.009	0.109	1.015
$\bar{\alpha}^2$	-1.0	-0.982	0.132	-0.994	-0.995	0.131	-1.007
μ^1	1.0	1.031	0.166	1.019	0.966	0.065	0.957
μ^2	1.0	1.035	0.186	1.019	1.000	0.077	0.993
Ch_{11}	2.0	2.007	0.115	1.993	2.015	0.111	1.983
Ch_{22}	2.0	1.998	0.165	2.046	1.996	0.110	2.010
Ch_{21}	-0.5	-0.422	0.223	-0.445	-0.530	0.105	-0.505
σ	0.25	0.250	0.005	0.250	0.334	0.184	0.318
Time (min)		306.0	212.3	228.0	870.1	175.7	843.5
Iterations		22.1	14.2	17.0	23.4	2.8	23.5
Successes			20			20	
			$M = 1000$				
$\bar{\alpha}^1$	1.0	1.008	0.106	1.013	1.015	0.108	1.023
$\bar{\alpha}^2$	-1.0	-0.979	0.129	-0.988	-0.996	0.132	-0.993
μ^1	1.0	1.043	0.175	0.994	0.964	0.071	0.955
μ^2	1.0	1.043	0.191	1.006	1.005	0.083	1.010
Ch_{11}	2.0	2.011	0.118	1.997	2.011	0.096	2.033
Ch_{22}	2.0	2.013	0.163	2.045	1.985	0.121	1.987
Ch_{21}	-0.5	-0.404	0.230	-0.409	-0.534	0.113	-0.517
σ	0.25	0.249	0.003	0.249	0.354	0.215	0.297
Time (min)		511.2	242.1	436.5	1768.5	285.2	1811.9
Iterations		18.5	8.4	16.0	23.5	3.7	23.5
Successes			20			20	

Table 2.7: Full learning model Monte Carlo results for $K = 0$

N=100							
		$M = 100$			$M = 500$		
Variable	True Value	Mean	SD	Median	Mean	SD	Median
μ^1	1.0	0.802	0.587	0.600	0.999	0.688	0.742
μ^2	1.0	0.696	0.583	0.670	0.960	0.753	0.693
Ch_{11}	2.0	1.505	0.578	1.367	1.808	0.500	1.723
Ch_{22}	2.0	1.676	0.889	1.605	1.623	0.526	1.511
Ch_{21}	-0.5	-0.786	1.124	-0.993	-0.136	1.261	-0.751
σ	0.25	0.379	0.337	0.249	0.235	0.061	0.239
Time (min)		54.0	62.0	27.5	91.3	102.3	61.8
Iterations		143.3	150.7	74	90.7	110.2	59.5
Successes			19			20	
N=500							
		$M = 100$			$M = 500$		
μ^1	1.0	1.059	0.346	1.029	1.132	0.339	1.060
μ^2	1.0	0.968	0.326	0.846	1.038	0.330	1.041
Ch_{11}	2.0	2.008	0.281	2.054	2.142	0.307	2.086
Ch_{22}	2.0	1.922	0.364	1.863	2.012	0.366	2.039
Ch_{21}	-0.5	-0.493	0.542	-0.581	-0.344	0.596	-0.461
σ	0.25	0.238	0.051	0.248	0.251	0.017	0.254
Time (min)		64.3	62.6	41.4	239.5	98.2	203.3
Iterations		67.0	69.1	42	53.1	22.6	43.5
Successes			20			20	

Table 2.8: Full learning model Monte Carlo results for $K = 1$

Variable	True Value	N=100, $M = 100$ Mean	SD	Median	$M = 500$ Mean	SD	Median
α^1	1.0	0.801	0.808	0.667	0.707	0.823	0.743
α^2	-1.0	-0.934	0.606	-0.909	-0.762	0.658	-0.763
μ^1	1.0	0.816	0.777	0.499	1.010	0.993	0.707
μ^2	1.0	1.097	0.975	0.797	1.177	1.007	0.827
Ch_{11}	2.0	1.720	0.739	1.616	2.117	0.771	1.962
Ch_{22}	2.0	1.812	1.106	1.728	1.785	0.697	1.874
Ch_{21}	-0.5	-0.336	1.369	-0.536	0.056	1.540	-0.316
σ	0.25	0.513	0.687	0.264	0.243	0.042	0.248
Time (min)		320.9	389.5	114.6	452.5	270.1	367.2
Iterations		186.2	186.3	88.0	81.8	50.8	67.5
Successes			20			20	
N=500		$M = 100$			$M = 500$		
α^1	1.0	0.898	0.286	0.844	0.945	0.319	0.872
α^2	-1.0	-0.976	0.258	-0.987	-1.000	0.267	-1.055
μ^1	1.0	1.013	0.343	0.932	1.019	0.280	0.960
μ^2	1.0	1.053	0.303	1.037	1.044	0.260	1.060
Ch_{11}	2.0	1.950	0.379	1.989	2.081	0.262	2.089
Ch_{22}	2.0	1.871	0.366	1.919	1.934	0.338	1.903
Ch_{21}	-0.5	-0.545	0.481	-0.657	-0.504	0.456	-0.629
σ	0.25	0.275	0.073	0.258	0.249	0.009	0.247
Time (min)		516.4	591.8	331.3	1923.0	1211.7	1399.5
Iterations		66.2	69.6	44.0	52.2	29.7	43.0
Successes			20			19	

CHAPTER 3

DO CONSUMERS' BELIEFS CONVERGE TO EMPIRICAL DISTRIBUTIONS WITH REPEATED PURCHASES? (CO-AUTHOR FORREST SPENCE)

3.1 Introduction

Price dispersion is a feature of many markets and even occurs in markets for homogeneous goods or services (Stigler 1961). One possible reason for the persistence of price dispersion is that consumers have limited information over prices and acquiring information may be costly. In markets with limited information and costly search, an individual may not purchase from the seller with the lowest price if she is unaware of that price. Theoretical models of consumer search incorporate the search decision into a model of consumer demand by assuming that individuals have beliefs about the empirical distribution of prices in the market and must incur a cost to reveal price information from one or more retailers before deciding whether to purchase the good or service (e.g., Reinganum, 1979; Burdett and Judd, 1983). The decision to search depends upon the magnitude of the search costs as well as the individual's subjective beliefs about the distribution of prices. When estimating models of consumer search, researchers may impose assumptions on individuals' beliefs in order to recover estimates of search costs. In this paper, we test the validity of these assumptions using data on the observed distribution of prices for the online textbook market and data on individuals' subjective beliefs about this distribution.

There is a growing literature focusing on the development and estimation of structural models of consumer search. These models have been used to explain observed price dispersion for homogeneous goods (Hortaçsu and Syverson 2004; Hong and Shum 2006), test competing models of consumer search (De los Santos, Hortacsu, and Wildenbeest 2012b), and to recover demand estimates in markets where price uncertainty is important (Koulayev 2012; Moraga-González, Sándor, and Wildenbeest 2009). A critical assumption used in these studies is that consumers have rational expectations, (i.e. the price of a product is a random variable, but consumers know the parameters that govern the distribution of prices. However, if consumers have biased beliefs about the parameters of the empirical distribution of prices, this will lead to biased estimates of search costs. In particular, if consumers' beliefs about prices are biased upward, the rational expectations assumption will bias search cost estimates upwards and bias price elasticity estimates towards zero (low levels of search can be explained by either high search costs or low expected benefits from search). By comparing subjective beliefs to actual observed price distributions, we are able to test the validity of this assumption.

In addition to testing the validity of the rational expectations assumption, we also investigate the degree to which experienced consumers have more accurate beliefs than their less experienced counterparts. Recent research has supported this idea by incorporating learning into consumer search models.[1] In these models, consumers learn about the parameters of the empirical price distribution within a single purchasing decision through a sequential search process (De los Santos, Hortacsu, and Wildenbeest 2012a; Koulayev 2009, 2013). We focus instead on learning across purchasing decisions; in particular we examine the hypothesis that more experienced consumers have acquired information about the empirical price distribution through repeated participation in the market.[2]

We use data on the empirical distribution of textbook prices from online retailers and consumers' subjective beliefs about this distribution. In order to obtain data on individuals' subjective beliefs, we provide an online questionnaire to 1,224 undergraduate students with multiple textbook purchasing scenarios in order to elicit their beliefs about prices. For each hypothetical textbook purchasing scenario, students are given the price of a textbook from the campus bookstore and are asked about their expectations of the lowest price available from an online retailer. Additional questions are asked to elicit consumers' beliefs about the variability of the lowest price. For example, if a consumer

[1]Earlier studies examined learning and search through experimental designs (e.g., Sonnemans, 1998; Einav, 2005)

[2]The research questions we address in this paper are further supported by research in the labor literature, which uses subjective beliefs about future earnings to explain college major choice (Arcidiacono, Hotz, and Kang 2012; Stinebrickner and Stinebrickner 2011; Wiswall and Zafar 2012). These studies show that incorporating students' subjective beliefs leads to significantly different estimates than those obtained under the assumption of rational expectations. In addition to this, Wiswall and Zafar (2012) show that college students' beliefs about future earnings become more consistent with the actual earnings distribution as they proceed through school (i.e., become more familiar with their field).

reports an expected online price of \$100, she is then asked about the likelihood that the actual price is below \$95.

Our results show that inexperienced consumers have price expectations that are significantly greater than the mean of the empirical price distribution for both new and used textbooks. Therefore, we can reject the hypothesis that inexperienced consumers know the parameters of the price distribution for the online textbook market. Individuals with higher levels of experience, measured by the number of prior online textbook purchases, typically have price expectations that are closer to the empirical mean. For used books, individuals tend to underestimate the variation of the empirical distribution, and beliefs about the variation of the price distribution do not appear to become more accurate with experience. Overall, the evidence is consistent with learning, at least for learning about the mean of the price distribution.

The following section provides theoretical motivation for this project and expands on our goals. Section 3 describes the data and Section 4 presents results. Section 5 discusses the issue of selection, and section 6 concludes.

3.2 Theoretical Motivation

We use the following simple model of consumer search to motivate the empirical section of this paper. Individuals can purchase a given product from two locations. Assume for simplicity that the search cost is zero for one of the locations, so the individual knows the price of the product at this location. The price of the product at the other location is unknown by the individual, and there is a cost associated with determining this price. Denote the price at the zero search cost location as p^* and the price at the location with a search cost as p, which is a random variable with cumulative density function, $F(p)$. The individual can either purchase the product from the first location or pay some cost, c, to search and discover the price at the other location. If the individual decides to search, he does not incur an additional search cost should he choose to purchase the product from the first location (i.e., search with recall).

The decision rule for the search problem is given by Equation (3.1). An individual chooses to search if,

$$\int_0^{p^*} (p^* - p) d\tilde{F}_i(p) > c_i \tag{3.1}$$

where $\tilde{F}_i$ denotes an individual's beliefs about the empirical price distribution. The LHS of Equation (3.1) is the expected benefit of search. A individual integrates over the difference between the known (p^*) and unknown price (p), given his beliefs about the distribution of the unknown price. The domain of integration is bounded above by p^* because an individual can costlessly revisit the first location (i.e., the benefit from search is weakly positive).

The RHS of Equation (3.1) is an individual specific search cost c_i. The majority of the structural consumer search literature attempts to recover the distribution of individuals' search costs. In order to do so, the econometrician must make assumptions regarding individuals' beliefs, $\tilde{F}$. A common assumption regarding individuals' beliefs is that there is no learning over the parameters of the distribution, and individuals have rational expectations. In other words, individuals are assumed to know the parameters of the distribution of p.[3]

In this paper, we focus on the first two moments of individuals' beliefs. Determining if these moments match the corresponding moments of the empirical price distribution is important for the estimation of search costs. If consumers overestimate the mean of the empirical price distribution, then the model will generate an upwardly biased distribution of search costs under the rational expectations assumption. Similarly, if consumers underestimate the variance of the empirical price distribution, search cost estimates will also be biased upward.[4]

An alternative to rational expectations is to allow uncertainty and learning over the parameters

---

[3] An alternative to making a parametric assumption on the empirical price distribution and consumers' beliefs is to instead assume that consumers form beliefs non-parametrically based on the empirical CDF of observed prices:

$$F(p) = \tilde{F}_i(p) = \frac{1}{N} \sum_{k=1}^{N} I[p_k < p]$$

where N is the number of observed prices. If consumers' beliefs are biased relative to the empirical distribution, this leads to similar biases in search costs that are discussed under the assumption of a parametric distribution for prices and beliefs.

[4] Misspecification of beliefs also leads to biases in price elasticity estimates. If individuals' beliefs are biased such that they underestimate the benefit of search (relative to the assumed, true benefit), then the model will recover price elasticities that are lower relative to the true elasticities.

of the price distribution. When individuals search and observe a price draw, they can use this information to update their beliefs according to a learning process (e.g., Bayesian). Even in the learning framework, however, some variant of the rational expectations assumption is commonly used to restrict individuals' initial prior beliefs as the initial priors are typically not separately identified. In the empirical section of the paper we test whether inexperienced individuals have biased beliefs about the parameters of the price distribution. We also examine whether individuals' beliefs are consistent with learning by testing whether more experienced individuals have beliefs that are closer to the parameters of the empirical price distribution.

3.3 Data

We collected data on subjective beliefs about the distribution of prices in the online textbook market through online questionnaires sent to students at the University of North Carolina at Chapel Hill (UNC).[5] The questionnaires asked individuals about their previous textbook purchasing behavior and presented them with hypothetical textbook purchasing scenarios. We supplement the responses to these textbook purchasing scenarios with price data scraped from an online marketplace for a large number of textbooks. Before providing a summary of both datasets, we will provide more detailed information about the textbook purchasing scenarios.

3.3.1 Textbook Purchasing Scenarios

Each questionnaire contained three randomly assigned hypothetical textbook purchasing scenarios from a total of twelve potential scenarios.[6] Figure 3.1 is a screenshot of the information provided in one particular scenario.[7]

After being presented with information about the scenario, respondents were provided with the (actual) price of a new copy of the textbook from the campus bookstore, and were asked to give their

[5] Appendix Section B.1.1 contains the text from the online questionnaire provided to consumers. Individuals who agreed to participate in the survey were sent a link to the questionnaire.

[6] These twelve textbooks include four textbooks each from physical sciences, social sciences, and humanities. Of the four textbooks within these general fields, two are from introductory level courses. More information on the characteristics of the textbooks used in the hypothetical purchasing scenarios can be found in Appendix B.1.2.

[7] For each scenario, we randomly assigned respondents to a full information case (title, author, publisher, picture, etc.) or a limited information case. As opposed to the full information case, as seen in Figure 3.1, the limited information case only provided information on the title, author, and course.

expectations about the lowest price they would find for a new copy of this textbook if they searched only one online retailer.[8] Respondents were then presented with the actual price of a *used* copy of the textbook from the campus bookstore (including taxes) and were asked to give their expectations about the lowest price they would find for a used copy online if they searched one online retailer (including shipping fees).

In order to elicit information about individuals' beliefs about the higher order moments of the price distribution, we then asked respondents for the probability that the price realized after search would be less than X% or greater than Y% of their reported expected price for both new and used copies of the textbook. For example, in Figure 3.1, the new price of the textbook at the campus bookstore for the Fall 2012 semester was $87.00. If the respondent reported that her expectation of the lowest price for a new copy of the textbook from one online retailer was $50.00, then the next questions would ask her the probability that the price would be less than $45.00 and the probability that the price would be greater than $55.00. In practice, X was randomly drawn from {85, 90, 95} and Y was randomly drawn from {105, 110, 115}.

Given that individuals may not be accustomed to thinking about prices in a probabilistic manner, we first presented individuals with an example in order to help clarify the questions within the textbook purchasing scenarios. In the example, we asked individuals to consider the lowest price they might find for a pair of jeans if they searched one retailer at the mall. This example contained information about probabilities (e.g., that their response should be between 0 and 100 percent) and clarification about the nature of price uncertainty (i.e. that although their best guess might be $20, there is some chance that the price is actually lower or higher than $20).

3.3.2 Online Questionnaire Data

We conducted two waves of the survey. The first was during the Fall semester of 2012, and the second was during the Spring semester of 2013. For the Fall 2012 and Spring 2013 semesters, 820 and 798 respondents completed the background questions about their previous textbook purchasing

[8] The bookstore price provided to students explicitly included sales tax. Respondents were asked to include shipping fees when providing their expectation of the lowest price available. Respondents were also reminded to not actually search for the lowest price of the textbook.

experience, respectively. The sample used in analysis is composed of 739 respondents from the Fall 2012 semester and 726 respondents from the Spring 2013 semester.[9] 104 respondents (52 from both semesters) were dropped because they had been enrolled in college for more than 10 semesters and an additional 49 respondents (29 from Fall 2012 and 20 from Spring 2013) were dropped for reporting nonsensical answers (e.g., reporting an expected price of $100,000).[10] Appendix B.1.3 provides a more detailed description of within survey attrition.

Table 3.1 displays the number of semesters enrolled for the questionnaire respondents. This count includes both traditional fall and spring semesters and any summer sessions the students had previously been enrolled in. Individuals in later semesters are slightly over sampled due to the nature of how we recruited individuals for the study. We obtained the email addresses of individuals who participated in a separate, longer running data collection project and agreed to receive follow-up emails. This other project began in the Fall of 2011 and recruited new individuals each semester. Individuals who participated at the start of this other project would be at least in their third semester at the time of data collection (assuming continual enrollment). Appendix B.1.1 provides more detail on how individuals were recruited.

Respondents' previous textbook purchasing behavior and major choice are also reported in Table 3.1. A majority of respondents have purchased textbooks at the campus bookstore and from an online retailer. There is significant variation in how many textbooks respondents have purchased online; 33.6% of the individuals in the sample have purchased five textbooks or fewer from online retailers. Approximately a third of respondents reported either Economics or a STEM field as one of their stated majors.

3.3.3 Online Retailer Data

In order to construct an empirical distribution of prices for textbooks, we used a script in Perl to scrape .html files from Amazon.com. We collected daily price data for approximately 3,500 books

[9]There were 240 individuals who participated in both surveys.

[10]In practice, this was done by removing respondents who reported expectations less than 10% or greater than 150% of the bookstore price. In Appendix B.2.1 we report out main results for a more relaxed omission criterion. The results are substantially the same.

that were assigned at UNC during the 2012-2013 school year. Using these .html files, we used a separate script in Perl to parse the lowest prices available for both new and used copies of the books on each day.[11] Since the survey asked individuals about their expectations of the lowest available price for a particular book, we define the empirical distribution as the distribution of the lowest online price as a proportion of the price at the campus bookstore across textbooks. We use the daily price data for two intervals corresponding to the timing of the surveys. The Fall survey period is from November 30, 2012 to December 10, 2012, and the Spring survey period is from April 11, 2013 to April 26, 2013.[12] To construct the empirical price distribution, we use the average price of the textbook over the survey period. The price sample used in the analysis trims the top and bottom 0.5% of the prices for each survey period.[13]

The total cost of purchasing books online includes shipping fees, which vary depending on the speed of delivery. For items purchased on the Amazon Marketplace from third party sellers, we added the fee for standard shipping. Items purchased directly from Amazon qualify for free standard shipping as long as the item is purchased in as part of an order that exceeds a certain amount.[14] Most new textbooks will qualify for free shipping if purchased directly from Amazon, so we do not add any shipping fees to the price of these books.[15] We include sales tax in the campus bookstore prices. Sales taxes are not included in the online prices.[16]

---

[11]Further analysis could incorporate additional prices from these .htmls files such as the lowest price conditional on reported quality (e.g. very good, good, etc.).

[12]The online questionnaire was initially distributed on November 30, 2012 for the fall survey and April 11, 2013 for the spring survey. Nearly all of the surveys were completed during these intervals. We take these periods as the the time frames that individuals are forming their expectations over. This is potentially problematic as online textbook prices vary systematically across the year (e.g., they are generally higher in August than May.). Further analysis could examine additional time frames in the construction of the empirical distribution.

[13]The trimmed sample excludes books that have an online price listing that is either a very small fraction or a large multiple of the bookstore price. In some cases, particularly for books with low sales volumes, the automated pricing algorithms used by larger book resellers can generate these extreme prices.

[14]Orders that exceeded $25 qualified for free shipping at the time of the surveys.

[15]We do not include shipping for books that do not individually qualify for free shipping because they could be purchased as part of a larger order that does qualify for free shipping. The empirical analysis focuses on higher priced books that would qualify for free shipping. All of the books in the hypothetical textbook purchasing scenarios qualify for free shipping if purchased new from Amazon.

[16]At the time of this analysis, Amazon did not collect sales taxes. Individuals were responsible for paying the sales

Table 3.2 provides the ratio of prices of textbooks from Amazon.com relative to the price from the campus bookstore. The first row reports the prices of new books for the full sample of books for which we have data. On average, new prices on Amazon.com are approximately 85% of the bookstore price. The second to last column reports the average difference between the price of the textbook from the bookstore and an online retailer. For all textbooks in our sample, the savings in absolute terms is approximately $10.

The second row reports the new prices that includes new books listed on the Amazon marketplace by third party sellers. Including the marketplace listings increases the savings relative to the bookstore price. On average, used prices on Amazon.com are approximately 76% of the used bookstore price. This corresponds to an average difference of approximately $33. The median is lower than the mean for both new and used books, as the distributions are slightly skewed to the right. On average the prices during the Spring survey period were slightly lower than during the Fall period.

The next three rows of Table 3.2 provide summary statistics for textbooks which are priced greater than $100 for a new copy from the campus bookstore. Books with a list price below $100 include popular press titles that have a large market outside of being assigned for a college course. The restricted sample of books which are priced greater than $100 at the campus bookstore consists primarily of books that are commonly thought of as textbooks. Relative to the full sample, the potential savings from shopping online becomes greater for both new books and used books (i.e. in both percentage and magnitude terms, more expensive textbooks have greater savings in the online market). The variability of prices is less for both new and used books relative to the full sample. The final three rows provide summary statistics for the textbooks used in the hypothetical textbook purchasing scenarios.[17] On average, these prices are slightly lower than the sample of textbooks with a price of $100 or more at the campus bookstore, but the difference is not significant.

Ideally, how we define the empirical price distribution should match the price distribution of the individuals' beliefs, but there are a few reasons why this may not be the case. First, textbook

taxes for online purchases, however compliance was low. Sellers on the Amazon Marketplace are responsible for paying any applicable sales taxes, so sales taxes are already included in the listed prices.

[17] Note that the total number of textbooks in the purchasing scenarios is actually 12. However, online retailer data for one textbook is missing.

prices vary over time, and the time frame used to define the empirical distribution may not match the time frame of the individuals' subjective beliefs. Second, we only use price data from a single online retailer. We believe the prices from Amazon.com provide a reasonable approximation to the empirical distribution of prices that consumers face if they only search one online retailer. Of the individuals in our sample, 75% reported Amazon.com as the first website they would visit to search for a textbook. These issues of timing and alternative retailers affect the comparison between the individuals' subjective beliefs and the empirical price distribution. The comparison of individuals' subjective beliefs across different levels of experience is not affected as long as individuals with different levels of experience do not systematically differ in the time frame considered or in the choice of the online retailer.

3.4 Results

The first subsection presents results using the data on reported expectations. The following subsection incorporates additional data on beliefs to examine not only individuals' expectations but also individuals' beliefs about the variance of the empirical price distribution in the context of a parametric learning model. The online survey asks individuals to report what they thought the price of the textbook would be if they searched one online retailer. We interpret the responses to this question as corresponding to individuals' subjective beliefs about the mean of the price distribution of the lowest price for a particular textbook.

3.4.1 Expectations Results

In this section we present descriptive statistics of individuals' price expectations. Then, we test for differences in price expectations relative to the empirical prices across levels of experience in order to determine if consumers' expectations converge to the mean of the empirical price distribution. Finally, we perform regressions to control for additional characteristics of the respondents and the textbook scenarios.

The first columns of table 3.3 provide the summary statistics of the reported expectations of the lowest online price as a proportion of the bookstore price for individuals with different levels of online textbook purchasing experience. In the survey, individuals were asked about the number of textbooks they had ever purchased online, and they responded by selecting one of four possible categories. Individuals with no prior online textbook purchases expect the price of a used book online to be approximately 74% of the price of a used book at the college bookstore. This corresponds to an expected savings of \$31.53 on average across the hypothetical textbook purchasing scenarios. Individuals with prior online textbook purchases expect the online price to be lower, with higher levels of experience corresponding with a greater expected savings. On average, individuals with more than ten previous online purchases expect the price of a used book online to be approximately 65% of the price of a used book at the college bookstore. This corresponds to an average savings of \$41.65. The results from the spring survey display a similar pattern.

Table 3.4 repeats the analysis done in table 3.3 using level differences instead of the normalized price ratio. Consumers across experience levels expect for there to be an average savings of \$30 to \$40 for textbooks from online retailers. The patterns across experience levels are the same when using levels as using ratios. As consumers gain experience, they expect to find larger savings in the online market.

These results demonstrate that higher levels of experience are associated with lower expectations of online textbook prices. This relationship would be consistent with learning if the individuals with higher levels of experience report expectations that are closer to the true mean of the price distributions. The final two columns of table 3.3 report the difference between the average of the reported expectations and the mean of the empirical price distribution for the sample of scenario textbooks as well as the sample of textbooks with a list price greater than \$100. On average, the reported expectations become closer to the empirical mean at higher levels of experience. For the scenario textbooks, the difference between the mean of the reported expectations and the empirical mean is not significant at any level of experience.[18] For the sample of books with list price greater than \$100 for the fall survey, this difference is significant at the lowest levels of experience and is not significant at the higher levels of significance. For the spring survey, the difference between the mean of the reported expectations and the mean of the empirical price distribution is significant at all levels of experience for books with a list price greater than \$100. This is due to the mean of the

[18]This result is primarily due to the small sample size for the scenario textbooks.

empirical price distribution being lower during the spring survey period.

Table 3.5 reports the results for new books. Individuals with higher levels of experience report lower expected prices on average. Individuals with no experience expect the online price to be 83% of the price of the textbook from the campus bookstore (a $26.67 average savings). Individuals with eleven or more previous online purchases expect the online price to be 77% of the price from the campus bookstore (a $37.58 average savings). The difference between the empirical mean and the mean of the reported expectations decreases for higher levels of experience. However, unlike the results for used books, individuals with higher levels of experience have expectations that are on average significantly below the empirical prices. One explanation for this result is that the new price is the price for purchasing the book directly from Amazon.com. When the new price is defined as the minimum of the marketplace price and the price charged by Amazon.com, the reported expectations are significantly greater than the empirical mean for all levels of experience. Some individuals likely include the marketplace when forming their beliefs about the prices of new textbooks. One possible explanation for the relationship between experience and price expectations for new books is that individuals are learning about the availability of new textbooks by third party sellers.

Due to the nature of the data collection, we want to control for differences in the textbook purchasing scenarios that individuals are given and control for additional characteristics of the individual which may explain the differences in price expectations across levels of experience. Table 3.6 reports results from a regression of normalized price expectations on level of experience, textbook characteristics, scenario characteristics and additional individual controls. The scenario characteristics include indicators to control for the different possible scenarios, the survey period, and whether the textbook purchasing scenario was a full information case (details were provided on textbook characteristics such as years since revision, etc.). The additional individual controls include indicators for whether the individual has previously taken the course for which the textbook was assigned and whether the individual has previously been assigned the textbook in the scenario.

The regression estimates are consistent with the mean comparisons above. Individuals who have never made an online textbook purchase before have significantly higher price expectations than individuals who have purchased a textbook online. Price expectations evolve gradually, as individuals in the highest category of experience consistently have lower expectations.

The coefficients on indicators for whether the respondent had previously taken the course or been assigned the textbook are consistently negative, but only statistically significant for individuals who had previously taken the course. The coefficient on the number of years since the last revision is negative and significant, perhaps reflecting beliefs about a greater supply of textbooks in the secondary market. As the number of years since a textbook has been revised increases, the supply of textbooks in the secondary market increases, generally reducing the price of the textbook. Consumers seem to internalize this when making a textbook purchasing decision, which supports the results in Chevalier and Goolsbee (2009). Similarly, consumers have higher price expectations for textbooks that are the latest edition released (two of the twelve textbook scenarios were for previous editions).

Table 3.7 reports regression results with year in school dummies. These results show that the differences in beliefs are due to differences in direct online textbook purchasing experience rather than from indirect experience (e.g., word of mouth).

3.4.2 Distribution Results

In this subsection, we examine whether the patterns observed in the data are consistent with learning over additional parameters/moments of the empirical price distribution. In the hypothetical textbook purchasing scenarios, individuals report their expectations for prices as well as the probability that a draw from the price distribution is below a given threshold ($\mathbf{E}[H]$ and $F_p(p_L; \mu, \sigma)$). We use these two moments to calculate the expected parameters of each individual's beliefs (i.e., $\mathbf{E}[\mu]$ and $\mathbf{E}[\sigma]$), under the assumption that individuals believe that prices follow a log-normal distribution.[19] The log normal distribution has two properties that make it an appropriate distribution in the current context. First, the support of the distribution is non-negative real numbers and prices are bounded below by zero. The second feature is that the log normal distribution is skewed to the right, which is a feature of both the reported beliefs in the sample and the empirical distribution. The most important criteria is that the beliefs (i.e., prior and posterior distributions) of the distribution parameters

[19]The parameters of the log-normal distribution this is done by using the following equations for the mean and CDF of a log-normal random variable: $\mathbf{E}[H] = exp(\mu + \frac{\sigma^2}{2})$, and $F_p(p_L; \mu, \sigma) = \Phi(\frac{log(p)-\mu}{\sigma})$, where Φ is the standard normal CDF.

are conjugate distributions, which is necessary for tractably modeling a Bayesian learning process. The results are similar under alternative distributional assumptions.[20]

Assuming that individuals believe that the distribution of prices is log-normal, then individuals' prior distribution on μ and $\frac{1}{\sigma^2}$ is Normal-Gamma. If the individual searches, she observes a price which she uses to update her beliefs. As the number of price observations increases, the individual's mean prior on μ and σ converge to the true parameters, and the variance of the priors converge to zero. In terms of the search problem, evidence of individual learning requires that individuals with more experience in the market (i.e., more observations of prices) have more accurate beliefs about the true parameters of the price distribution and more certainty in their beliefs.

Denote the individual's expected parameters as μ_i and σ_i. In the analysis, we consider the distribution of the individual's expected parameters in the population. Define $\bar{\mu}^e$ and $\bar{\sigma}^e$ as the mean of individuals' beliefs with the same level of experience e (i.e. $\bar{\mu}^e = \frac{1}{N^e}\sum_i \mu_i * \mathbf{1}[e_i = e]$ and $\bar{\sigma}^e = \frac{1}{N^e}\sum_i \sigma_i * \mathbf{1}[e_i = e]$). Similarly, define $Var[\mu]^e$ and $Var[\sigma]^e$ as the variance among individuals' beliefs with experience level e. As the number of signals the individual receives increases, the expected parameters should converge to the true price distribution parameters. Since each individual's beliefs converge to the true parameters, $\bar{\mu}^e$ and $\bar{\sigma}^e$ should converge to the true parameters as e increases. The convergence of each individual's beliefs to the true parameters as experience increases implies that the variance among individuals' beliefs goes to zero. However, at low levels of experience, $Var[\mu]^e$ and $Var[\sigma]^e$ may increase depending on the variance among individuals' initial prior beliefs. If individuals have similar initial mean priors, then the signal noise would generate greater dispersion of individuals' beliefs for low levels of experience.

Table 3.8 reports the summary statistics for the reported probability that a draw from the price distribution is below some threshold for different levels of the threshold. The threshold is defined as a fraction of the individual's reported expectation. On average, individuals report that the likelihood of the lowest price being less than 85% of their expected lowest price is 0.298. For higher levels of the threshold, individuals assign a larger probability that the price is below the threshold.

[20]The results for the gamma and normal distributions are presented in the Appendix.

Some individuals report a probability of zero or 100 which cannot be justified given the distributional assumption. Similarly, reported probabilities close to zero or 100 will only fit the distribution for extreme values of the parameters. Once the parameter values are calculated, individuals with parameter values in the top or bottom 2.5% of parameter values for either parameter are dropped from the sample to reduce the impact of outliers.[21]

In order to make the interpretation of the results more straightforward, we use the individual's distribution parameters to calculate the mean and standard deviation of the individual's expected price distribution, which is defined as the distribution with the individual's expected parameter values.[22] Table 3.9 reports the sample mean and standard deviation of these moments of the individual's expected price distribution by level of experience. Differences in the mean values from the analysis in the previous section is due to the different samples that result from the different rejection criteria. The mean and standard deviation of the preferred specifications of the empirical distribution are presented for comparison.

For used books, the variability of the mean across individuals with the same level of experience does not decrease for individuals with the highest level of experience. So there is greater variability in the expected lowest price for individuals with the highest level of experience. One reason for the greater variability for the highest category of experience is that there may be greater variability in the underlying level of experience for individuals in this group since it includes a larger range of the number of previous textbook purchases. The mean of the standard deviation of the expected price distribution initially increases with experience (from 0.238 for individuals with no online purchases to 0.250 for individuals with 1 to 5 online purchases) and then decreases with experience for higher levels of experience. The variability of the standard deviation of the expected price distribution across individuals with the same level of experience tends to decrease for higher levels of experience,

[21] Probabilities of zero and 100 are replaced with 1 and 99 respectively. The individuals who report probabilities of zero or 100 are included in the 2.5%. For the normal distribution, the initial sample only includes individuals who report a probability less than 50%.

[22] Note that individuals' two responses for each scenario exactly identify their expectations of the mean and variance of the normalized price distribution.

which is consistent with learning. The significance levels reported for the mean are from a two-sample equality of means test that compares individuals within a particular experience group to everyone not in that group. The test for the equality of variances defines the comparison samples in the same way. For new books, the standard deviation of the expected price distributions and its variability within experience groups display similar patterns as for used books. However, the significance of these patterns is less.

Comparing the beliefs about the standard deviation of the price distribution to the empirical standard deviation suggests that individuals may underestimate the variability of prices for used books and overestimate the variability of prices for new books. For used books, however, individuals' beliefs about the mean of the standard deviation decrease at higher levels of experience, moving farther away from the empirical standard deviation. There are several possible explanations for this result. First, our construction of the empirical price distribution may overstate the variability of prices by including erroneous product listings (e.g., sellers listing old editions or international editions). Also, the empirical distribution we construct may not be representative of the books encountered by the typical student since we include all books that are assigned at the university. Another explanation is the inherent difficulty in eliciting beliefs about variance as individuals may not be accustomed to thinking in probabilistic terms.

Figure 3.2 shows the density function of the log-normal distribution for the mean of the individual parameter values as well as the empirical distribution. Moving from the group with no experience to the group with some experience (1 to 5 online textbook purchases), the price distribution shifts to the left and the variance increases slightly. The distributions for higher levels of experience are similar to the group with some experience but have lower variance. This is in contrast to the empirical distribution, which displays much more variability than the beliefs. Although individuals with experience are more accurate in predicting the mean of the distribution, even high experience individuals tend to place too little weight in the left tail of the price distribution. Figure 3.3 shows the densities for new books. As experience increases, the variance of the distributions decreases. Similar to used books, individuals tend to understate the variability of the empirical distribution but to a lesser degree.

Overall, the evidence is consistent with learning, although the evidence suggests incomplete learning. It may be the case that individuals are only learning over a single parameter. This would explain why individuals with more experience are better able to predict the mean price, but are no better (and are actually worse for used books) in incorporating the variance of the price distribution into their beliefs. Another possibility is that individuals in the sample do not have sufficient experience for the convergence properties of the learning process to be evident.

3.4.3 Price Beliefs by Major

In this section, we test whether there are differences in individuals' beliefs for STEM majors and non-STEM majors. The STEM majors include the natural sciences, math, and other quantitative fields (including Economics). Individuals with multiple majors are categorized as STEM majors if any of their majors are in a STEM field. Table 3.10 reports the average expected price for STEM and non-STEM majors. For both new and used books, there is not a significant difference between the price expectations for individuals with no prior online purchases. For used books, this difference becomes significant at low levels of experience as the price expectations of STEM majors decreases at a faster rate. At higher levels of experience, the price expectations of non-STEM majors appears to "catch up" to the price expectations of STEM majors and the difference is significant at the 10% level. For new books, the difference in price expectations between STEM and non-STEM majors is only significant at the highest level of experience.

Table 3.11 presents the average standard deviation of the expected price distribution for STEM and non-STEM majors by level of experience. STEM majors tend to have lower expectations about the variability of prices and there is little change in the expected price variation across different levels of experience. For non-STEM majors, the variation in the expected price distribution initially increases at the lowest level of experience and decreases at the higher levels of experience. This pattern holds for both new and used books.

The results suggest that individuals in non-STEM majors may incorporate new information about the price distribution differently from STEM majors. The mean of the expected price distribution is higher than the mean of the empirical distribution for individuals with no online purchasing experience regardless of major. If these individuals with no experience search for a textbook online, they are likely to observe a price that is lower than the mean of their expected price distribution. On average, STEM majors incorporate this initial experience by lowering the mean of their expected price distribution while non-STEM majors increase the variance of their expected price distribution.

The results from this section should be interpreted with some caution as there are other factors that may cause the reported beliefs about the price distribution to differ by major. First, STEM majors may more comfortable answering the kind of probabilistic questions that we ask in the survey. Second, the types of books purchased may be systematically different.

3.5 Learning vs. Selection

Although the evidence is consistent with learning, the differences in individuals' beliefs across levels of experience could also result from selection. If individuals have heterogeneous initial prior beliefs, then individuals who believe that the online price is similar to the bookstore price will not search and will not purchase their books online. Then, if the individuals whose initial priors are close to the true distribution are the ones who search and purchase online, the observed difference in beliefs would be the result of selection based on the initial difference in beliefs and not because of learning.

To distinguish between the effects of learning and selection, we examine the individuals who participated in the survey in both the fall and spring semesters. There were 240 individuals who participated in both surveys. Of these individuals, 89 reported an increase in their level of online textbook purchasing experience from the fall to the spring survey. If selection is generating the observed patterns in the data, then the individuals who report an increase in experience in the spring would have lower expected online prices in the fall than the individuals who do not have an increase in experience. Alternatively, in order for the data to be consistent with learning, then individuals who report an increase in experience should be more likely to report different beliefs in the spring, whereas the beliefs of individuals who do not report an increase in experience should be similar in both periods. For the prior online purchase experience measure, we restrict the analysis to the 22 individuals (between 47 and 56 scenarios) who report no experience in the fall survey. Since this measure of experience is an interval, individuals who remain in the same interval for both fall and spring may or may not have gained experience.[23] The inherent limitation of this test is that one period of learning may not generate a significant difference in beliefs for those whose experience increased. Therefore, this test is primarily a test of the hypothesis of no selection.

Table 3.12 reports the mean parameter values for a log-normal distribution of prices for the two groups for both surveys as well as the mean change in parameter values between surveys. The results of the two-sample t-tests comparing each of the mean values between groups are also reported. There is not a significant difference between the mean parameter values of the two groups in the fall semester for both new and used books. The only difference that is significant is the difference in the value of σ for used books in the fall compared to the spring. However, this change is significant for both groups. These results suggest that selection is not the primary cause of the differences in beliefs across experience levels. However, due to the limited sample size, no definitive conclusion can be drawn.

3.6 Conclusion

Although the evidence is consistent with learning, it appears that the learning process is incomplete. Even individuals with the highest levels of experience on average do not fully converge to the empirical distribution. Also, many individuals with high levels of experience have inaccurate beliefs (i.e., the variation across individuals' beliefs does not converge to zero). There are three primary explanations for the persistence of inaccurate beliefs. The first is that the level of experience where this convergence would occur is beyond what we measure in the data. The second is that the beliefs are converging to a distribution other than what is observed during the sample period. For much of the year, the prices of these textbooks online are relatively stable. For a few weeks prior to the start of the semester, prices rise sharply and peak around the first week of the semester. Since individuals are likely to purchase textbooks during this period, the signal that they receive will be from a distribution with a higher mean than what is observed during the sample period. If an individual only ever purchases books online during the first week of the semester (the time when online prices are greatest), then a high experience individual may expect that potential savings online are relatively

[23]The results of the test are the same if the sample is not limited to individuals who report no experience in the fall.

modest. Finally, there is likely to be some noise in the reported data as individuals may have differed in their interpretation of questions as well as the amount of consideration given to their responses.

One limitation of this analysis is the problem of external validity. Although the online market for textbooks is comparable to online markets for other goods, the individuals in our sample are not representative of consumers in other online markets. Relative to consumers in other online markets, individuals in our sample are likely to be younger with higher intellectual ability, but they may have less overall experience in online markets. If there are knowledge spillovers across online markets, i.e. if experience in one online market causes individuals to have more accurate beliefs about the prices in other online markets, then the observed bias in the initial beliefs may be more pronounced in the online textbook market, where individuals are likely to have less overall experience in online markets.

In this paper we use a novel dataset to examine subjective price beliefs and their relationship with experience in a market. We find that inexperienced consumers have biased beliefs, but that consumers appear to be learning about the empirical price distribution as they repeatedly participate in the market. This study also leaves open a wide avenue for future research. First, since we do not estimate a dynamic model of search and learning, we are not able to show how individuals incorporate their beliefs into the search decision. Thus, we are not able to determine whether individuals incorporate the benefits of the additional information obtained through search for future purchasing decisions in their decision to search. Also, if individuals have heterogeneous initial prior beliefs, one potential avenue of future research would be to determine the sources of this heterogeneity. Finally, future research is needed to justify the distributional assumptions on the empirical distribution as well as the prior beliefs.

Figure 3.1: Textbook Purchasing Scenario

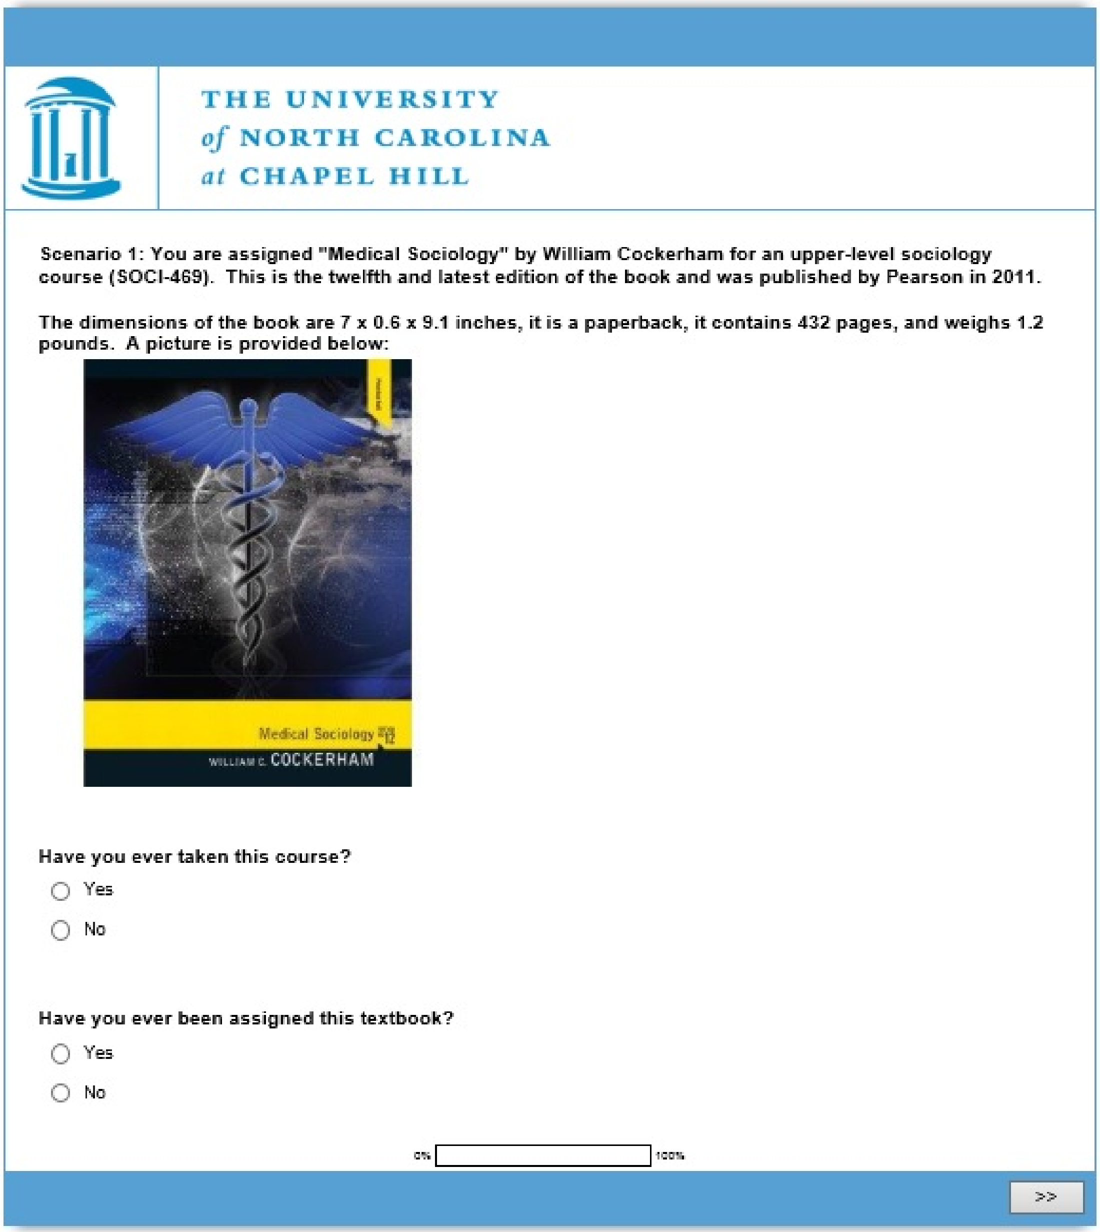

THE UNIVERSITY
of NORTH CAROLINA
at CHAPEL HILL

Scenario 1: You are assigned "Medical Sociology" by William Cockerham for an upper-level sociology course (SOCI-469). This is the twelfth and latest edition of the book and was published by Pearson in 2011.

The dimensions of the book are 7 x 0.6 x 9.1 inches, it is a paperback, it contains 432 pages, and weighs 1.2 pounds. A picture is provided below:

Have you ever taken this course?

- Yes
- No

Have you ever been assigned this textbook?

- Yes
- No

0% 100%

>>

Table 3.1: Respondent Characteristics

	Proportion
1 - 2 Semesters	0.143
3 - 4 Semesters	0.248
5 - 6 Semesters	0.242
7 or More Semesters	0.369
Ever Purchased at Campus Bookstore	0.960
Never Purchased Online	0.106
Purchased 1 - 5 Online	0.230
Purchased 6 - 10 Online	0.241
Purchased 11 or More Online	0.423
STEM Major	0.280
Economics Major	0.089
Other Major	0.631
N = 1465	

Table 3.2: Ratio of Amazon Prices to Bookstore Prices by Survey Period

		Mean Ratio	S.D.	Min	Median	Max	Mean Diff. ($)	N
Fall Survey Period								
All Books	New	0.857	0.156	0.467	0.850	1.432	9.37	2051
	New_{alt}	0.772	0.236	0.206	0.753	2.877	20.05	2220
	Used	0.758	0.383	0.091	0.715	4.207	16.24	2129
Bookstore Price > $100	New	0.825	0.142	0.475	0.829	1.222	28.33	405
	New_{alt}	0.678	0.169	0.206	0.673	1.377	57.29	429
	Used	0.657	0.246	0.097	0.659	1.348	46.03	390
Scenario Books	New	0.788	0.112	0.600	0.802	0.961	31.07	11
	New_{alt}	0.659	0.136	0.514	0.609	0.861	61.48	11
	Used	0.609	0.241	0.151	0.578	0.979	45.81	11
Spring Survey Period								
		Mean Ratio	S.D.	Min	Median	Max	Mean Diff. ($)	N
All Books	New	0.834	0.148	0.357	0.838	1.425	10.99	2023
	New_{alt}	0.755	0.295	0.202	0.737	3.970	21.64	2248
	Used	0.735	0.441	0.080	0.684	5.675	18.57	2161
Bookstore Price > $100	New	0.795	0.158	0.358	0.815	1.326	34.02	379
	New_{alt}	0.646	0.219	0.216	0.658	1.804	62.39	434
	Used	0.597	0.279	0.085	0.602	1.731	54.57	390
Scenario Books	New	0.729	0.120	0.524	0.768	0.879	46.57	11
	New_{alt}	0.607	0.169	0.275	0.593	0.853	66.80	11
	Used	0.614	0.255	0.123	0.581	0.919	50.26	11

Notes: The ratio reported is the lowest price on Amazon.com divided by the price of the same title (of equivalent quality) from the campus bookstore. New_{alt} refers to the lowest price listed by marketplace sellers for a new copy of the title.

Table 3.3: Mean Ratio Comparisons by Online Purchasing Experience, Used Books

Fall						
	Expectation / Bookstore Price				Mean Expectation Bias	
Experience	N	Mean Ratio	S.D.	Median	Scenario Books	Books > $100
No online purchases	256	0.735	0.204	0.750	0.126	0.078***
1-5 online purchases	478	0.714	0.191	0.744	0.105	0.057***
6-10 online purchases	477	0.663	0.172	0.683	0.054	0.006
11+ online purchases	810	0.645	0.184	0.645	0.036	-0.012
Spring						
	Expectation / Bookstore Price				Mean Expectation Bias	
Experience	N	Mean Ratio	S.D.	Median	Scenario Books	Books > $100
No online purchases	182	0.744	0.198	0.761	0.130	0.147***
1-5 online purchases	439	0.710	0.170	0.745	0.097	0.113***
6-10 online purchases	480	0.703	0.173	0.741	0.089	0.106***
11+ online purchases	888	0.660	0.176	0.675	0.046	0.063***

Notes: The ratio reported is an individual's expectation of the lowest price from an online retailer divided by the price of the same title from the campus bookstore. Expectation Bias refers to the difference between this ratio and the ratio of the observed online price to the bookstore price.

* refers to t–test p–values $< .1$; ** $< .05$; *** $< .01$; $H_0 =$ No difference between ratios.

Table 3.4: Mean Difference Comparisons by Online Purchasing Experience, Used Books

Fall						
	Bookstore Price - Expected Price				Mean Expectation Bias	
Experience	N	Mean Diff.	S.D.	Median	Scenario Books	Books > $100
No online purchases	256	31.55	29.86	28	14.26	14.48***
1-5 online purchases	478	35.02	30.51	27.5	10.79	11.01***
6-10 online purchases	477	40.37	30.79	31	5.44	5.66
11+ online purchases	810	41.73	30.78	31	4.08	4.30
Spring						
	Bookstore Price - Expected Price				Mean Expectation Bias	
Experience	N	Mean Diff.	S.D.	Median	Scenario Books	Books > $100
No online purchases	182	32.16	29.52	27.5	18.10	22.41 ***
1-5 online purchases	439	35.22	28.56	28	15.04	19.33 ***
6-10 online purchases	480	35.97	27.75	28.5	14.29	18.59 ***
11+ online purchases	888	41.28	31.09	31	8.98	13.29 ***

Notes: The difference reported is an individual's expectation of the lowest price from an online retailer subtracted from the price of the same title from the campus bookstore. Expectation Bias refers to the difference between this difference and the difference of the observed bookstore price to the online price.

* refers to t–test p–values $< .1$; ** $< .05$; *** $< .01$; $H_0 =$ No difference between differences.

Table 3.5: Mean Ratio Comparisons by Online Purchasing Experience, New Books

Fall								
	Expectation / Bookstore Price				Mean Expectation Bias			
					Scenario Books		Books > \$100	
Experience	N	Mean Ratio	S.D.	Median	New	New_{alt}	New	New_{alt}
No online purchases	256	0.834	0.164	0.853	0.046	0.174***	0.008	0.156***
1-5 online purchases	479	0.819	0.155	0.851	0.032	0.160***	-0.006	0.142***
6-10 online purchases	480	0.778	0.171	0.817	0.010	0.119**	-0.048***	0.100***
11+ online purchases	814	0.768	0.159	0.798	-0.020	0.109**	-0.058***	0.090***
Spring								
	Expectation / Bookstore Price				Mean Expectation Bias			
					Scenario Books		Books > \$100	
Experience	N	Mean Ratio	S.D.	Median	New	New_{alt}	New	New_{alt}
No online purchases	184	0.835	0.183	0.870	0.103**	0.225***	0.037**	0.186***
1-5 online purchases	444	0.818	0.150	0.856	0.089**	0.210***	0.023**	0.171***
6-10 online purchases	486	0.788	0.152	0.822	0.059	0.180***	-0.008	0.141***
11+ online purchases	892	0.772	0.155	0.795	0.043	0.164***	-0.023**	0.126***

Notes: The ratio reported is an individual's expectation of the lowest price from an online retailer divided by the price of the same title from the campus bookstore. Expectation Bias refers to the difference between this ratio and the ratio of the observed online price to the bookstore price.

* refers to t–test p–value $< .1$; ** $< .05$; *** $< .01$; $H_0 =$ No difference between ratios.

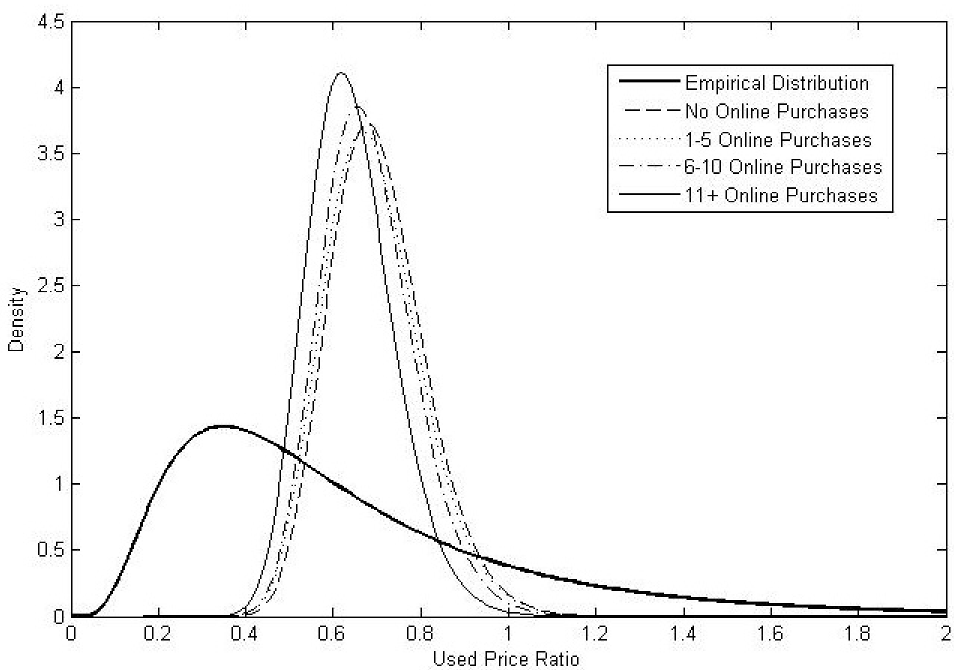

Figure 3.2: Used price pdf versus empirical dist. by level of experience

Table 3.6: (Online Expectation / Bookstore Price) Regressed on Prev. Purchases

	(1) New	(1) Used	(2) New	(2) Used
1 - 5 Online Purchases	-0.014	-0.026	-0.015	-0.028*
	(0.015)	(0.017)	(0.014)	(0.017)
6 - 10 Online Purchases	-0.051***	-0.058***	-0.052***	-0.059***
	(0.015)	(0.017)	(0.014)	(0.017)
11+ Online Purchases	-0.063***	-0.087***	-0.063***	-0.088***
	(0.014)	(0.016)	(0.014)	(0.016)
Previously Taken Course	-0.037***	-0.027**	-0.022*	-0.025*
	(0.012)	(0.013)	(0.012)	(0.013)
Previously Assigned Book	0.006	-0.005	-0.008	-0.017
	(0.016)	(0.017)	(0.016)	(0.017)
Introductory Course	-0.009	0.003	.	.
	(0.006)	(0.008)	.	.
Latest Edition	0.007	0.028**	-0.006	-0.003
	(0.012)	(0.012)	(0.015)	(0.016)
Years Since Last Revision	-0.001*	-0.004***	-0.002*	-0.002**
	(0.001)	(0.001)	(0.001)	(0.001)
Hardback	-0.009	-0.005	0.004	-0.016
	(0.010)	(0.011)	(0.011)	(0.013)
Book Fixed Effects	No		Yes	

Notes: Clustered standard errors (on the individual) given in parentheses. Also included: full information indicator, Spring indicator, pages, and weight.

* refers to p–value < .1; ** < .05; *** < .01

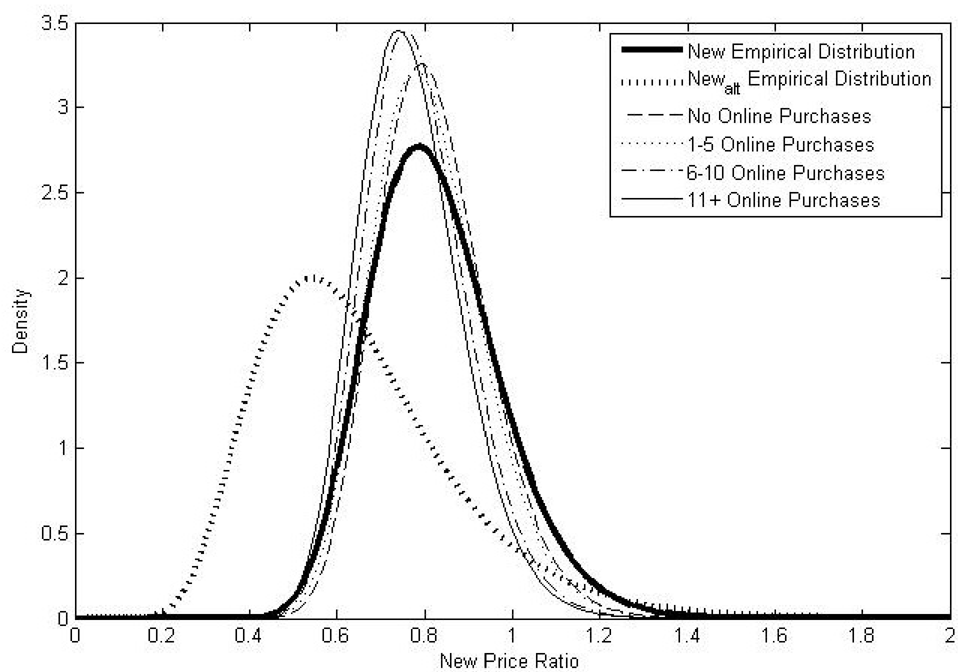

Figure 3.3: New price pdf versus empirical dist. by level of experience
New$_{\text{alt}}$ is the minimum price for a new textbook from Amazon or Amazon Marketplace.

Table 3.7: (Online Expectation / Bookstore Price) Regressed on Experience

	(1)		(2)	
	New	Used	New	Used
1 - 5 Online Purchases	-0.019	-0.026	-0.020	-0.027
	(0.014)	(0.017)	(0.014)	(0.017)
6 - 10 Online Purchases	-0.058***	-0.054***	-0.060***	-0.058***
	(0.015)	(0.016)	(0.014)	(0.016)
11+ Online Purchases	-0.076***	-0.084***	-0.077***	-0.087***
	(0.014)	(0.016)	(0.014)	(0.016)
Second Year	0.018	-0.018	0.019	-0.013
	(0.013)	(0.015)	(0.013)	(0.015)
Third Year	0.026*	-0.012	0.026*	-0.007
	(0.013)	(0.015)	(0.014)	(0.015)
Four and Above	0.040***	-0.005	0.040***	-0.002
	(0.016)	(0.015)	(0.013)	(0.015)
Additional Controls	No		Yes	

Notes: Clustered standard errors (on the individual) given in parentheses. Additional controls include scenario f.e.s, full information ind., Spring indicator, pages, weight, prev. taken, prev. assigned, latest edition ind., hardback ind. and years since revision. Second year denotes an indicator for individuals in their 3^{rd} or 4^{th} semester, etc.

* refers to p–value $<$ .1; ** $<$ .05; *** $<$ .01

Table 3.8: Reported Probability that Lowest Price $< b *$ Expected Lowest Price

Used						
b	N	Mean	S.D.	Min	Median	Max
0.85	1352	0.298	0.187	0	0.250	0.95
0.90	1368	0.328	0.189	0	0.300	1.00
0.95	1283	0.359	0.198	0	0.300	1.00
New						
b	N	Mean	S.D.	Min	Median	Max
0.85	1359	0.271	0.176	0	0.250	1.00
0.90	1376	0.312	0.190	0	0.300	1.00
0.95	1293	0.339	0.195	0	0.300	1.00

Table 3.9: Mean and Variance Comparisons (Log-Normal Assumption)

Used					
Empirical Distribution					
	N	Mean	S.D.		
$p_{Bookstore} > \$100$	390	0.657	0.246		
Scenario Books	11	0.609	0.241		
Beliefs					
Experience	N	Mean $\mathbf{E}_i(p)$	S.D. $\mathbf{E}_i(p)$	Mean $\sqrt{\mathbf{Var}_i(p)}$	S.D. $\sqrt{\mathbf{Var}_i(p)}$
No online purchases	370	0.713***	0.166*	0.238	0.384
1-5 online purchases	819	0.703***	0.166	0.250**	0.362**
6-10 online purchases	884	0.682	0.167*	0.230	0.332
11+ online purchases	1577	0.652***	0.171***	0.204***	0.284***
New					
Empirical Distribution					
	N	Mean	S.D.		
Bkstr. Price $> \$100$	405	0.825	0.142		
Scenario Books	11	0.788	0.112		
Beliefs					
Experience	N	Mean $\mathbf{E}_i(p)$	S.D. $\mathbf{E}_i(p)$	Mean $\sqrt{\mathbf{Var}_i(p)}$	S.D. $\sqrt{\mathbf{Var}_i(p)}$
No online purchases	371	0.813***	0.151	0.232	0.300
1-5 online purchases	815	0.809***	0.145**	0.246**	0.346*
6-10 online purchases	872	0.783	0.151	0.222	0.304
11+ online purchases	1535	0.769***	0.149	0.215*	0.289

Notes: The significance levels reported for the mean values are from a two-sample equality of means test. The significance levels for the standard deviations are from Brown and Forsythe's alternative formulation of Levene's robust two-sample equality of variances test.

* refers to p–value $< .1$; ** $< .05$; *** $< .01$

Table 3.10: Mean Comparison by Major

Used Books					
	STEM majors		Non-STEM majors		
Experience	N	Mean $\mathbf{E}_i(p)$	N	Mean $\mathbf{E}_i(p)$	p-value
No online purchases	117	0.703 (0.167)	252	0.718 (0.166)	0.429
1-5 online purchases	314	0.681 (0.174)	505	0.716 (0.160)	0.004
6-10 online purchases	323	0.667 (0.175)	560	0.692 (0.161)	0.036
11+ online purchases	598	0.642 (0.171)	979	0.657 (0.171)	0.093
New Books					
No online purchases	118	0.820 (0.143)	259	0.811 (0.154)	0.568
1-5 online purchases	320	0.801 (0.149)	508	0.814 (0.144)	0.245
6-10 online purchases	321	0.787 (0.150)	563	0.779 (0.155)	0.426
11+ online purchases	596	0.754 (0.148)	966	0.777 (0.152)	0.003

Notes: The p-value is from a two-sample equality of means test.
Standard deviation in parentheses.

Table 3.11: Variance Comparison by Major (Log-Normal Assumption)

Used Books					
	STEM majors		Non-STEM majors		
Experience	N	Mean $\sqrt{\mathbf{Var}_i(p)}$	N	Mean $\sqrt{\mathbf{Var}_i(p)}$	p-value
No online purchases	117	0.220 (0.387)	252	0.248 (0.384)	0.528
1-5 online purchases	314	0.207 (0.224)	505	0.276 (0.424)	0.003
6-10 online purchases	323	0.215 (0.292)	560	0.240 (0.353)	0.249
11+ online purchases	598	0.209 (0.292)	979	0.202 (0.279)	0.642
New Books					
No online purchases	118	0.203 (0.290)	259	0.252 (0.336)	0.153
1-5 online purchases	320	0.213 (0.292)	508	0.278 (0.421)	0.009
6-10 online purchases	321	0.231 (0.326)	563	0.222 (0.322)	0.699
11+ online purchases	596	0.201 (0.270)	966	0.225 (0.316)	0.114

Notes: The p-value is from a two-sample equality of means test.
Standard deviation in parentheses.

Used Books								
	Fall			Spring			Difference	
Group	N	mean μ_i	mean σ_i	N	mean μ_i	mean σ_i	mean μ_i	mean σ_i
Increase Exp.	48	-0.447 (0.321)	0.297 (0.307)	47	-0.412 (0.356)	0.199 (0.169)	0.034	-0.098*
Same Exp.	55	-0.479 (0.363)	0.342 (0.372)	54	-0.446 (0.299)	0.230 (0.255)	0.034	-0.111*
Difference		0.033	-0.045		0.034	-0.032		

New Books								
	Fall			Spring			Difference	
Group	N	mean μ_i	mean σ_i	N	mean μ_i	mean σ_i	mean μ_i	mean σ_i
Increase Exp.	45	-0.243 (0.163)	0.216 (0.148)	50	-0.294 (0.274)	0.240 (0.276)	0.051	-0.024
Same Exp.	56	-0.302 (0.289)	0.277 (0.279)	53	-0.288 (0.217)	0.246 (0.221)	-0.014	0.031
Difference		0.059	-0.060		-0.005	-0.006		

Table 3.12: Parameter Values by Change in Experience

APPENDIX A

APPENDIX FOR EXPLAINING YOUTH SMOKING INITIATION IN THE CONTEXT OF A RATIONAL ADDITION MODEL WITH LEARNING

A.1 Data Appendix

Table A.1 presents the summary statistics by year for the sample of individuals who are observed in every wave of the survey. The proportion of individuals who smoke increases over the first few waves. The proportion of smokers peaks at around 36% in 2002 and remains in the low 30's for the rest of the sample period.

A.2 Estimation Appendix

A.2.1 CCP Representation and Finite Dependence

When the preference shock is GEV, the future value term in the conditional value function has a closed form solution. With type I EV errors, the future value term can be expressed as the one period ahead CCP and conditional value function of any alternative. The closed form expression of the future value term is:

$$\mathbf{E}[\max_j v^j_{t+1}] = log(\sum_{k=1}^{J} e^{v^k_{t+1}}) + e.c. \tag{A.1}$$

where $e.c.$ is Euler's constant.[1] To express the future value term in terms of the conditional value function and CCP of alternative 1, consider the probability of choosing alternative 1 in period $t+1$:

$$P^1_{t+1} = \frac{e^{v^1_{t+1}}}{\sum_{k=1}^{J} e^{v^k_{t+1}}} \tag{A.2}$$

Now, take the log of both sides:

$$log(P^1_{t+1}) = v^1_{t+1} - log(\sum_{k=1}^{J} e^{v^k_{t+1}}) \tag{A.3}$$

[1] Individual subscripts are suppressed for simplicity.

Table A.1: Summary statistics by year for individuals observed every period ($N = 5,385$)

	Year					
	1997		1998		1999	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Ever smoked	0.373	0.484	0.473	0.499	0.532	0.499
Current smoker	0.164	0.370	0.239	0.427	0.279	0.448
# of cigs/day	0.560	2.572	1.274	4.566	1.674	5.137
Age	14.23	1.472	15.87	1.430	16.84	1.439
Employed	0.441	0.497	0.501	0.500	0.522	0.500
Income	243.5	749.5	626.3	1,720	1,189	3,168
Married	0.000	0.019	0.004	0.064	0.013	0.112
	2000		2001		2002	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Ever smoked	0.579	0.494	0.618	0.486	0.648	0.478
Current smoker	0.316	0.465	0.335	0.472	0.361	0.480
# of cigs/day	2.122	5.616	2.451	6.168	2.703	6.535
Age	17.91	1.435	18.90	1.427	19.90	1.402
Employed	0.611	0.488	0.697	0.459	0.750	0.433
Income	2,191	4,565	3,785	6,281	4,877	7,702
Married	0.026	0.159	0.052	0.222	0.074	0.262
	2003		2004		2005	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Ever smoked	0.697	0.460	0.668	0.471	0.684	0.465
Current smoker	0.350	0.477	0.361	0.480	0.357	0.479
# of cigs/day	2.936	7.019	2.772	6.349	2.880	6.631
Age	22.85	1.420	20.85	1.424	21.88	1.418
Employed	0.830	0.376	0.785	0.411	0.804	0.397
Income	10,832	13,255	6,327	8,995	8,485	11,807
Married	0.108	0.310	0.143	0.350	0.182	0.386
	2006		2007		2008	
Variable	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Ever smoked	0.709	0.454	0.717	0.450	0.722	0.448
Current smoker	0.354	0.478	0.343	0.475	0.328	0.470
# of cigs/day	2.818	6.462	2.740	6.488	2.760	6.827
Age	23.81	1.420	24.77	1.432	25.78	1.426
Employed	0.854	0.353	0.870	0.336	0.867	0.340
Income	14,024	15,555	17,402	18,905	20,711	20,735
Married	0.217	0.412	0.248	0.432	0.277	0.447

Substituting the log-sum term into the future value term gives the CCP representation of the future value term:

$$\mathbf{E}_t[\max_j v^j_{t+1}] = u^1_{t+1} + \mathbf{E}V_{t+2} - log(P^1_{t+1}) + e.c. \tag{A.4}$$

When forming the choice probabilities in the likelihood function, all that matters is the difference in conditional value functions. Finite dependence occurs when two sequences of choices lead to the same future state in expectation. Then when taking the difference in conditional value functions, the remaining future value terms in the CCP representation cancel. The state variables are the individuals beliefs and the prior period's decision. The expectation in the current period of future mean priors is simply the mean of the current period priors for any future sequence of signals (i.e., $\mathbf{E}_t[m_{t+k}] = m_t,\ \forall k$ and $\cup_j \{d^j_{t+1}, \ldots, d^j_{t+k-1}\}$). The variance of the priors only depends on the number and intensity of the signals; the timing of the signals does not matter. So the expected distribution of a future period's beliefs will be the same along any two sequences that generate the same number and intensity of the signals. The other state variable is the prior period's decision, which will be the same as long as the two sequences end with the same alternative. The following table gives the sequences that generate finite dependence.

period	$t-1$	t	$t+1$	$t+2$	$t+3$
sequence 1	0	a_j	0	0	
sequence 2	0	0	a_j	0	
	For any $a_j > 0$				
sequence 1	$a_{j'}$	a_j	1	$a_{j'}$	0
sequence 2	$a_{j'}$	0	$a_{j'}$	a_j	0
	For any $a_j, a_{j'} > 0$				

Consider the simpler case, which is when the individual did not smoke in the prior period. The conditional value function in period t for any alternative $j > 1$ and $j = 1$ are:

$$\begin{aligned} v^j(d_{t-1}^1 = 1, \Gamma_t) &= u^j(m_t) + \beta * \mathbf{E}[V(\Gamma_{t+1}|d_t^j = 1)] \\ v^1(d_{t-1}^1 = 1, \Gamma_t) &= \beta * \mathbf{E}[V(\Gamma_{t+1})|d_t^1 = 1] \end{aligned} \tag{A.5}$$

The CCP representation of the future value term in the conditional value function for alternative $j > 1$ is:

$$\begin{aligned} \mathbf{E}[V(\Gamma_{t+1}|d_t^j = 1)] = u^1(\mathbf{E}[m_{t+1}|d_t^j = 1]) - \mathbf{E}_m[log(P^1(m_{t+1}|d_t^j = 1))] \\ + \beta u^1(\mathbf{E}[m_{t+2}|d_t^j = 1, d_{t+1}^1 = 1]) - \beta\mathbf{E}_m[log(P^1(m_{t+2}|d_t^j = 1, d_{t+1}^1 = 1))] \\ + \beta^2\mathbf{E}[V(\Gamma_{t+3}|d_t^j = 1, d_{t+1}^1 = 1, d_{t+2}^1 = 1)] \end{aligned} \tag{A.6}$$

and the CCP representation of the future value term in the conditional value function for alternative $j = 1$ is:

$$\begin{aligned} \mathbf{E}[V(\Gamma_{t+1})|d_t^1 = 1] = u^j(\mathbf{E}[m_{t+1}|d_t^1 = 1]) - \mathbf{E}_m[log(P^j(m_{t+1}|d_t^1 = 1))] \\ + \beta u^1(\mathbf{E}[m_{t+2}|d_t^1 = 1, d_{t+1}^j = 1]) - \beta\mathbf{E}_m[log(P^1(m_{t+2}|d_t^1 = 1, d_{t+1}^j = 1))] \\ + \beta^2\mathbf{E}[V(\Gamma_{t+3}|d_t^1 = 1, d_{t+1}^j = 1, d_{t+2}^1 = 1)] \end{aligned} \tag{A.7}$$

When calculating the choice probability in the likelihood function, all that matters is the difference between these conditional value functions. The $t + 3$ expected future value term is the same for the alternative $j > 1$ and $j = 1$ conditional value functions, so it will cancel out in the difference term. All that remains are the flow utilities for periods t, $t + 1$, and $t + 2$ as well as CCPs for periods $t + 1$ and $t + 2$. Note that the CCPs are functions of the mean prior beliefs (m), which depend on the realized value of the signal. If no signal is received, then the CCP can be evaluated using the current period beliefs. If, however, a signal is received, then calculating the expectation requires integrating over possible realizations of the signal. Approximating these integrals numerically adds to the computational burden of the estimation procedure, but the computational requirements are much less than would be needed to fully solve the dynamic learning problem.

An additional advantage of the CCP representation of the value function with finite dependence is that it is not necessary to estimate a separate closing function for the value function in the final period that the individuals are observed, which is the case when solving for the value function using backwards recursion.

The experimentation decision is an optimal stopping problem, which is one of the original class of problems where the CCP representation was applied. The conditional value function of experimenting is:

$$v_t^E = u_t^E + \mathbf{E}_t[V_t | d_t^E = 1] \tag{A.8}$$

In periods after the individual experiments, he no longer faces an experimentation decision. So, the conditional value function for experimenting does not include a future value term for future experimentation decisions. The expectation over the value of the consumption decision is over the iid preference shock as well as potential realizations of the value of α. The conditional value function for not experimenting is the discounted expected value of the next period's experimentation decision. The CCP representation of this future value term is:

$$\begin{aligned} v_{n,t}^{NE} &= \beta \mathbf{E}_t[V_{n,t+1}^E] \\ &= \beta \mathbf{E}_t[v_{n,t+1}^E - log(P_{n,t+1}^E) + ec] \\ &= \beta \mathbf{E}_t[u_{n,t+1}^E + \mathbf{E}_{t+1}[V_{n,t+1} | d_{n,t+1}^E = 1] - log(P_{n,t+1}^E) + ec] \end{aligned} \tag{A.9}$$

Note that the CCP representation does not contain a future value term for the experimentation decision for period $t+2$. Now both conditional value functions contain an expected value of the consumption decision in the period that the individual experiments. The expected value of the consumption decision in the conditional value function for experimenting can be expressed as a function of the CCP of not smoking and the conditional value function of not smoking:

$$\mathbf{E}_t[V_{n,t}|d^E_{n,t} = 1] = \mathbf{E}_t[u^1_{n,t} + \beta\mathbf{E}_t[V_{n,t+1}|d^1_{n,t} = 1] - log(P^1_{n,t}) + ec] \tag{A.10}$$

The future value term will cancel with the future expected value of the consumption decision in the conditional value function of not experimenting. Then the difference in the conditional value functions of not experimenting and experimenting is:

$$v^{NE}_{n,t} - v^E_{n,t} = -u^E_{n,t} + \mathbf{E}[log(P^1_{n,t})] - ec + \beta(u^E_{n,t} - \mathbf{E}[log(P^E_{n,t})] + ec) \tag{A.11}$$

The probability that an individual experiments is:

$$P^E_{n,t} = \frac{1}{1 + exp(v^{NE}_{n,t} - v^E_{n,t})} \tag{A.12}$$

A.2.2 Estimation Procedure

This section describes the details of the estimation procedure. The estimation procedure uses the EM algorithm to estimate the parameters that maximize the likelihood function (equation 1.20). The integrals in the likelihood function are approximated numerically, so the likelihood function becomes a simulated likelihood function in estimation.[2] The procedure begins with initial guesses for the parameters and the CCPs as well as M vectors of draws from the standard normal distribution for each individual, $\{z^m_n\}^M_{m=1}$. These draws are used to form a sample of $N*M$ simulated individuals. Each iteration proceeds according to the following steps:

1. Calculate the value of the unobserved state variables for each individual using the current estimates of the population distribution parameters and the M draws using the following equations

[2] When the EM algorithm is used to maximize a simulated expectation (the likelihood being maximized is the expected conditional likelihood), it is called a simulated EM (SEM) algorithm.

and the corresponding elements of the vector z:

$$\theta_n^m = \bar{\theta} + Ch' * z_n^m \tag{A.13}$$

$$\{\lambda_{n,t}^m\}_{t=1}^T = \sigma_\lambda * z_n^m,\ \{\psi_{n,t}^m\}_{t=1}^T = \sigma_\psi * z_n^m, \text{and } \{\eta_{n,t}\}_{t=1}^T = \sigma_\eta * z_n^m \tag{A.14}$$

where Ch is the Cholesky decomposition of Σ. These values of the addictive parameters and the noisy component of the signals are used to calculate the individual's prior beliefs for each period.

2. E step, part 1: Use the prior iteration parameter values and CCPs, denoted $\hat{P}$, to update π:

$$\pi(\theta_n^m, \Lambda_n^m) = \frac{\prod_t L_{n,t}(\theta_n^m, \Lambda_n^m, \hat{P})}{\sum_m \prod_t L_{n,t}(\theta_n^m, \Lambda_n^m, \hat{P})} \tag{A.15}$$

3. E step, part 2: Use the updated values of π to update the CCPs. There are several methods for updating the CCPs. The method used in this paper is to estimate a weighted multinomial logit model of the observed choices on a flexible polynomial of the state variables (both observed and unobserved), where the values of π are the weights. The coefficients from this multinomial logit are used in order to approximate the CCPs at the relevant combinations of state variables in the solution to the individual's problem. This method for updating the CCPs is analogous to least squares value function interpolation. The only heterogeneity in utility from experimentation is in observable characteristics. Therefore, the CCPs for the experimentation decision can be estimated outside of the main estimation routine. Similarly, state transition probabilities that do not depend on the unobserved heterogeneity, or that only depend on the unobserved heterogeneity through the smoking choice, can also be estimated in a first stage.

4. M step: Using the updated CCPs and π, the parameters are updated by maximizing the simulated log-likelihood function:

$$\tilde{\mathscr{L}}(\gamma, \xi, \sigma_\psi^2, \sigma_\lambda^2, \sigma_\eta^2, \bar{\theta}, \Sigma) = \sum_n \frac{1}{M} \sum_m \pi(\theta_n^m, \Lambda_n^m) \mathscr{L}(\theta_n^m, \Lambda_n^m, \hat{P}, \gamma, \xi | \sigma_\psi^2 \sigma_\lambda^2, \sigma_\eta^2, \bar{\theta}, \Sigma) \quad \text{(A.16)}$$

   The parameters of the population distribution of unobserved heterogeneity can be estimated separately and have a closed form solution (Train 2007). The updated parameters are simply the weighted mean (for $\bar{\theta}$) and variance (for Σ, σ_λ^2, and σ_η^2) of the values of θ_n^m and Λ_n^m, where the weights are the values of π. The remaining parameters are estimated using simulated maximum likelihood.

These steps are repeated until the parameters converge. The criteria for convergence can either be based on changes in the parameter values or changes in the likelihood function. In practice, there are a wide range of criteria used to determine the convergence of the SEM algorithm. Also, the algorithm may not converge to the global maximum, so to confirm any potential maximum, the algorithm must be rerun using different starting values. The convergence criteria used for preliminary estimation results are that the parameters change by less than one half of one percent, which is the criteria suggested by Train (2007). A feature of the EM algorithm is that the likelihood function weakly increases from one iteration to the next. Performing the full maximization in the M-step yields the largest possible increase in the likelihood but may be computationally intensive. The computational burden is particularly great if an the derivative and Hessian must be approximated using finite differences. In order to reduce the computational burden, I use an alternative version of the EM algorithm. This alternative version of the EM algorithm replaces the full optimization of the M-step, which gives the greatest possible likelihood improvement, with a procedure that is simply guaranteed to improve the likelihood function. This version of the EM algorithm is called a Generalized EM algorithm (GEM) and is commonly implemented by replacing the maximization in the M-step with a single Newton-Raphson iteration. GEM algorithms share similar convergence properties as the

EM algorithm, although they converge at a slower rate. Even though GEM algorithms require more iterations to converge, each iteration requires much fewer evaluations of the likelihood function.

This estimation procedure is still computationally demanding, although standard procedures would likely be infeasible.[3] Evaluating the likelihood for a single simulated individual only takes a fraction of a second, but with $N * M$ simulated individuals, a single evaluation of the likelihood function can take hours.[4]

[3] Using Simulated Maximum Likelihood to estimate the version of the model without learning took around 12 hours for a modest number of draws.

[4] There are two factors that influence the number of calculations needed to evaluate the likelihood function for a single simulated individual. The most significant determinant of the number of necessary calculations is the number of draws used to approximate the future value terms. Increasing the number of draws by a given factor increases estimation time by nearly the same factor. The second determinant of estimation time is the number of terms used in the interpolation of the CCPs. Increasing the number of terms by using a higher order polynomial approximation increases the calculations needed to evaluate the likelihood function. However, the most significant effect of increasing the number of interpolation terms comes in the increase in the time it takes to update the CCPs.

APPENDIX B

APPENDIX FOR DO CONSUMERS' BELIEFS CONVERGE TO EMPIRICAL DISTRIBUTIONS WITH REPEATED PURCHASES?

B.1 Data Appendix

B.1.1 Online Questionnaire Data

Individuals who participated in the online questionnaire were respondents from a list of emails generated through participation in a separate online questionnaire conducted at UNC during the 2011-2012 and 2012-2013 academic calendar years. These individuals agreed to participate in follow-up surveys at the completion of the separate questionnaire. This separate questionnaire was distributed by instructors to their students, who had the option to participate. Additional details about this separate questionnaire are available in Spence (2013).

Online Questionnaire

[The following is a subset of the questions provided to textbook consumers using Qualtrics online survey software. Notes are provided in brackets.]

Textbook Purchasing Questionnaire

The following survey seeks to gain understanding into how consumers choose which retailers to consider when faced with purchasing decisions. Over the course of this survey you will be presented with a number of hypothetical textbook purchasing decisions. You will be asked about your price expectations from online retailers and your beliefs about the time costs involved with searching within an online market. You will not actually have to price any textbooks from online retailers or visit any website outside of this survey.

Directions: Please answer all questions to the best of your ability. Use the right arrow button at the bottom of the screen to advance to the next page. You may also use the left arrow at the bottom of the screen to move back at any time and change a previous answer. If you are uncomfortable answering a specific question you can either skip that question or exit the survey. Thank you for participating!

How many semesters in total, including this one, have you attended UNC and any other college? (Count a summer session as a semester)

Semesters: ________________

What is (are) your major(s)? Please write Undecided if you do not currently have a stated major.

Major(s): ___________________

Which of the following have you ever purchased a textbook from? (Please check all that apply)

☐ UNC Student Stores (campus bookstore)

☐ Ram Book and Supply

☐ Another college bookstore

☐ Amazon.com

☐ Half.com or Ebay.com

☐ Ecampus.com

☐ Chegg.com

☐ Another online retailer

☐ Another student (directly)

Which of the following have you ever rented a textbook from? (Please check all that apply)

☐ I have never rented a textbook

☐ UNC Student Stores (campus bookstore)

☐ Ram Book and Supply

☐ Another college bookstore

☐ Amazon.com

☐ Half.com or Ebay.com

☐ Ecampus.com

☐ Chegg.com

☐ Another online retailer

☐ Another student (directly)

Please write in the other online retailers you have ever rented or purchased a textbook from:

____

When do you normally purchase (or order) your textbooks?

☐ More than 2 weeks before the semester starts

☐ 1 - 2 weeks before the semester starts

☐ A few days before the semester starts

☐ The day the semester starts

☐ A few days after the semester starts

☐ 1 - 2 weeks after the semester starts

☐ More than 2 weeks after the semester starts

Have you purchased or rented any textbooks for an upcoming summer session?

☐ Yes

☐ No

How many textbooks have you ever purchased or rented online?

☐ 1 - 5

☐ 6 - 10

☐ 11 or more

When purchasing or renting a textbook online, have you ever used a website that shows the lowest prices available from multiple online retailers?

☐ Yes

☐ No

On average, when you purchase a textbook online, how many different online retailers do you visit? Number of Retailers Visited: ________________

Do you have an Amazon Prime membership?

☐ Yes

☐ No

Do you have a Paypal account?

☐ Yes

☐ No

How many online purchases do you typically make in a three month period?

Number of Purchases: ________________

If you were given an isbn number or textbook title and wanted to purchase or rent this textbook online, what is the first website you would visit?

Website Name: ___________

Hypothetically, if you only visited one online retailer, how many minutes do you think it would take to look up **one** textbook and purchase it? (Include the time to search, find the option you want, enter your information, and complete the transaction)

Minutes: ___________

Hypothetically, if you only visited one online retailer, how many minutes do you think it would take to look up **three** textbooks and purchase them? (Include the time to search, find the option you want, enter your information, and complete the transaction)

Minutes: ___________

You will now be given a number of hypothetical textbook purchasing decisions. In each case, you will be given information about the textbook and asked to give your best guess about what the price of this textbook is from an online retailer. This survey is concerned about what your expectations are about prices from online retailers, so please do not actually search for the price of the textbook online. Before presenting you with the hypothetical purchasing decisions, you will be provided with an example of what the questions will be like.

Example: If you searched only one clothing store in the mall (ex. Old Navy), what do you think is the lowest price you could find for a pair of jeans in your size? Please enter your answer as a number. Note that this question does not have a right or wrong answer, it simply asks for your best guess.

$___________ [Denoted “Example Expectation” in following questions]

Example Continued: Given that you dont know the lowest price of a pair of jeans with certainty, there is some chance that the lowest price is lower than $[Example Expectation] and some chance that the lowest price is greater than $[Example Expectation]. In the following questions, you will be asked about your beliefs about the chance that the lowest price you could find would be below $[Example Expectation] and also the chance the lowest price you could find would be above $[Example Expectation].

What do you think is the chance that the lowest price of the pair of jeans is less than $[90% of Example Expectation]? Please enter the chance as a percentage (i.e. a number between 0 and 100). For example: I think there is a 30% chance that the lowest price of the pair of jeans is less than $[90% of Example Expectation];

Percent Chance: ___________

What do you think is the chance that the lowest price of the pair of jeans is more than $[110% of Example Expectation]? Please enter the chance as a percentage (i.e. a number between 0 and 100). For example: "I think there is a 35% chance that the lowest price of the pair of jeans is greater than $[Example Expectation]."

Percent Chance: ____________

You will now be given three textbook purchasing scenarios, each similar to the previous example.

[The following is an example using one of the possible twelve textbooks. Respondents were given three scenarios randomly drawn from three groups of four textbooks (one from each group).]

Scenario: You are assigned "Economics: Principles and Policy" by William Baumol and Alan Blinder for an introductory economics course (ECON-101). [The following information on textbook characteristics was randomly assigned to respondents with 50% likelihood. The likelihood of receiving this information for the subsequent scenarios did not depend on whether the information on characteristics was shown for previous textbook purchasing scenarios.] This is the twelfth and latest edition of the textbook, it was published by South-Western College Publishing, and was last revised in 2012. The dimensions of the book are 8.4 x 1.5 x 11.1 inches, it is a hardcover, it contains 880 pages, and it weighs 4.4 pounds. A picture is provided below:

[Picture presented such as the one shown in the screenshot in Section 3]

Have you ever taken this course?

☐ Yes

☐ No

Have you ever been assigned this textbook?

☐ Yes

☐ No

You know that a new copy of this book costs $212 (including taxes) at the UNC Student Stores. If you searched one online retailer, what do you think the price of a new copy at this online retailer would be (include shipping costs)? Reminder: Please do not actually search for this price. Provide your best guess instead.

$ ____________ [Denoted "New Expectation" in future questions]

What do you think is the probability that the lowest price for a new copy of this book is less than $[85%, 90%, or 95% of New Expectation] (including shipping costs) at the online retailer?

Percent Chance: ____________

What do you think is the probability that the lowest price for a new copy of this book costs more than $[105%, 110%, or 115% of New Expectation] (including shipping costs) at the online retailer? Note that your answer to this question added to your answer from the previous question should not exceed 100.

Percent Chance: ____________

You know that a used copy of this book costs $159 (including taxes) at the UNC Student Stores. If you searched one online retailer, what do you think the price of a used copy at this online retailer would be (include shipping costs)?

$ ____________ [Denoted "Used Expectation" in future questions]

What do you think is the probability that the lowest price for a used copy of this book costs less than $[85%, 90%, or 95% of Used Expectation] (including shipping costs) at the online retailer?

Percent Chance: ____________

What do you think is the probability that the lowest price for a used copy of this book costs more than $[105%, 110%, or 115% of Used Expectation] (including shipping costs) at the online retailer? Note that your answer to this question added to your answer from the previous question should not exceed 100.

Percent Chance: ____________

B.1.2 Textbook Purchasing Scenarios

Table B.1 provides information on the textbooks used in the hypothetical textbook purchasing scenarios. Respondents that completed the survey faced three scenarios; in each scenario, one textbook from each group was randomly assigned to the respondent. The first group is composed of social science textbooks; the second group is composed of hard science textbooks; the third group is composed of humanities textbooks. In the Fall 2012 semester, individuals were presented these scenarios in the previous ordering (social sciences, hard sciences, then humanities). The Spring 2013 questionnaire assigned individuals to groups at random (i.e. roughly one third of respondents completed a scenario with a hard sciences textbook, then social sciences, then humanities).

Textbooks were chosen to provide variation in the following characteristics: the number of total editions of the textbook, whether the textbook is the latest edition, the year of publication, whether the course was designed for an introductory or upper-level course, the type of cover (hardback vs.

Table B.1: Textbook Scenarios

Book #	Group	Title	Author	Edition	Latest Edition
1	1	Economics: Principles and Policies	Baumol and Blinder	12	Yes
2	1	Introductory Econometrics: A Modern Approach	Wooldridge	4	No
3	1	Experience Sociology	Croteau and Hoynes	1	Yes
4	1	Medical Sociology	Cockerham	12	Yes
5	2	Chemistry: The Essential Science	Brown et al.	12	Yes
6	2	Animal Psysiology: Adaptation and Environment	Schmidt-Nielsen	5	Yes
7	2	Earth: Portrait of a Planet	Marshak	3	No
8	2	Data Structures and Algorithm Analysis in Java	Weiss	3	Yes
9	3	Norton Anthology of Short Fiction	Bausch and Cassill	7	Yes
10	3	Medicine and Morality in Haiti	Brodwin	1	Yes
11	3	Western Civilization, Volume 1	Perry	10	Yes
12	3	Voces de Hispanoamerica	Chang-Rodrigues and Filer	4	Yes

Book #	Year Published	Subject	Course	Cover	Pages	Weight (lbs)	ISBN	New Bookstore Price
1	2012	Economics	101	Hard	880	4.4	9780538453577	$286
2	2008	Economics	570	Hard	896	3.4	9780324581621	$262
3	2012	Sociology	101	Paper	576	2.8	978007319353	$141
4	2011	Sociology	469	Paper	432	1.2	9780205054183	$87
5	2011	Chemistry	101	Hard	1200	5.4	9780321696724	$253
6	1997	Biology	451	Hard	617	3.4	9780521570985	$104
7	2007	Geology	101	Paper	880	4.4	9780393935189	$146
8	2011	Computer Science	410	Hard	640	2.1	9780132576277	$151
9	2006	English	123	Paper	1776	2.8	9780393926118	$81
10	1996	Anthropology	470	Paper	260	1.0	9780521575430	$52
11	2012	History	151	Paper	496	1.6	9781111831707	$172
12	2012	Spanish	400	Hard	706	2.2	9781111837921	$213

paperback), the number of pages, and the weight. In the following tables, course number refers to the numbering at UNC for the Fall 2012 and Spring 2013 semesters. New bookstore price refers to the price from UNC's campus bookstore during these semesters.

B.1.3 Survey Attrition and Estimation Sample

Table B.2 summarizes the number of respondents at various points within the survey. For the Fall 2012 semester, 979 respondents began the questionnaire and 734 (75%) completed the questionnaire. For the Spring 2013 semester, 1002 respondents began the questionnaire and 703 (70%) completed the questionnaire. We only exclude individuals who did not complete the background questions. This leaves 820 respondents from the Fall and 798 respondents from the Spring. Of these individuals, we exclude 104 respondents (52 from both semesters) because they had been enrolled in college for more than 10 semesters and/or summer sessions and an additional 49 respondents (29 from Fall 2012 and 20 from Spring 2013) for reporting nonsensical answers (e.g., reporting an expected price of $100,000).

Table B.2: Survey Attrition

	Fall 2012 Respondents	Percent Remaining	Spring 2013 Respondents	Percent Remaining
Began the questionnaire	979	100	1002	100
Completed the background questions	820	83.8	798	79.6
Completed at least one scenario	759	77.5	761	75.9
Completed at least two scenarios	741	75.7	716	71.5
Completed the questionnaire	734	75.0	703	70.2

B.2 Robustness Checks

This section investigates the robustness of the results presented in the paper by providing results from a number of other specifications. Explicitly, we explore the robustness of our results by varying the following:

B.2.1 The criteria for begin omitted from the sample and the number of scenarios used for each respondent.

B.2.2 Testing the distributional assumption.

Table B.3: Results Using Alternative Samples

	Used Books								
	Main Sample			Extended Sample			Single Scenario		
Experience	N	Mean	SD	N	Mean	SD	N	Mean	SD
No Online Purchases	438	0.739	0.201	450	0.732	0.211	151	0.736	0.209
1-5 Online Purchases	917	0.712	0.181	939	0.709	0.189	311	0.700	0.190
6-10 Online Purchases	957	0.683	0.173	976	0.677	0.181	324	0.674	0.181
11+ Online Purchases	1698	0.653	0.180	1755	0.647	0.191	577	0.638	0.186
	New Books								
	Main Sample			Extended Sample			Single Scenario		
Experience	N	Mean	SD	N	Mean	SD	N	Mean	SD
No Online Purchases	440	0.833	0.172	452	0.832	0.186	152	0.828	0.177
1-5 Online Purchases	923	0.819	0.153	945	0.814	0.161	313	0.804	0.166
6-10 Online Purchases	966	0.783	0.161	985	0.779	0.165	326	0.761	0.176
11+ Online Purchases	1706	0.770	0.157	1763	0.768	0.166	581	0.752	0.168

B.2.1 Omission Criteria

Respondents are omitted from our sample in the main body of the paper for two reasons:

1. Being enrolled in more than 10 semesters of college.

2. Reporting expectations less than 10% of the bookstore price or greater than 150% of the bookstore price.

The first criteria is used to focus on traditional college students. The second criteria is used to eliminate respondents who we believe did not take the questionnaire seriously (for example, individuals who reported expectations of $0 or $100,000). To make sure that our results are not biased because of these omissions, we relax the second omission criteria. We also report evidence that our omission criteria is not correlated with our measures of experience.

We proceed to report the main findings from the paper for a less stringent omission criteria. Specifically, we only omit respondents who report expectations less than 1% of the bookstore price or greater than 200% of the bookstore price. This results in four respondents being omitted from the Fall sample and one respondent being omitted from the Spring sample for reporting expectations below 1% of the bookstore price, and four respondents being omitted from the Fall sample and four respondents being omitted from the Spring sample for reporting expectations greater than 200% of

Table B.4: Distribution Results

	Used			New		
Experience	N	mean $\mathbf{E}_i(p)$	mean $\sqrt{\mathbf{Var}_i(p)}$	N	mean $\mathbf{E}_i(p)$	mean $\sqrt{\mathbf{Var}_i(p)}$
	Normal					
No online purchases	291	0.720*** (0.159)	0.126 (0.096)	300	0.837*** (0.134)	0.139 (0.099)
1-5 online purchases	619	0.714*** (0.155)	0.130* (0.094)	662	0.822*** (0.129**)	0.141* (0.096)
6-10 online purchases	667	0.701* (0.150***)	0.125 (0.095)	696	0.796 (0.130)	0.130 (0.088*)
11+ online purchases	1202	0.669*** (0.160***)	0.119** (0.089)	1259	0.784*** (0.134*)	0.133 (0.096)
	Gamma					
No online purchases	411	0.747*** (0.192)	0.219 (0.281)	412	0.836*** (0.164)	0.237 (0.288)
1-5 online purchases	852	0.717*** (0.175)	0.244** (0.294**)	870	0.819*** (0.152*)	0.243* (0.312)
6-10 online purchases	900	0.689 (0.170***)	0.234 (0.292)	908	0.785* (0.158)	0.224 (0.276)
11+ online purchases	1601	0.656*** (0.174)	0.207*** (0.252***)	1604	0.774*** (0.153)	0.217* (0.269)
	Log-Normal - Restricted Sample					
No online purchases	283	0.726*** (0.158)	0.113 (0.068)	299	0.836*** (0.135)	0.129 (0.074***)
1-5 online purchases	610	0.715*** (0.156)	0.116*** (0.067**)	649	0.822*** (0.126***)	0.131*** (0.071)
6-10 online purchases	647	0.701* (0.151***)	0.110 (0.063)	678	0.797 (0.128)	0.121 (0.065**)
11+ online purchases	1168	0.665*** (0.161***)	0.104*** (0.060**)	1218	0.783*** (0.134**)	0.121* (0.068)
	Gamma - Restricted Sample					
No online purchases	317	0.754*** (0.197)	0.108 (0.060)	325	0.847*** (0.164)	0.126** (0.072***)
1-5 online purchases	637	0.715*** (0.174)	0.112*** (0.064**)	690	0.821*** (0.151*)	0.125*** (0.070**)
6-10 online purchases	676	0.694 (0.166***)	0.107 (0.063)	708	0.788** (0.153)	0.114* (0.063**)
11+ online purchases	1210	0.660*** (0.174)	0.101*** (0.059**)	1263	0.780*** (0.150)	0.114*** (0.063**)

Notes: Standard deviations in parenthesis.
The significance levels reported for the mean values are from a two-sample equality of means test. The significance levels for the standard deviations are from Brown and Forsythe's alternative formulation of Levene's robust two-sample equality of variances test.

* refers to p–value $< .1$; ** $< .05$; *** $< .01$

the bookstore price.

Table B.3 reports the mean ratio of expectations to bookstore prices for the main sample as well as the extended sample. Including outliers does not significantly change the estimates of mean price expectations. Also included in table B.3 are the results that only use the first hypothetical textbook purchasing scenario that an individual responded to (out of a potential of six for individuals who completed the questionnaire in the fall and spring semester). The price expectations for the first scenario are lower than for the full sample, but the relationship between experience and price expectations is the same for both groups.

Results from regressions of the ratio of expectations to bookstore prices on measures of experience and other covariates also remain quantitatively similar to the results reported in the main body of the paper (not reported).

Table B.4 reports the results for the normal and gamma distributions. Also included are the results for the log-normal and gamma distributions using a more restrictive sample. Since the normal distribution requires dropping individuals who report a greater than 50% probability of being below the threshold, the restricted samples are constructed using a similar rejection criteria. The samples are constructed by dropping individuals who report a 50% or greater probability of being below the threshold. Then the parameter values are calculated for each individual, and the final sample includes individuals whose parameter values are not in the top or bottom 2.5% of values for either parameter. The results for the log-normal and gamma using the restricted samples does not change the mean of the price expectations by a large amount, but the variability of the expected price distribution falls substantially. With a similar sample construction, the normal distribution is closest to the variability of the empirical distribution. The higher mean variability of the expected price distribution using the less restrictive sample for the log-normal and gamma distributions is driven by the individuals who report a high probability of the price being below the threshold. Ultimately, the results are similar regardless of the distribution used.

B.2.2 Price Distribution

In this section we provide some evidence supporting the use of the log normal price distribution as well as discussing some limitations of the distribution in fitting certain features of the empirical price distribution. Tests for normality reject the assumption of normality for both the distribution of prices and the log of prices for most specifications of the empirical distribution. For used books with a list price greater than $100, the assumption of normality cannot be rejected. For new books, the normal distribution is able to fit the data better than the log-normal distribution. For used books, the normal distribution only fits better for the relatively expensive books. In order for the log normal distribution to fit the long right tail of the price distribution, the result is that it places too little weight on the left tail relative to the empirical distribution. The analysis in this paper is not dependent on a particular distributional assumption. In structural search models, however, an incorrect distributional assumption on the individual's beliefs about the price distribution or about the empirical price distribution can significantly bias estimates. Figure B.1 displays histograms of the empirical prices for the Fall survey period. Figures B.2 and B.3 display kernel density estimates for used and new prices. Finally figures B.5 and B.4 display the time series of the mean daily price with the 95% confidence interval. During the survey periods, prices are relatively stable. Used prices rise considerably leading up to

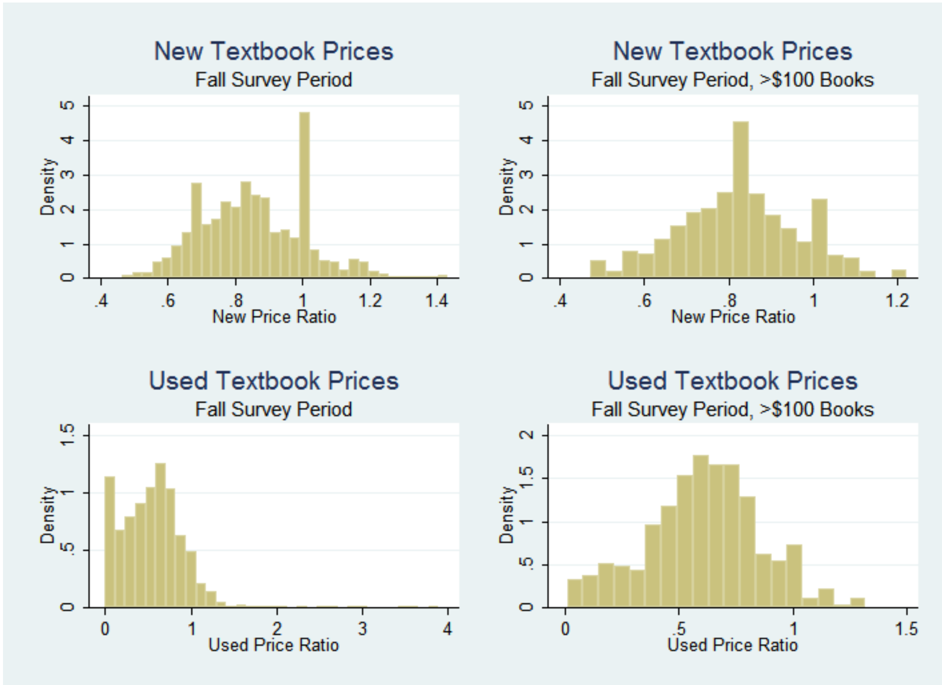

Figure B.1: Histograms of Prices

the start of the semester, however, new prices remain fairly stable throughout the year.

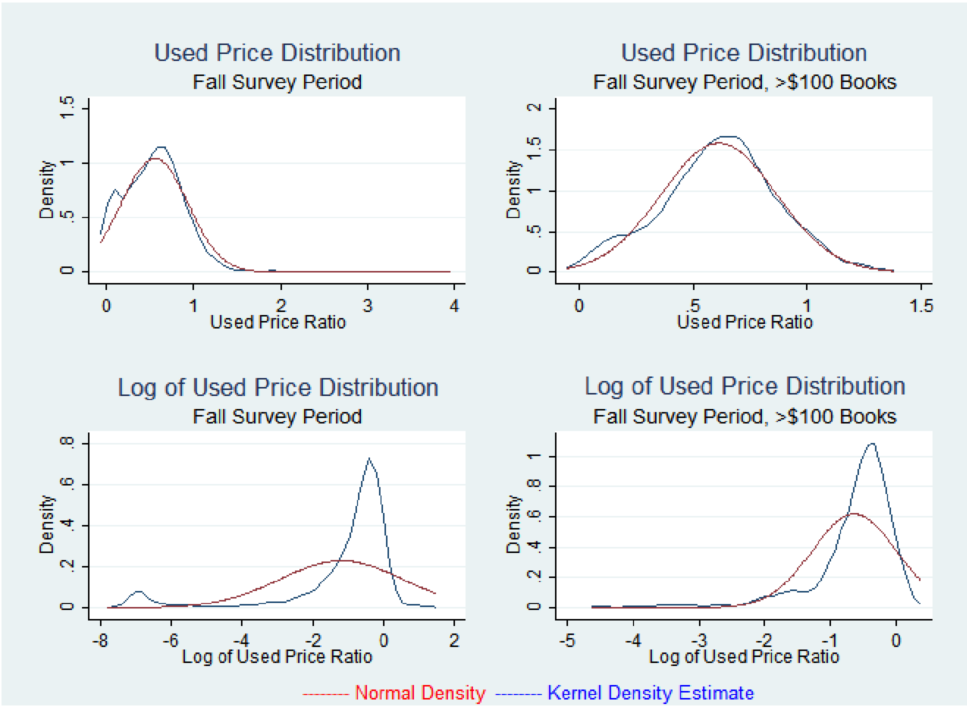

Figure B.2: Kernel Density Estimate

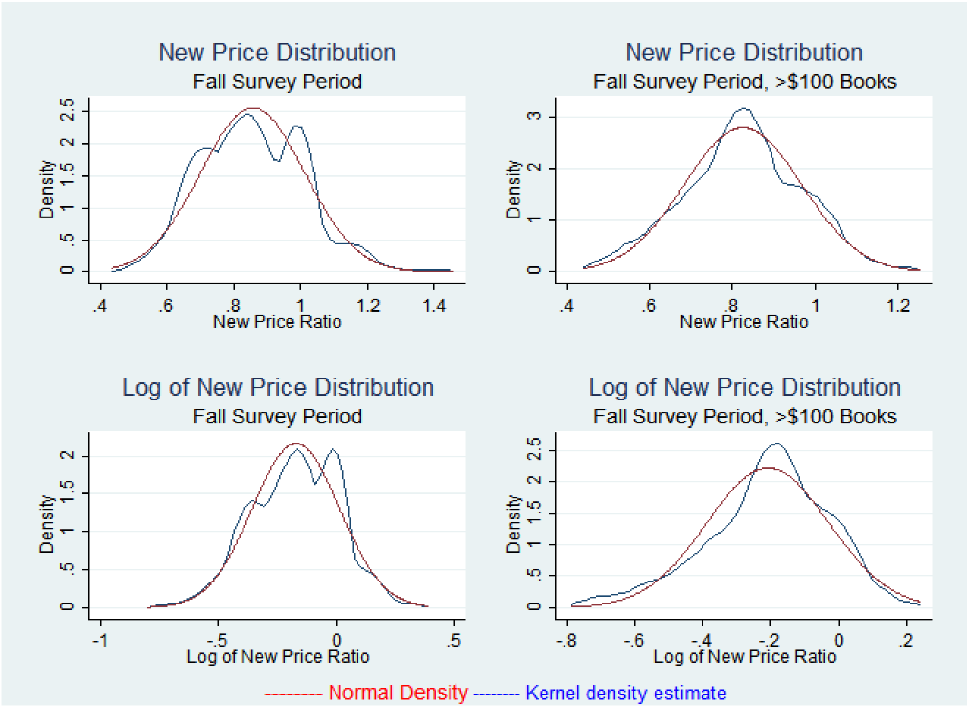

Figure B.3: Kernel Density Estimate

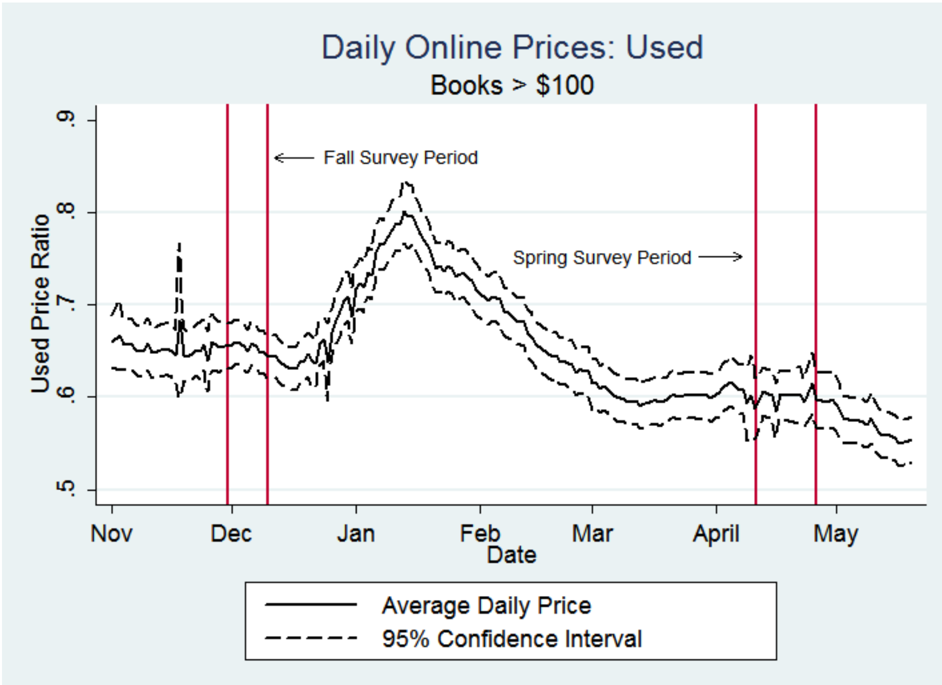

Figure B.4: Daily Used Prices

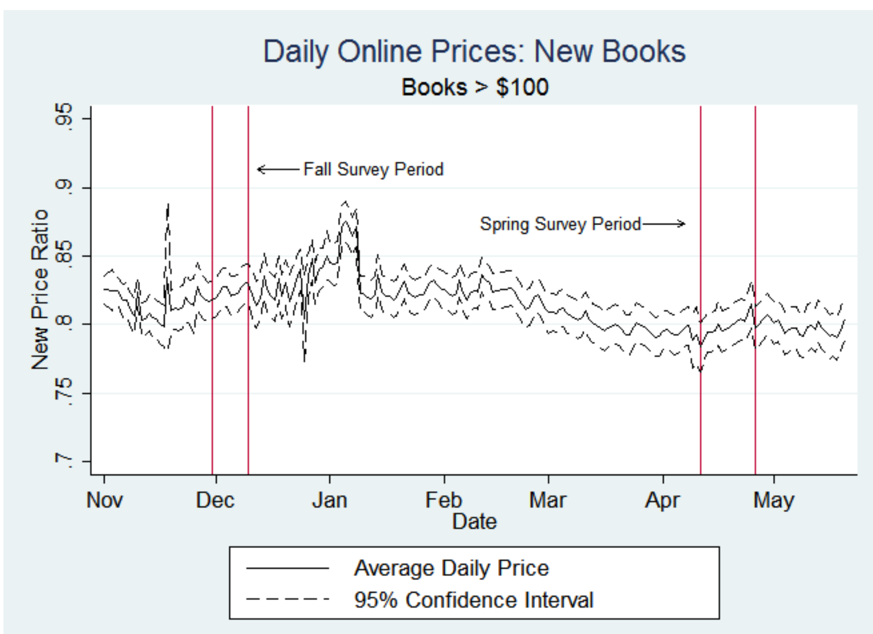

Figure B.5: Daily New Prices

REFERENCES

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

Arcidiacono, P., J. Hotz, and S. Kang (2012). Modeling College Major Choice using Elicited Measures of Expectations and Counterfactuals. *Journal of Econometrics 166*(1), 3–16.

Arcidiacono, P. and J. Jones (2003). Finite Mixture Distributions, Sequential Likelihood and the EM Algorithm. *Econometrica 71*(3), 933–946.

Arcidiacono, P. and R. Miller (2011). Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity. *Econometrica 7*(6), 1823–1868.

Arcidiacono, P., H. Sieg, and F. Sloan (2007). Living Rationally Under the Volcano? An Empirical Analysis of Heavy Drinking and Smoking. *International Economic Review 48*(1), 37–65.

Becker, G., M. Grossman, and K. Murphy (1994). An Empirical Analysis of Cigarette Addiction. *American Economic Review 84*(3), 396–418.

Becker, G. and K. Murphy (1988). A Theory of Rational Addiction. *Journal of Political Economy 96*(4), 675–700.

Benabou, R. and J. Tirole (2004). Willpower and Personal Rules. *Journal of Political Economy 112*(4), 848–886.

Bernheim, B. and A. Rangel (2004). Addiction and Cue-Triggered Decision Processes. *American Economic Review 94*(5), 1558–1590.

Chaloupka, F. (1991). Rational Addictive Behavior and Cigarette Smoking. *Journal of Political Economy 99*(4), 722–742.

Chaloupka, F. and K. Warner (2000). The Economics of Smoking. In J. Newhouse and A. Cuyler (Eds.), *Handbook of Health Economics*, pp. 1539–1627. New York: North Holland.

Chevalier, J. and A. Goolsbee (2009). Are Durable Goods Consumers Forward-Looking? Evidence from College Textbooks. *Quarterly Journal of Economics 124*(4), 1853–1884.

Ching, A., T. Erdem, and M. Keane (2011). Learning Models: An Assessment of Progress, Challenges, and New Developments. Working Paper.

Choo, E. (2000). Rational Addiction and Rational Cessation: A Dynamic Structural Model of Cigarette Consumption. Working Paper.

Darden, M. (2011). *A Dynamic Stochastic Model of Lifetime Smoking Behavior*. Ph. D. thesis, University of North Carolina at Chapel Hill.

De los Santos, B., A. Hortacsu, and M. Wildenbeest (2012a). Search with Learning. *Working Paper*.

De los Santos, B., A. Hortacsu, and M. Wildenbeest (2012b). Testing Models of Consumer Search using Data on Web Browsing and Purchasing Behavior. *American Economic Review 102*(6).

DeCicca, P., D. Kenkel, and A. Mathios (2002). Putting out the Fires: Will Higher Taxes Reduce the Onset of Youth Smoking? *Journal of Political Economy 110*(1), 144–169.

DeCicca, P., D. Kenkel, A. Mathios, Y. Shin, and J. Lim (2008). Youth Smoking, Cigarette Prices, and Anti-smoking Sentiment. *Health Economics 17*, 733–749.

Einav, L. (2005, May). Informational Asymmetries and Observational Learning in Search. *Journal of Risk and Uncertainty 30*(3), 241–259.

Emery, S., M. White, and J. Pierce (2001). Does Cigarette Price Influence Adolescent Experimentation? *Journal of Health Economics 20*, 261–270.

Erdem, T. and M. Keane (1996). Decision-Making under Uncertainty: capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science 15*, 1–20.

Fletcher, J., P. Deb, and J. Sindelar (2009). Tobacco Use, Taxation and Self Control in Adolescence. NBER Working Paper Series, No. 15130.

Gilleskie, D. and K. Strumpf (2005). The Behavioral Dynamics of Youth Smoking. *Journal of Human Resources 40*(4), 822–866.

Glied, S. (2002). Youth Tobacco Control: Reconciling Theory and Empirical Evidence. *Journal of Health Economics 21*, 117–135.

Goel, R. and M. Nelson (2006). The Effectiveness of Anti-smoking Legislation: A Review. *Journal of Economic Surveys 20*(3), 325–355.

Gordon, B. and B. Sun (2009). A Dynamic Structural Model of Addiction, Promotions, and Permanent Price Cuts. Working Paper.

Gruber, J. and B. Koszegi (2001). Is Addiction "Rational"? Theory and Evidence. *Quarterly Journal of Economics 116*(4), 1261–1303.

Hong, G. and M. Shum (2006). Using Price Distributions to Estimate Search Costs. *RAND Journal of Economics 37*(2), 257–275.

Hortaçsu, A. and C. Syverson (2004). Product Differentiation, Search Costs, and Competition in the Mutual Fund Industry: A Case Study of S&P 500 Index Funds. *Quarterly Journal of Economics 119*(2), 403–456.

Hotz, V. and R. Miller (1993). Conditional Choice Probabilities and the Estimation of Dynamic Models. *Review of Economic Studies 60*(3), 497–529.

Hotz, V., R. Miller, S. Sanders, and J. Smith (1994). A Simulation Estimator for Dynamic Models of Discrete Choice. *Review of Economic Studies 61*(2), 265–289.

James, J. (2015). A Tractable Estimator for General Mixed Multinomial Logit Models. Working Paper.

James, J. and B. Matsumoto (2013). CCP Estimation with Continuous Unobserved Heterogeneity. Working Paper.

Keane, M. and K. Wolpin (1994). Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence. *Review of Economics and Statistics 76*(4), 648–672.

Koulayev, S. (2009). Estimating Search with Learning. *Net Institute Working Paper No. 08-29*.

Koulayev, S. (2012). Estimating Demand in Search Markets: The Case of Online Hotel Bookings. *Federal Reserve Bank of Boston Working Paper No. 09-16*.

Koulayev, S. (2013). Search with Dirichlet Priors: Estimation and Implications for Consumer Demand. *Journal of Business and Economic Statistics*.

Laviolette, S. and D. van der Kooy (2004). The Neurobiology of Nicotine Addiction: Bridging the Gap from Molecules to Behavior. *Nature Reviews: Neuroscience 5*, 55–65.

McFadden, D. and K. Train (2000). Mixed Multinomial Logit Models for Discrete Response. *Journal of Applied Econometrics 15*, 447–470.

Moraga-González, J., Z. Sándor, and M. Wildenbeest (2009). Consumer Search and Prices in the Automobile Market. *Working Paper*.

Nestler, E. and G. Aghajanain (1997). Molecular and Cellular Basis of Addiction. *Science 278*(5335), 58–63.

Orphanides, A. and D. Zervos (1995). Rational Addiction with Learning and Regret. *Journal of Political Economy 103*(4), 739–758.

Orphanides, A. and D. Zervos (1998). Myopia and Addictive Behavior. *The Economic Journal 108*(446), 75–91.

Sonnemans, J. (1998, April). Strategies of Search. *Journal of Economic Behavior and Organization 35*(3), 309–332.

Spence, F. (2013). Does Consumer Inexperience Generate Welfare Losses? Evidence from the Textbook Market. *Working Paper*.

Stigler, G. (1961, June). The Economics of Information. *Journal of Political Economy 69*(3), 213–225.

Stinebrickner, T. and R. Stinebrickner (2011). Math or Science? Using Longitudinal Expectations Data to Examine the Process of Choosing a College Major. *NBER Working Papers 16869*.

Suranovic, S., R. Goldfarb, and T. Leonard (1999). An Economic Theory of Cigarette Addiction. *Journal of Health Economics 18*, 1–29.

Thaler, R. and H. Shefrin (1981). An Economic Theory of Self-Control. *Journal of Political Economy 89*(2), 392–406.

Train, K. (2007). A Recursive Estimator for Random Coefficient Models. Working Paper.

Tworek, C., R. Yamaguchi, D. Dloska, S. Emery, D. Barker, G. Giovino, P. O'Malley, and F. Chaloupka (2010). State-level Tobacco Control Policies and Youth Smoking Cessation Measures. *Health Policy 97*, 136–144.

USDHHS (2012). *Preventing Tobacco Use Among Youth and Young Adults: A Report of the Surgeon General*. Atlanta: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health.

USDHHS (2013). *The Health Consequences of Smoking - 50 years of progress*. Washington D.C.: U.S. Department of Health and Human Services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health.

Winston, G. (1980). Addiction and Backsliding: A Theory of Compulsive Consumption. *Journal of Economic Behavior and Organization 1*, 295–324.

Wiswall, M. and B. Zafar (2012). Determinants of College Major Choices: Identification from an Information Experiment. *Working Paper*.