A Theoretical and Empirical Study of Addiction

Thesis by

Angela A. Hung

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2001

(Submitted February 20, 2001)

Dedication

To my family, for making everything possible

Acknowledgments

The completion of my thesis would not have been possible without the help and support from the faculty, staff, and my fellow classmates in the Division of Humanities and Social Sciences.

I thank my committee, Colin Camerer, Dave Grether, Matt Jackson and Simon Wilkie, for their time and endless patience. Your guidance has been the most important influence in shaping my professional career.

Thank you, to Colin for serving as my advisor and role model (well, at least in my professional life) since my first year at Caltech. Your advice, support, and friendship have been instrumental in helping me through my years in graduate school.

I would also like to thank seminar participants at the 2000 World Congress of the Econometric Society, Carnegie Mellon University, University of Miami, University of California at Irvine, and CentER at Tilburg University for helpful comments and insight on the papers in this dissertation.

Thank you to Information Resources, Incorporated for providing the data used in this thesis. Financial support from the Division of Humanities and Social Sciences, the John Randolph Haynes and Dora Haynes Foundation, the Russell Sage Foundation, and the Li Ming Scholarship is greatly appreciated.

And of course, thank you to Jeffrey Dominitz, my partner in life, for your unwavering love and support at every step of the way.

Abstract

Consumption of addictive substances poses a challenge to economic models of rational, forward-looking agents. This dissertation presents a theoretical and empirical examination of consumption of addictive goods.

The theoretical model draws on evidence from psychology and neurobiology to improve on the standard assumptions used in intertemporal consumption studies. I model agents who may misperceive the severity of the future consequences from consuming addictive substances and allow for an agent's environment to shape her preferences in a systematic way suggested by numerous studies that have found craving to be induced by the presence of environmental cues associated with past substance use. The behavior of agents in this behavioral model of addiction can mimic the pattern of quitting and relapsing that is prevalent among addictive substance users.

Chapter 3 presents an empirical analysis of the Becker and Murphy (1988) model of rational addiction using data on grocery store sales of cigarettes. This essay empirically tests the model's predictions concerning consumption responses to future and past price changes as well as the prediction that the response to an anticipated price change differs from the response to an unanticipated price change. In addition, I consider the consumption effects of three institutional changes that occur during the time period 1996 through 1999.

Contents

List of Figures

List of Tables

Chapter 1 Introduction

Schelling (1978) proposes that consumption of addictive goods is an anomaly in consumer theory because "consumers [are] getting negative satisfaction out of something they spend a lot of money to consume" (p. 293). Unless these consumers prefer "negative" satisfaction, this anomaly poses a challenge to modern economic theory. It appears, however, that this phenomenon can be understood once the standard restrictive assumptions of dynamic models of consumption are relaxed.

This thesis presents an economic model of consumption of addictive goods. Unlike previous economic models of addiction, this model can generate behavior that resembles the pattern of quitting and relapse that is extremely prevalent among addictive substance users. The standard economic model of addiction is the "rational addiction" model of Becker and Murphy (1990). This thesis includes an empirical analysis of this model.

Habit formation models, such as Pollak (1970), Ryder and Heal (1973) and Boyer (1978), relax the assumption of intertemporal separation of utility. They allow utility of current consumption to depend on past consumption. The macroeconomic theory research by Ryder and Heal (1973) and Boyer (1978) finds that the assumption of intertemporally dependent preferences can substantially change the optimal growth path of an economy. The more relevant study is Pollak's study, which finds that the relaxation of the assumption of intertemporal separability also leads to different optimal behavior on the individual level. For example, long-run demand functions differ from short-run demand functions. Pollak, however, implicitly assumes that agents are myopic.

The rational choice model of addiction was first introduced by Stigler and Becker (1977) in an attempt to demonstrate that many behavioral phenomena, including addiction, can be modeled without the assumption of a change in tastes over time. The model of rational addiction was further developed by Becker and Murphy (1988). Like Pollak, Becker and Murphy focus on individual behavior. However, they assume that individuals are fully aware of the effect of their current consumption on future consumption. They find that consumption patterns consistent with addiction result from forward-looking utility maximization with stable preferences.

Orphanides and Zervos address the criticism of the Becker and Murphy model that addicts in their model are "happy addicts" in that they choose their addiction. In the Orphanides and Zervos model, agents are uncertain as to whether or not they will experience negative side effects as a result of past consumption. By the time the individual realizes his true type, he may already be addicted. Their model captures the same characteristics of addictive goods as Becker and Murphy, but it also offers an explanation for such things as experimentation with addictive substances; the simultaneous existence of casual users, addicts, and non-users; and the role of drug education programs.

The work presented in the subsequent chapters follows directly from this line of research, but it is also strongly influenced by research in other disciplines. The models most similar to the one that I develop are those that stress the importance of environmental cues in explaining the consumption patterns of those who consume addictive goods (Laibson, 1999 and Loewenstein, 1999). Both papers draw on neurobiological and psychological evidence that environmental cues associated with past consumption of addictive substances can induce craving. In Laibson's model, past behavior in a certain environment only affects current utility if the agent is currently in that environment. His model demonstrates how it is possible for an agent to be addicted in one environment, but not another. In addition to the characteristics of addiction that Becker and Murphy explain, Laibson's model can also explain short-term impatience with regard to consumption of addictive goods. Loewenstein finds that utility derived from consuming the addictive good decreases over time, while, simultaneously, the craving the agent experiences increases in severity if he abstains from consuming the good in the presence of the environmental cue.

There are also models of addiction arising from self-control problems. In some models agents have two personalities with distinct preferences (see Schelling, 1978; Thaler and Shefrin, 1981; or Winston, 1980). In other models, agents are simply over-attentive to present well-being at the expense of future well-being (O'Donoghue and Rabin, 1999 or Gruber and Koszegi, 1999). Both types of self-control models predict that individuals may choose to constrain their choice set in order to control or prevent addiction.

"Adjustment cost" theories of addiction (Jones, 1999 or Suranovic, Goldfarb and Leonard, 1999) model utility from current consumption as depending on a reference level of consumption, which is a function of past consumption. The disutility from decreasing current consumption below the reference level is greater in magnitude than the utility from increasing consumption beyond the reference level.

Chapter 2 presents a behavioral model of addiction. The model adds to the literature by explaining the cycle of quitting and relapse that is extremely common in substance abusers. All the previous economic models of addiction can easily explain consumption of addictive goods. Agents receive immediate positive utility from consumption, but the negative effects are delayed. In some models, this delay is exacerbated by self-control problems or uncertainty or underestimation of these negative consequences. In my model, the negative effects from consumption of the addictive good is not only delayed, but also underestimated. Not all the previous models have a well-motivated explanation for why agents would choose to quit. In Becker and Murphy, for example, agents will quit only as the result of an exogenous shock to the measure of past consumption. In the model presented in the next chapter, agents may quit when they realize the true negative consequences.

The real puzzle is why a person would resume consumption of an addictive substance after deciding to quit. Addiction research in neurobiology suggests that seemingly neutral environments are the main cause of relapse. A person's physiological system learns to predict the onset of addictive substances through environmental cues after repeated drug use. Even long after quitting, experiencing environmental cues that were once associated with drug use will in a sense remind the individual's system of past drug use. This "reminder" manifests itself as craving for the addictive substance. In my model, I allow for an agent's environment to shape her preferences in a systematic way suggested by these findings. Every period, the agent is in one of two possible environments. Past consumption in a given environment only enters into the utility function if the agent is in that environment. Therefore, even though the environments may initially be neutral in that they have no direct effect on utility, preference can come to depend on environment.

The main results are driven by the environmentally dependent preferences, separability between environments, and multiple steady states. The multiple steady states are possible because of the complementarity between current and past consumption. Under this framework, I show how an agent can choose to quit her addiction in one environment (e.g., hospital, jail), but not in the other (e.g., home). Therefore, when she is in the first environment, she consumes very little or none of the addictive good, but when she is in the other environment, environmental cues trigger craving, and she resumes consuming large quantities.

Chapter 3 presents an empirical analysis of the Becker and Murphy model of rational addiction using data on grocery store sales of cigarettes. Thus far, Becker and Murphy's rational model has been the standard model of addiction in economics. There have been a few empirical tests of the rational addiction model that pertain to a variety of addictive substances and activities, such as cocaine (Grossman and Chaloupka, 1998), alcohol (Grossman, Chaloupka, and Sirtalan, 1998), casino gambling (Nichols, 1999) and cigarettes (Becker, Grossman and Murphy, 1994; Chaloupka, 1991; Keeler, Hu, Barnett and Manning, 1993; and Gruber and Kőszegi, 1999).

The previous empirical tests of rational addiction that study cigarette addiction typically use either state cigarette and tobacco tax receipts or survey data. A serious problem with using cigarette and tobacco tax receipts to measure consumption is that, for most states, state-level tobacco taxes are paid by tobacco distributors, rather than tobacco consumers. Therefore, state-paid tobacco taxes more accurately reflect distributors' demand for cigarette and tobacco tax stamps, rather than consumer demand for cigarettes. As for survey data, there may be concern that survey respondents may deny or downplay their consumption of such goods as cigarettes, alcohol, or illegal drugs due to social conformity.

The dataset that I use avoids these data problems. The data, compiled by Information Resources Incorporated from grocery store scanner data, describe weekly sales in 20 markets that span the states of California, Arizona, Colorado, Nevada, and Washington.

The previous tests of the rational addiction model have focused on the model's predictions concerning consumption responses to future and past price changes. This essay also allows an empirical test of the prediction that the response to an anticipated price change differs from the response to an unanticipated price change. According to the rational model, if a price change will cause an agent to change his consumption of the addictive good, then if the price change is anticipated, as in the case of an announced future tax increase, the agent will change his consumption after the announcement, but before the implementation of the price change.

I consider the consumption effects of three institutional changes that occur during the time period 1996 through 1999. The first is the ban on smoking in bars and taverns in California as part of the state's comprehensive "Smoke-Free Workplace" law. Secondly, as a result of the settlement that the five largest tobacco companies signed with 46 states in November 1998, these companies raised wholesale tobacco prices by 45 cents per pack, the largest cigarette price increase in history. Lastly, in the November 1998 election, California voters approved a 50 cent tax increase on cigarettes.

Chapter 2 A Behavioral Theory of Addiction

2.1 Introduction

Consumption of addictive goods has been studied by researchers in such diverse fields as psychology, biochemistry, neurobiology, epidemiology, and sociology. Herrnstein and Prelec (1992) argue that this broad range of disciplines reflects the complexity of the issues involved. Recent work in economics adds a unique perspective to the study of addiction.

The rational choice model of addiction was first introduced by Stigler and Becker (1977) in an attempt to demonstrate that many behavioral phenomena, including addiction, can be modeled without the assumption of a change in tastes over time. The model of rational addiction was further developed by Becker and Murphy (1988) to explain how a perfectly rational forward looking agent may develop a harmful addiction. Becker and Murphy (B-M) present an infinite horizon continuous time problem where utility depends on current consumption of addictive and non-addictive goods as well as a stock of past consumption of the addictive good. Agents are aware of the negative effect of their current consumption of a (harmfully) addictive good on future utility via future craving. The key to this model lies in the relaxation of the usual assumption of intertemporal separability. Consumption patterns consistent with addiction result from forward-looking utility maximization with stable preferences. The framework incorporates characteristics associated with addiction such as tolerance, reinforcement, and withdrawal, and it offers an explanation for behaviors such as bingeing or quitting "cold-turkey."

Orphanides and Zervos (O-Z) extend the B-M framework to an infinite horizon discrete time problem in which there is uncertainty about types. The population consists of addictive types, who may experience negative side effects as a result of past consumption, and non-addictive types, who are not adversely affected by past consumption. In their model, the negative side effects from past consumption are irregular. An addictive type may, therefore, believe that he is a non-addictive type and begin to consume as a non-addictive type would. By the time the individual realizes his true type, he may already be addicted. O-Z refer to these individuals as "regretful" addicts–if they had known with certainty that they were addictive types, they would have consumed less or none of the addictive good. Their model captures the same characteristics of addictive goods as B-M, but it also offers an explanation for such things as experimentation with addictive substances; the simultaneous existence of casual users, addicts, and non-users; and the role of drug education programs.

There are also models of addiction that deviate from the rational paradigm. Laibson (1999) presents a model of "cue-based consumption," in which agents perfectly forecast their preferences, but neutral environments, or cues, can eventually affect both welfare and behavior. Utility in Laibson's model is qualitatively similar to the utility function used in B-M. As in the rational models, past behavior affects current utility and marginal utility. However, in Laibson's model, past behavior in a certain environment only affects current utility if the agent is currently in that environment. His model demonstrates how it is possible for an agent to be addicted in one environment, but not another. The Laibson framework formalizes a biological micro-foundation for why an agent may have significantly different preferences in different environments. In addition to the characteristics of addiction that B-M explain, Laibson's model can also explain short-term impatience with regard to consumption of addictive goods.

Like Laibson, Loewenstein (1999) stresses the importance of environmental cues in explaining the consumption patterns of those who consume addictive goods. Loewenstein argues that drug craving falls into the category of "visceral factors," which includes such other motivational states as hunger, thirst, or sexual arousal. Visceral factors in general, and craving in particular, are defined by a direct, negative impact on utility together with the ability to focus attention on alleviating this aversive effect so that the relative desirability of other goods or actions is severely diminished. Loewenstein's visceral factor account of addiction places great weight on environmental cues because these cues can induce craving. He finds that the utility derived from consuming the addictive good decreases over time, while, simultaneously, the craving the agent experiences increases in severity if he abstains from consuming the good in the presence of the environmental cue.

This line of economic research may be traced back to Schelling (1978) who, in his essay on "Egonomics, or the art of self-management," argues that agents often behave as if they are two people–one who is "straight" and one who is "wayward." Thaler and Shefrin (1981) explicitly model an agent as having two sets of preferences at any given point in time. One set of preferences represents short-run preferences, while the other represents long-run preferences. In order to maximize long-run preferences, an agent may choose to restrict his short-run choice set. For example, in the case of alcohol abuse, an alcoholic may take Antabuse, which will make him severely ill if he then consumes alcohol. Their model of self control implies that "people will rationally choose to impose constraints on their own behavior."

O'Donoghue and Rabin (1999) explicitly model agents as having self-control problems. They show how self-control problems can affect the consumption of addictive goods by comparing agents who have no self-control problems, agents who have self-control problems but are not aware of this, and agents who have self-control problems and are aware of their problems. Unlike the previously mentioned models, O'Donoghue and Rabin's (O-R) model considers both stationary and dynamic preferences as well as finite and infinite horizons. Gruber and Koszegi (1999) generalize the O-R framework from the case of a binary consumption decision to continuous consumption and include prices so that they can analyze optimal government policy.

Despite this range of research, there are two prevalent features of addiction and addicts themselves that these rational addiction models do not capture at all, and that the behavioral models of Laibson, O'Donoghue and Rabin, and others fail to formalize completely. The initial choice to consume an addictive good may be voluntary, but, after sustained drug use, the addict's physiological system is altered in such a way that the individual's preferences change (for example, see Leshner, 1997 or O'Brien and McLellan, 1996). The rational models capture these changes simply by the inclusion of past consumption in the current utility function. In the Becker-Murphy, Laibson and O'Donoghue-Rabin models, the forward-looking agent perfectly foresees the future effects of his current consumption (although agents in the O-R model may not perfectly predict their own self-control problems). As in the B-M model, individuals in the Orphanides-Zervos model know the extent of the future effects, should they occur, but they do not know when, and even if, they will occur. However, in contrast to all of these models, addictions are frequently believed to result from underestimation of future cravings (see for example, Loewenstein et al., 1999 or Loewenstein, 1999).

The second omitted feature is the pattern of quitting and relapse that is frequently seen in addicts. It is estimated that 50-70% of addicts who complete a treatment program fail to abstain (i.e., relapse at least once) within the following year (O'Brien and McLellan, 1996). In the Becker-Murphy model, the decision to quit or to relapse can be explained by the addition to the model of an exogenous shock that directly affects the measure of past consumption. However, this ad hoc extension of the model is not empirically testable, offers no room for policy analysis, and leaves unclear the interpretation of a shock to the consumption stock variable. Furthermore, if the shocks are sufficiently regular, then the model should include the agent's beliefs about the process generating these shocks. This type of change to the model could significantly change the dynamics and results.

Alternatively, the B-M model can explain bingeing cycles, which may also be interpreted as quitting-relapsing cycles, by introducing two separate consumption stock variables which depreciate at different rates. As with the exogenous shock to the consumption stock variable, the justification for having two different consumption stocks is not made clear.

The O-Z model predicts that quitting can occur at most once, at the time that the agent realizes he is an addictive type. The O-R model generates, in the infinite horizon case with stationary preferences, agents who may begin an addiction and agents who may end an addiction, but these agents do not quit and then relapse.

Research outside of economics offers insight into the phenomenon. Numerous studies have found that the presence of environmental cues that have been associated with past consumption of an addictive good can induce craving, even after the addict has quit the substance: "Even after detoxification and long periods of abstinence, relapse frequently occurs despite sincere efforts to refrain. People or situations previously associated with drug use may provoke a relapse" (O'Brien 1997, p. 66). Although Laibson's model can not explain quitting and therefore relapse, it can explain this important link between environment and behavior that is often observed in cases of relapse.

The model presented in this paper captures both of the aforementioned features of addiction: agents may not perfectly forecast the effect of their current consumption on future utility, and, among those agents who do begin to consume the addictive good, some may exhibit consumption behavior consistent with a pattern of quitting and relapse. First, the agent knows that current consumption of an addictive good will decrease future utility and increase future marginal utility from consumption of the addictive good, but she does not know the strength of these effects. In particular, after sufficient experience consuming the addictive good, these effects become more severe, but the agent does not fully anticipate this change.

Behavior that can be interpreted as a pattern of quitting and relpase is generated by building on Laibson's (1999) framework in which environment may play a role in shaping preferences, together with allowing for imperfect foresight that is similar to the uncertainty in Orphanides and Zervos (1995). Using this framework, I show how an addict may choose to quit her addiction in one environment, yet being placed in an environment in which she had frequently used the addictive substance may trigger such strong craving that the addict will resume consumption of the addictive good. The misperception of the tolerance function can generate quitting, while the link between environments and preferences can generate relapse.

This paper focuses on substance addiction. The wealth of information from other disciplines has given economists a number of insights as to how preferences for addictive substances may be modelled. Of course, consumption of addictive goods is a somewhat anomalous example of consumption behavior. However, it is easy to see how such an analysis could apply to a wide range of consumption goods for which preferences display some degree of habit formation, albeit not nearly as strong as that of addictive goods.

2.2 Evidence from Psychology and Neurobiology

The complexity of issues involved in studying addiction is revealed not only in the number of disciplines that are involved in its research, but also in the variety of definitions of addiction (for examples of the wide range of definition, see Pomerleau and Pomerleau, 1988). There does appear to be a strong consensus, however, about the underlying behavioral mechanisms involved in the addiction process.

2.2.1 Conditioned Responses

An organism's physiological system relies on internal equilibrium (Koob and LeMoal, 1997). Disturbances to stability are mediated by *homeostatic mechanisms*, mechanisms that work to return the organism to its equilibrium. For example, even though the external temperature may fluctuate, one's body maintains a constant internal body temperature through adjustments to heart rate and blood pressure. However, some disturbances, such as those caused by the administration of an unfamiliar chemical, require more complex strategies.

Classical conditioning is the experimental study of anticipatory responses. In a typical conditioning paradigm, two stimuli are repeatedly paired so that eventually one stimulus predicts the second stimulus. Conditioning then allows an organism's physiological system to prepare for the second stimulus. Classical conditioning studies have revealed that there are two effects of repeated drug administration: the responses elicited by the administration of the drug, which are labelled as feedback responses, and responses elicited by the anticipation of administration of the drug, labelled feedforward responses (see Siegel et al., 1988, or Eikelboom and Stuart, 1982).

In terms of addiction research, the first stimulus is typically an environmental "cue." Cues for laboratory animals can include sounds, temperatures, or such visual cues as colors. Cues in the addict's world "can include mood states (positive as well as negative), specific persons, locations, events or times of year, mild alcohol intoxication, interpersonal strife previously soothed by cocaine euphoria, or abuse objects (for example, money, white powder, glass pipes, mirrors, syringes, and single-edged razor blades)" (Gawin, 1991, p. 1582).

Homeostatic responses tend to be compensatory. That is, they work to counteract the direct effect of the substance in order to restore stability. For example, nicotine raises blood sugar. The compensatory response that is generated works to lower blood sugar. When the administration of a drug is anticipated, feedforward mechanisms are activated, and therefore the compensatory response is operational before the actual administration of the drug. If the drug is subsequently administered, then the effects of the drug appear to be diminished. This "progressively diminished response to a drug over the course of successive administrations defines tolerance" (Siegel et al., 1988, p. 88).

The consensus among researchers is that tolerance and withdrawal are both manifestations of the same mechanism. Tolerance manifests itself when the drug is administered, and withdrawal occurs when the drug is withheld. Continuing with the nicotine example, suppose a nicotine addict frequently follows drinking an alcoholic beverage by smoking a cigarette. Eventually, the consumption of an alcoholic beverage signals the physiological system that nicotine will soon follow. If the nicotine addict subsequently smokes a cigarette, then by the time nicotine has entered the system, the blood sugar level has already decreased. Therefore, the net increase in blood sugar from the administration of nicotine is not as large as if the system had not anticipated the nicotine. On the other hand, if the addict does not subsequently smoke the cigarette, the decrease in blood sugar will cause the addict to feel hunger or irritability, traits often associated with nicotine withdrawal.

Repeated administration of a drug does not necessarily imply tolerance. Instead, tolerance is environmentally specific. It results from repeated administration of a drug in the presence of environmental cues. Because "overdose" is often simply a failure of tolerance, studies of overdose can shed light on tolerance. In a study by Siegel et al. (1982), rats were given regular and increasing doses of heroin in a specific environment. The rats were subsequently given a high dose of heroin in either the familiar environment or an unfamiliar environment. Survival rates were significantly higher for rats that received the heroin in the familiar environment than for rats who were given heroin in an unfamiliar environment. Similarly, in a small study of survivors of heroin overdose, the majority stated that the overdose occurred when the drug was administered without the usual environmental cues (Siegel, 1984).

Lastly, consider the experience of U.S. enlisted Army men who served in Vietnam. While serving overseas, a large proportion of soldiers, the majority of whom had little prior experience with narcotics, became addicted to heroin and/or opium (Robins 1993). In a study of over 600 men who left Vietnam in September of 1971, Robins

(1974) found that 45% of enlisted men had tried narcotics while in Vietnam. About 20% of the sample of soldiers reported that they had felt addicted to heroin or opium in Vietnam. Despite prior warning of a mandatory urine test at the time of departure, 11% of enlisted men tested positive. However, one year after discharge, the relapse rate among the addicted servicemen was only 5%. In contrast, young men who had not served in Vietnam who were treated in a Federal Narcotics Hospital during the same time period as the Vietnam study had a six month relapse rate of 67%. In terms of the conditioning framework, because the addicted servicemen were removed from the environment that they had associated with opiate use, and returned to an environment with very few past drug cues, it is not surprising that their relapse rate is so low.

The meaning of the term "craving" is less clear. Unlike tolerance and withdrawal symptoms which can be measured by changes in observed outcomes such as heart rates, chemical levels in the brain, or blood sugar levels in both humans and animals, craving is subjective and usually measured by human self-reports. For purposes of this paper, craving is taken to be "a strong desire for the alleviation of unpleasant withdrawal symptoms" (Marlatt, 1987, p. 42).

2.2.2 Perceptions of the Effects of Addictive Substances

The degree of tolerance and the intensity of withdrawal symptoms vary widely among users and substances (Goldstein and Kalant, 1990). Because effects of substances vary widely across substances, it is not surprising that people have misperceptions or judgment biases about those effects.

For example, when asked their perceptions of the risks associated with heavy drinking and drunk driving, those who abstain from drinking alcohol perceive more risks than those who drink. (Agostinelli and Miller, 1994). Champion and Bell (1980) also report an inverse relationship between substance use and perceptions of danger of addictive substances in a study in Australia that includes high school and college students, nurses, prisoners, probationers, and juvenile delinquents. In a survey of college students, Rohsenow (1983) finds that social drinkers expect that other people will be more strongly affected, for both positive and negative effects, by alcohol than they expect themselves to be affected.

Predicting future, or long term, effects of substance use may be more difficult. Loewenstein (1999) argues that craving falls into the category of "visceral factors," which includes such other motivational states as hunger, thirst, or sexual arousal. He claims that people tend to underestimate not only the strength of visceral influences, but also their own susceptibility to them.

Even after negative effects of substance use manifest themselves, users can ignore or deny their existence. For example, "as cocaine addiction develops, a transition to high-dose long-duration bingeing occurs, in which the intensely pleasurable effects are experienced alone, and increasingly apparent negative contingencies go unrecognized" (Gawin, 1991, p. 1581). Addicts may not even realize that their consumption may be excessive, presumably due to tolerance. In a study of nurses and high school and college students, of respondents who were categorized by the researchers as heavy or excessive substance users, only 33.3% perceived their own use to be heavy or excessive (Champion and Bell, 1980).

2.3 Basic Model

The first basic feature of the model is the variable environment. In each period, the agent can find himself in one of two environments, environment A or environment B.[1] After observation of the environment, the agent allocates his/her resources between two goods: c, a non-addictive consumption good, and a, a potentially addictive consumption good. Assume that the choice variables c and a are continuous. Let a_t^A and c_t^A denote consumption of the potentially addictive good and the non-addictive good, respectively, when the environment at time t is environment A. Likewise, let a_t^B and c_t^B denote consumption of the potentially addictive good and the non-addictive good, respectively, under environment B.

Past consumption of the potentially addictive good in environment A is summarized by a stock variable x^A. Each period in which the environment is A ($\omega_t = A$), the compensatory process evolves and the stock variable is updated according to

$$x_{t+1}^A = \alpha x_t^A + \beta a_t^A$$

where $\alpha, \beta \in (0, 1)$. When the environment is B ($\omega_t = B$), this stock variable is unchanged:

$$x_{t+1}^A = x_t^A$$

[1]The main results are easily extended to $n > 2$ environments.

Past consumption of the potentially addictive good in environment B is summarized by a stock variable x^B which evolves as follows:

$$x_{t+1}^B = \alpha x_t^B + \beta a_t^B$$

when $\omega_t = B$, and

$$x_{t+1}^B = x_t^B$$

when $\omega_t = A$.

Initial stocks, x_0^A, x_0^B are exogenous. As in Loewenstein, et al. (1999), the consumption stocks can be thought of as levels of addiction. When the agent is in environment A, the addiction level associated with environment B, x^B, is dormant: the addiction level does not evolve, and utility is unaffected by x^B. Likewise, when the environment is B, the addiction level associated with environment A is dormant.

For now, I assume that the probability of the environment is exogenous:

$$\omega_t = \begin{cases} A & \text{with probability } \mu \\ B & \text{with probability } 1-\mu \end{cases}$$

where $\mu \in [0, 1]$.

The second basic feature of the model is the potential misperception by the individual of the underlying physiological changes caused by her consumption of the addictive good. The tolerance and withdrawal that an agent may experience is represented by a "tolerance function" $v(a, x)$ where $v(a, x) \leq 0$ for all a, x, with strict inequality if and only if $x > 0$. The agent misperceives the tolerance function as

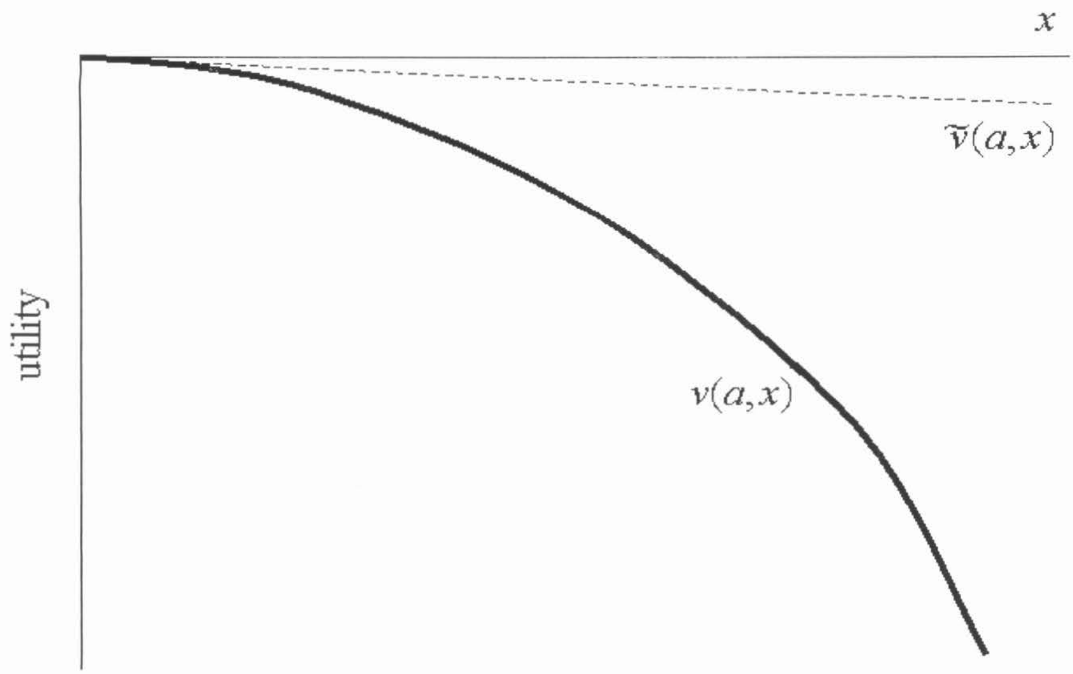

Figure 2.1: Example of true vs. perceived tolerance function

$\widetilde{v}(a,x)$ where, as with the true tolerance function, $\widetilde{v}(a,x) \leq 0$ for all a, x, with strict inequality if and only if $x > 0$. I assume that given a, $v(a,x)$ is weakly steeper, and more negative than $\widetilde{v}(a,x)$ for all x (see Figure 2.1). After sufficient experience with the addictive good, that is, when the addiction level, x, reaches some *threshold level*, $\widehat{x}$, the agent realizes the true tolerance function. I assume that this threshold level is exogenous and may be environment-specific.

If $x_t^j < \widehat{x}^j$, the agent believes that current and future instantaneous utility when

$\omega_t = j$, $j \in \{A, B\}$ is given by

$$\widetilde{U}(c_t^j, a_t^j, x_t^j) = u(c_t^j, a_t^j) + \widetilde{v}(a_t^j, x_t^j)$$

where $u_1, u_2 \geq 0$. However, when $x_t^j \geq \widehat{x}^j$, the agent learns that his true instantaneous utility when $\omega_t = j$ is given by

$$U(c_t^j, a_t^j, x_t^j) = u(c_t^j, a_t^j) + v(a_t^j, x_t^j)$$

The basic assumptions of the model are

1. $u\,(c, a)$ is twice continuously differentiable in c and a, and $v\,(a, x)$ and $\widetilde{v}(a, x)$ are twice continuously differentiable in a and x.

2. u is increasing and strictly concave in c and a; $u_1(c, a) \geq 0, u_2(c, a) \geq 0$ and $u_{11}(c, a) < 0$, $u_{22}(c, a) < 0$, $u_{11}(c, a) + u_{22}(c, a) < 2u_{12}(c, a)$.

3. The tolerance functions are negative, with $v(a, x)$ more negative than $\widetilde{v}(a, x)$: for all a, x, $v(a, x) \leq \widetilde{v}(a, x) \leq 0$, with $v(a, x) = 0$ iff $x = 0$ and $\widetilde{v}(a, x) = 0$ iff $x = 0$.

4. The tolerance functions are strictly increasing in a and decreasing in x, with $v(a, x)$ steeper than $\widetilde{v}(a, x)$ with respect to x: $v_1(a, x), \widetilde{v}_1(a, x) > 0$ for $a > 0$ and $x > 0$ and $v_2(a, x) \leq \widetilde{v}_2(a, x) \leq 0$.

5. $v(a, x)$ and $\widetilde{v}(a, x)$ are concave in a: $v_{11}(a, x), \leq 0$, $\widetilde{v}_{11}(a, x), \leq 0$.

6. a and c are complements: $u_{12} \geq 0$, Furthermore, the cross-partial derivatives between a and x are positive: $v_{12}, \widetilde{v}_{12} \geq 0$.

Income, y, and prices are assumed constant. Let c be the numeraire and let p be the price of the potentially addictive good. Define $\{X\}$ to be the indicator function that takes on the value 1 if the statement X is true; otherwise, it equals zero. The problem faced by an individual with discount rate δ and infinite time horizon is:

$$\max_{\{a_t^A, a_t^B, c_t^A, c_t^B\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \delta^t \left[\{\omega_t = A\} \left(u(c_t^A, a_t^A) + v(a_t^A, x_t^A) \right) + \{\omega_t = B\} \left(u(c_t^B, a_t^B) + v(a_t^B, x_t^B) \right) \right] \tag{2.1}$$

subject to

$$\begin{aligned} c_t^j + pa_t^j &\leq y \\ c_t^j &\geq 0 \end{aligned}$$

as well as the stochastic process on ω_t and the stock evolution equations:

$$x_{t+1}^A = \alpha x_t^A + \beta a_t^A$$

$$x_{t+1}^B = x_t^B$$

when $\omega_t = A$, and

$$x_{t+1}^B = \alpha x_t^B + \beta a_t^B$$

$$x_{t+1}^A = x_t^A$$

when $\omega_t = B$.

However, the problem that the individual believes that he needs to solve is

$$\max_{\{a_t^A, a_t^B, c_t^A, c_t^B\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \delta^t \left[\{\omega_t = A\} \left(u(c_t^A, a_t^A) + \widetilde{v}(a_t^A, x_t^A) \right) + \{\omega_t = B\} \left(u(c_t^B, a_t^B) + \widetilde{v}(a_t^B, x_t^B) \right) \right] \tag{2.2}$$

subject to the same constraints.

2.3.1 Discussion

Addiction Levels and Environments

Under the assumptions of the model, the budget constraint holds with equality ($c_t = y - pa_t$), and therefore we can focus on the consumption path of the addictive good. This result obtains because, for simplicity, saving and borrowing are not allowed. Therefore, all the intertemporal considerations enter the model through the consumption stock of the addictive good. Unlike the standard consumption problem in which the agent builds a stock of assets through savings, the agent in this model builds a stock based on past consumption of the addictive good. In the standard model, higher capital stock implies higher utility (under the usual assumptions of positive marginal utility of consumption and non-satiation). In the model of addiction, utility is decreasing in the consumption stock and, furthermore, there is no free disposal of the stock.

Previous models of addiction and habit formation (Ryder and Heal (1973), Becker and Murphy (1988), and Orphanides and Zervos (1995)) have also used a stock variable to summarize past consumption. However, the present model, which borrows from the framework used by Laibson (1999), accounts for the fact that environment can play an important role in shaping preferences for addictive goods. Neurobiological evidence presented in Section 2.1 suggests that tolerance and withdrawal do not necessarily occur purely as a result of past consumption, as the previous models assume. Instead, repeated consumption in a particular environment results in tolerance and withdrawal that are specific to that environment. This is captured in the present model through the use of environment-specific consumption stocks, or addiction levels, x_t^A and x_t^B.

Note that in this framework, the environment is neutral–the environment has no direct impact on utility. Utility and tolerance functions are constant across environments. Any effects of the environment enter only through the addiction levels and the threshold levels, $\widehat{x}^A$ and $\widehat{x}^B$.

Tolerance Function

The tolerance function is essential to incorporating withdrawal and tolerance into the model.

In Laibson's model, the addiction levels, or what he refers to as "the compensatory processes," enter into the utility function by directly offsetting consumption of the addictive good. That is, instantaneous utility is of the form:

$$f(c) + g(a - \lambda x)$$

whereas in my model, the addiction levels enter utility only through the tolerance function $v(a,x)$:

$$u(c,a)+v(a,x)$$

Note that this model nests Laibson's model.

Past consumption of the addictive good causes current disutility ($v\,(a,x)<0$ and $\widetilde{v}(a,x)<0$ iff $x>0$). If the agent abstains in the current period, she experiences withdrawal symptoms in the form of disutility. As the level of addiction increases, this disutility becomes more pronounced ($v_2(a,x), \widetilde{v}_2(a,x)\leq 0$).

However, these aversive effects can be "eased" by current consumption of the good ($v_1(a,x), \widetilde{v}_1(a,x)\geq 0$). Furthermore, as the stock increases, the appeal of the good as a mediator of craving increases, as represented by the positive cross partial derivative ($v_{12}, \widetilde{v}_{12}\geq 0$). In other words, preferences display what Becker and Murphy refer to as "adjacent complementarity"[2], as in the "rational" models of addiction.

Lastly, note how tolerance and withdrawal operate through the same mechanism: if the agent chooses to consume the addictive good, the net utility derived from a fixed dose of the addictive good is diminished by past consumption: for all feasible $\overline{c},\overline{a}$ $u(\overline{c},\overline{a})>u(\overline{c},\overline{a})+v(\overline{a},x)>u(\overline{c},\overline{a})+v(\overline{a},x')$ where $0<x<x'$.

In their paper, O'Donoghue and Rabin discuss two characteristics of addictive goods: they are habit-forming, and they involve "internalities." Both these characteristics are represented by the tolerance function.

[2]The term "adjacent complementarity" appears to have been coined by Ryder and Heal (1973) and referred to complementarity between consumption on adjacent dates, rather than consumption on distant dates. As Becker and Murphy use it, "adjacent complementarity" simply refers to complementarity between current consumption and the stock of past consumption.

Current consumption of a habit forming good will increase future marginal utility of consumption. In terms of my model, this is represented by the assumption that the cross partial derivative between the stock variable and current consumption is strictly positive. A good has internalities if current consumption affects the future level of instantaneous utility from consumption of the good. For example, tolerance is a negative internality. That is, current consumption decreases the future utility level from consumption. The assumptions on the tolerance function imply that the good in question has negative internalities.

Laibson's model implicitly assumes that the addictive good in question is one that is habit forming and has negative internalities. In the present model, these two facets of the addictive good can be separated, even though this feature is not taken advantage of in this paper. In order to generalize the model to goods that are not necessarily harmfully addictive substances, this sort of separability is necessary. There are goods that may be habit forming, but have positive internalities (exercise, for example). Alternatively, there are goods that may have negative internalities, but are not habit forming (overeating at a meal, for example).

Misperception of Tolerance Function

The assumptions on the the relationship between the true tolerance function and the misperceived tolerance function imply that the agent *underestimates* the negative consequences from current consumption. Of course, there are individuals who misperceive the future effects of current consumption in the opposite direction. That is, they *overestimate* the future effects. Recall that the studies of Agostinelli and Miller

(1994) and Champion and Bell (1980) find an inverse relationship between substance or alcohol use and perceived risk of using addictive substances. For this paper, I focus only on agents who underestimate the future consequences from using addictive substances, in part because those who overestimate the effects are very unlikely to become addicts. However, a generalization of the model to include consumption of goods that are not necessarily harmfully addictive substances might want to incorporate those who overestimate the future effects of consumption.

One interpretation of the misperception of the true tolerance function is as follows: the function $v(a, x)$ is approximated by $\widetilde{v}(a, x)$, given a, when x is close to x_0. For example, $\widetilde{v}(a, x)$ may be the linear Taylor approximation to $v(a, x)$, as in Figure 1. While the agent has little experience at consuming the addictive product, the tolerance and withdrawal symptoms that result from past consumption are relatively minor. However, the agent suffers from *projection bias*[3]–she underestimates changes in future utility from the present. She assumes that the process that governs the negative side effects from consumption of the addictive good will continue into the future as it has in the past. She does not realize that, after sufficient consumption of the addictive good, changes in her physiological systems cause tolerance and withdrawal symptoms to increase dramatically. Under this interpretation, the threshold level can be thought of as the point at which the true tolerance function $v(a, x)$ and the perceived tolerance function $\widetilde{v}(a, x)$ begin to diverge.

An alternative interpretation involves a heterogeneous population of agents. Suppose the population consists of two groups: one group for whom $\widetilde{v}(a, x)$ is the true

[3]For further discussion on projection bias, see Loewenstein, et al. (1997).

tolerance function, and another group for whom $v(a, x)$ is the true tolerance function. The population of interest in this paper are those for whom $v(a, x)$ is the true tolerance function, but they initially believe that they are of the group for whom $\widetilde{v}(a, x)$ is the true tolerance function. Initially these agents believe with probability one that their tolerance function is $\widetilde{v}(a, x)$. After sufficient consumption of the addictive good–that is, once $x > \widetilde{x}$, the agent updates her beliefs and believes with probability one that her tolerance function is given by $v(a, x)$.

In this framework, regardless of the interpretation, the agent's "learning" of the true tolerance function is very simple–the agent is completely unaware of the true tolerance function before his addiction level reaches the threshold level, after which he perfectly foresees the future effects of current consumption. In Section 5, I consider two alternative frameworks in which the agent slowly learns, or adjusts to, the true tolerance function.

Addiction

Although there are numerous definitions of addiction (for an overview, see Pomerleau and Pomerleau, 1988), most agree that addiction is characterized by prolonged compulsive use, tolerance, and physical and/or psychic dependence. In terms of the model, it seems that a reasonable baseline is the consumption levels and consumption stock of a hypothetical individual who does not experience tolerance and dependence, an individual for whom instantaneous utility is simply $u(c, a)$ rather than $u(c, a)+v(a, x)$.

Assume then, that an agent's consumption is compulsive if and only if

$$a > a^{NA}$$

where a^{NA} is the optimal consumption of the hypothetical individual:

$$-pu_1(y - pa^{NA}, a^{NA}) + u_2(y - pa^{NA}, a^{NA}) = 0$$

An individual is "addicted" if and only if his addiction level is higher than the baseline defined by a^{NA}; that is, if and only if

$$x_t > x_t(a^{NA}) = \alpha^t x_0 + \sum_{i=0}^{t-1} \alpha^i \beta a^{NA}$$

Susceptibility to Addiction

Clearly, individuals are not homogeneous in their susceptibility to addiction. The probability that an individual becomes addicted is influenced by a host of exogenous factors that may be genetic, social, or environmental. Many of these factors can be captured in the model.

Consider the initial addiction levels, x_0^A, x_0^B. These can reflect any genetic tendency toward addiction. In the extreme case of children who are born addicted to a substance, they can reflect the degree to which the child is addicted at birth. Alternatively, suppose the decision-making process regarding addictive substances begins not at birth, but later in life, such as adolescence. In this case, the initial addiction levels can represent any first-hand or second-hand experience with addictive substances prior to the adolescent years. In this latter example, it may be plausible that the initial addiction levels may vary across environments; however, different interpretations may suggest otherwise. Therefore, I assume $x_0^A = x_0^B$.

An individual's genetic tendency toward addiction may also be reflected in his true tolerance function. For example, the second derivative of the tolerance function may increase in magnitude as genetic tendency toward addiction increases. Or the appeal of the addictive substance as a mediator of craving (the cross partial derivative of v between a and x) may increase with genetic tendency.

The threshold levels, $\widehat{x}^A, \widehat{x}^B$, or the degree of misperception of the true tolerance function may be in part determined by the individual's personality. As discussed above, the misperception of the true effects of addictive substances may be due to projection bias. The threshold levels represent the point at which the agent realizes the true tolerance function.

I allow the threshold levels, $\widehat{x}^A, \widehat{x}^B$, to capture any difference in susceptibility across environments. For example, events that elicit dysphoric or euphoric states may sensitize an individual to the direct effects of addictive substances (Pomerleau and Pomerleau, 1988). In certain environments, it may take an agent longer to realize the degree of his substance abuse than in others. Alternatively, this phenomenon could be captured by having the *perceived* tolerance function vary across environments. Initially, I assume that the thresholds vary, but the perceived tolerance functions do not.

2.4 Optimal Behavior and Dynamics

2.4.1 Single Environment Model

To facilitate the analysis, I begin by analyzing optimal behavior in the world in which there is only one possible environment. Before I characterize the optimal policy function for the single environment (SE) model, consider the first order conditions in order to gain an understanding of optimal behavior.

Initially, the agent believes that she must solve:

$$\widetilde{V}^{SE}(x_0) = \max_{\{a_t,c_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \delta^t \left[u(y - pa_t, a_t) + \widetilde{v}(a_t, x_t)\right] \tag{2.3}$$

where x_0 is given. Assuming an interior solution, the first order condition to the problem in Equation (2.3) is

$$u_2(y - pa_t, a_t) + \widetilde{v}_1(a_t, x_t) = pu_1(y - pa_t, a_t) + \sum_{i=1}^{\infty} \delta^i \alpha^{i-1} \beta \widetilde{v}_2(a_{t+i}, x_{t+i}) \tag{2.4}$$

Each period, the individual weighs the benefit from consuming the addictive good, current marginal utility, against what Becker and Murphy call the full price of the addictive good. The full price includes the price of the addictive good, as well as the marginal effects of current consumption on future utility. The agent realizes that current consumption has a detrimental effect on future utility. However, when $x < \widehat{x}$ the agent does not realize the extent of these effects.

Note that, even though the agent may know the true utility level at time t, the agent does not necessarily know the true marginal utility that he will experience as a result of continued consumption at time t. In making the consumption decision at time t, the agent is unaware of the true marginal effect of current consumption on current utility: the agent's first order condition has $u_2(y - pa_t, a_t) + \widetilde{v}_1(a_t, x_t)$ rather than $u_2(y-pa_t, a_t)+v_1(a_t, x_t)$ on the left-hand side. Depending on the form of $\widetilde{v}(a_t, x_t)$, the agent could overestimate, or even underestimate, the current marginal utility from current consumption.[4]

When $x > \widehat{x}$, the agent realizes the true tolerance function, and the problem that the individual must now solve is

$$V^{SE}(x_\tau) = \max_{\{a_t,c_t\}_{t=\tau}^{\infty}} E_0 \sum_{t=\tau}^{\infty} \delta^t \left[u(c_t, a_t) + v(a_t, x_t)\right] \tag{2.5}$$

where τ denotes the first period after the change in the tolerance function. The appropriate first order condition is:

$$u_2(y - pa_t, a_t) + v_1(a_t, x_t) = pu_1(y - pa_t, a_t) + \sum_{i=1}^{\infty} \delta^i \alpha^{i-1} \beta v_2\left(a_{t+i}, x_{t+i}\right)$$

The individual's maximization problems can be recast as stationary dynamic programming problems. While $x < \widehat{x}$, the Bellman equation is:

$$\widetilde{V}^{SE}(x) = \max_{a \in [0, \frac{y}{p}]} \left[u(y - pa, a) + \widetilde{v}(a, x) + \delta \widetilde{V}^{SE}(\alpha x + \beta a)\right] \tag{2.6}$$

[4]Recall that, even though there are assumptions on the relationship between v_2 and $\widetilde{v}_2$, there are no restrictions on the relationship between v_1 and $\widetilde{v}_1$.

After the agent realizes the true tolerance function, the Bellman equation is given by:

$$V^{SE}(x) = \max_{a\in[0,\frac{y}{p}]} \left[u(y - pa, a) + v(a, x) + \delta V^{SE}(\alpha x + \beta a)\right] \tag{2.7}$$

The assumptions on utility and the Theorem of the Maximum ensure the existence, uniqueness, and differentiability of $\widetilde{V}^{SE}(x)$ and $V^{SE}(x)$, as well as the existence of non-empty upper semi-continuous policy correspondences

$$\begin{aligned}
\widetilde{\phi}(x) &= \left[x'|\widetilde{V}^{SE}(x) = u(y - \frac{p}{\beta}(x' - \alpha x), \frac{1}{\beta}(x' - \alpha x)) + \widetilde{v}(\frac{1}{\beta}(x' - \alpha x), x) + \delta \widetilde{V}^{SE}(x')\right] \\
\phi(x) &= \left[x'|V^{SE}(x) = u(y - \frac{p}{\beta}(x' - \alpha x), \frac{1}{\beta}(x' - \alpha x)) + v(\frac{1}{\beta}(x' - \alpha x), x) + \delta V^{SE}(x')\right]
\end{aligned}$$

Standard dynamic programming techniques can not be used to characterize $\widetilde{\phi}(x)$ and $\phi(x)$ for two reasons. First, utility is not assumed to be strictly concave in $c, a,$ and x. Second, utility is decreasing in the stock variable, in contrast to the standard production or consumption dynamic programming problems. Furthermore, there is no free disposal of the stock.

Both $\widetilde{\phi}(x)$ and $\phi(x)$, the optimal stock evolutions, can be characterized by the following proposition, which draws heavily from Orphanides and Zervos, 1994:

Proposition 1 *For an agent with value function given by Equation (2.6) or (2.7), (i) every optimal path is a monotonic sequence; (ii) any optimal path converges to a steady state; and (iii) there exists exactly one critical level between any two consecutive stable steady states.*

Monotonicity follows from the fact that time t marginal utility with respect to time t stock is increasing in time $t+1$ stock. Convergence to a steady state then follows because the optimal path is a bounded monotonic sequence. The critical value may or may not be an unstable steady state. When the critical value is not an unstable steady state, the possibility of multiple optimal paths arises. Such a critical value exists because of the complementarity between current consumption of the addictive good and the stock of past consumption. All proofs are in the appendix.

Suppose the optimal paths associated with equations 2.6 and 2.7 have one critical value ($\widetilde{x}^*$ and x^*, respectively) between two stable steady states, as pictured in Figure 2.2[5]. The optimal consumption correspondence for the misperceived maximization problem, $\widetilde{a}^*(x) = \frac{1}{\beta}(\widetilde{\phi}(x) - \alpha x)$, is represented by the medium bold correspondence, and $a^*(x) = \frac{1}{\beta}(\phi(x) - \alpha x)$, the consumption policy for the agent's true problem, is represented by the heavy bold correspondence. The steady states that correspond to Equation (2.6) are given by $\underline{\widetilde{x}}$ and $\overline{\widetilde{x}}$, while the steady states that correspond to Equation (2.7) are $\underline{x}$ and $\overline{x}$. For ease of exposition, denote the lower steady states $\underline{\widetilde{x}}$ and $\underline{x}$ as the no-addiction steady states, and the higher steady states as the addiction steady states.

When the critical value associated with Equation (2.6) is less than the critical level associated with equation 2.7 ($\widetilde{x}^* < x^*$), as in this example, whether the agent ultimately enters a state of addiction depends on the initial addiction level, x_0, and the

[5]In this example, $u(c,a) = \ln c + \ln a$, $v(a,x) = 5.6ax - 6x$ and $\widetilde{v}(a,x) = 5.6ax - 5.6x$. I assume $y = p = 1$, $\alpha = 0.5$, $\beta = 1$ and $\delta = 0.9$. The policy correspondence associated with the true tolerance function, $a^*(x)$, has two stable steady states $\underline{x} = 0.23$ and $\overline{x} = 1.79$ and a critical value $x^* = 0.94$. The policy correspondence associated with the perceived tolerance function, $\widetilde{a}^*(x)$, has two stable steady states $\underline{\widetilde{x}} = 0.26$ and $\overline{\widetilde{x}} = 1.81$ and a critical value $\widetilde{x}^* = 0.83$

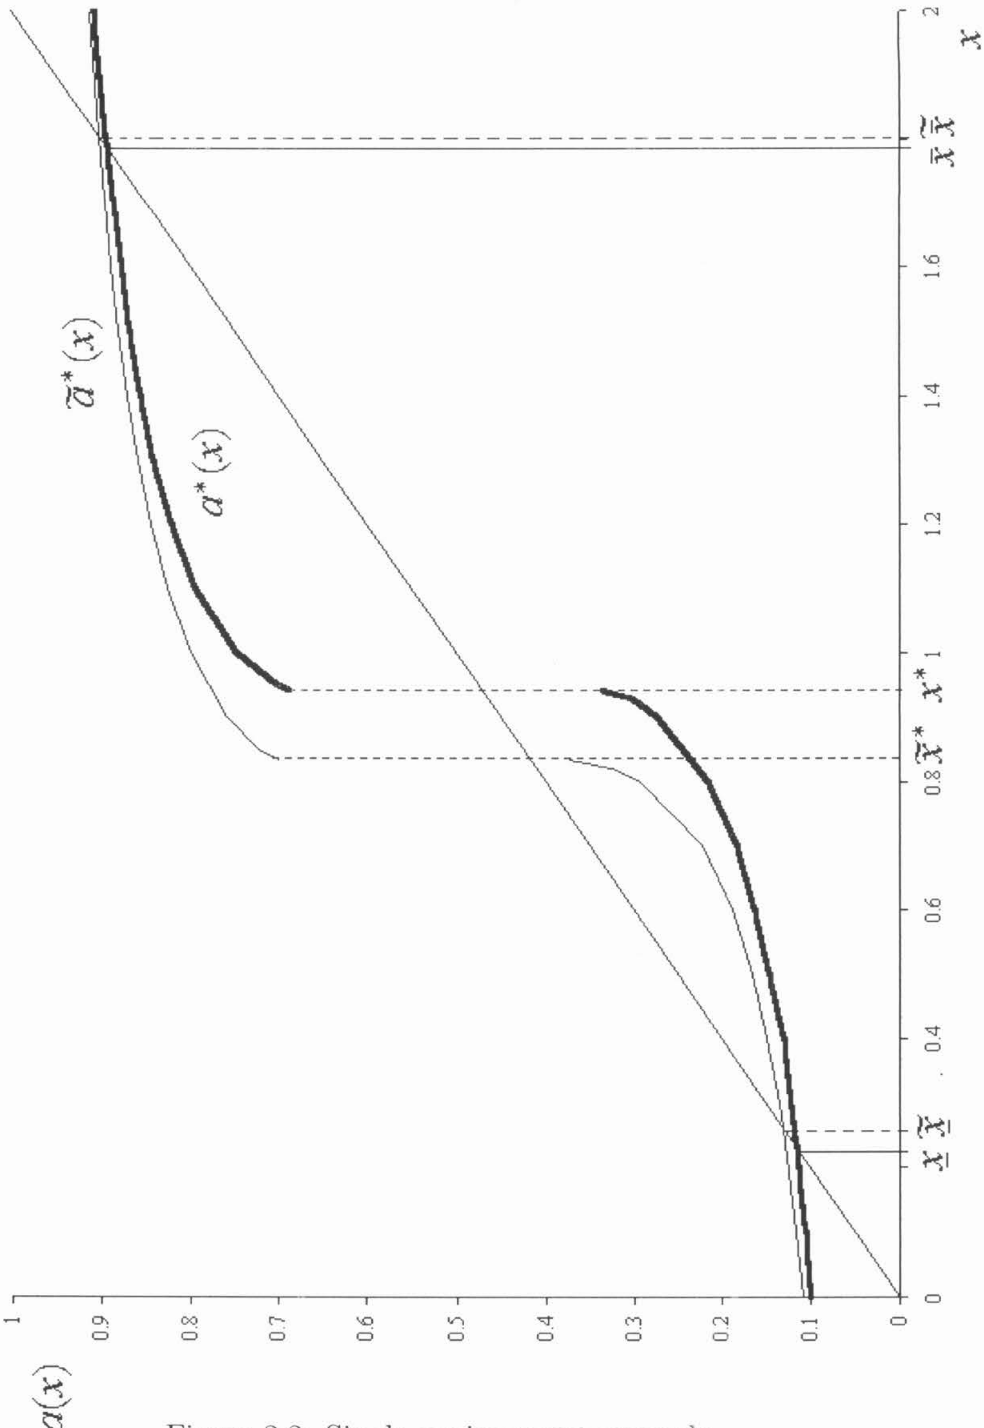

Figure 2.2: Single environment example

threshold, $\widehat{x}$. If $x_0 < \widetilde{x}^* < x^*$, the agent will never become addicted, regardless of the threshold. The optimal stock converges toward $\underset{\sim}{\underline{x}}$ or $\underline{x}$ depending on the relationship between x_0 and $\widehat{x}$. On the other hand, $x_0 > x^* > \widetilde{x}^*$ implies that the optimal stock will converge toward $\widetilde{\overline{x}}$ or $\overline{x}$, both of which yield harmful addictions.

If $\widetilde{x}^* < x_0 < x^*$, the agent will become addicted if and only if $\widehat{x} > x^*$. Initially, the agent's consumption stock will begin to approach the higher steady state $\widetilde{\overline{x}}$ (assuming that $\widehat{x} > x^0$). However, before this steady state is reached, the threshold level $\widehat{x}$ will be reached. At $\widehat{x}$, the agent will realize the true tolerance function, and his optimal path will shift. If $\widehat{x} > x^*$, then the agent will continue consuming the addictive good and his addictive level will monotonically approach the steady state labelled as $\overline{x}$. On the other hand, if $\widehat{x} < x^*$, then the agent's consumption of the addictive good will decrease and his stock will converge to the steady state labelled by $\underline{x}$. Therefore, if $\widehat{x} > x^*$ the individual will move toward a state of addiction, whereas, if $\widehat{x} < x^*$, the individual will essentially "quit." In this example, at the addiction steady state, the agent's instantaneous utility every period is -2.80 whereas instantaneous utility at the no-addiction steady state is -2.23.

The case when $\widetilde{x}^* < x_0 < x^*$ and $\widehat{x} > x^*$ illustrates how it is possible for the agent's consumption of the addictive good to converge to the steady state associated with the state of addiction, although had he known the true tolerance function, his consumption would never had progressed as far. These addicts are similar to the "regretful addicts" of Orphanides and Zervos (1995).

The agent for whom $\widetilde{x}^* < x_0 < x^*$ and $\widehat{x} < x^*$ "experiments" with the addictive good, but does not ultimately become addicted. When he is still naive to the true effects of the addictive good, he begins to consume the addictive good. When he then learns of the true effects of his past consumption, he decreases his consumption.

This example also illustrates that the agent's lifetime utility is weakly decreasing in the threshold level $\widehat{x}$. The earlier that the agent realizes the true tolerance function, the better off the agent is made, because he can then maximize his true lifetime consumption problem, rather than the incorrect, perceived problem.

Notice that in the single environment model, as in the O-Z model, agents can quit their addiction at most once, with no relapse, as illustrated by the previous case.

2.4.2 Dual Environment Model

Now consider the problem given by Equations (2.1) and (2.2). When $x_t^A < \widehat{x}^A$ and $x_t^B < \widehat{x}^B$, the first order conditions for the solution to the problem given by Equation (2.2) are

$$u_2(y - pa_t^A, a_t^A) + \widetilde{v}_1(a_t^A, x_t^A) = pu_1(y - pa_t^A, a_t^A) + \sum_{i=1}^{\infty} \mu \delta^i \alpha^{i-1} \beta \widetilde{v}_2(a_{t+i}^A, x_{t+i}^A) \quad (2.8)$$

when the environment is A, and

$$u_2(y - pa_t^B, a_t^B) + \widetilde{v}_1(a_t^B, x_t^B) = pu_1(y - pa_t^B, a_t^B) + \sum_{i=1}^{\infty} (1 - \mu) \delta^i \alpha^{i-1} \beta \widetilde{v}_2(a_{t+i}^B, x_{t+i}^B) \quad (2.9)$$

when the environment is B. After the agent realizes the true tolerance function, the first order conditions are

$$u_2(y - pa_t^A, a_t^A) + v_1(a_t^A, x_t^A) = pu_1(y - pa_t^A, a_t^A) + \sum_{i=1}^{\infty} \mu\delta^i\alpha^{i-1}\beta v_2(a_{t+i}^A, x_{t+i}^A) \quad (2.10)$$

when the environment is A, and

$$u_2(y - pa_t^B, a_t^B) + v_1(a_t^B, x_t^B) = pu_1(y - pa_t^B, a_t^B) + \sum_{i=1}^{\infty} (1-\mu)\,\delta^i\alpha^{i-1}\beta v_2(a_{t+i}^B, x_{t+i}^B) \quad (2.11)$$

when the environment is B.

Once the environment is revealed, the agent compares marginal utility of consumption with the full price of consumption, as before. However, in the dual environment (DE) model, the full price of consumption is less than in the SE model. Future considerations are further discounted because current consumption will not affect every future period. Suppose that the environment is A at time t. Current consumption is only relevant in the future periods in which the environment is also A. Therefore, marginal utility at time $t+\tau$ with respect to time t consumption, $\alpha^{\tau-1}\beta v_2(a_{t+\tau}^A, x_{t+\tau}^A)$, is also discounted by the probability that the environment will be A at time $t+\tau$.

The corresponding Bellman equation to Equation (2.2) is:

$$\begin{aligned} \widetilde{V}(x^A, x^B) &= \max_{a^A, a^B} \mu\left[u(y - pa^A, a^A) + \widetilde{v}(a^A, x^A) + \delta\widetilde{V}(\alpha x^A + \beta a^A, x^B)\right] \quad (2.12) \\ &\quad + (1-\mu)\left[u(y - pa^B, a^B) + \widetilde{v}(a^B, x^B) + \delta\widetilde{V}(x^A, \alpha x^B + \beta a^B)\right] \end{aligned}$$

and the Bellman equation corresponding to the true problem, given by Equation (2.1), is:

$$\begin{aligned} V(x^A, x^B) &= \max_{a^A, a^B} \mu \left[u(y - pa^A, a^A) + v(a^A, x^A) + \delta V(\alpha x^A + \beta a^A, x^B) \right] \quad (2.13) \\ &\quad + (1-\mu) \left[u(y - pa^B, a^B) + v(a^B, x^B) + \delta V(x^A, \alpha x^B + \beta a^B) \right] \end{aligned}$$

Notice that these value functions reflect the agent's welfare before the realization of the environment. Again, the assumptions on utility, together with the Theorem of the Maximum, ensure that $\widetilde{V}(x^A, x^B)$ and $V(x^A, x^B)$ exist and are unique.

These value functions can be expressed in terms of the single environment value functions.

Proposition 2

$$\widetilde{V}(x^A, x^B) = \frac{\mu}{1-\delta(1-\mu)} \widetilde{V}^{SE}(x^A | \widetilde{\delta} = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu} \widetilde{V}^{SE}(x^B | \widetilde{\delta} = \frac{\delta(1-\mu)}{1-\delta\mu})$$

and

$$V(x^A, x^B) = \frac{\mu}{1-\delta(1-\mu)} V^{SE}(x^A | \widetilde{\delta} = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu} V^{SE}(x^B | \widetilde{\delta} = \frac{\delta(1-\mu)}{1-\delta\mu})$$

where $\widetilde{V}^{SE}(x|\widetilde{\delta} = \xi)$ and $V^{SE}(x|\widetilde{\delta} = \xi)$ are the solutions to Equations (2.6) and (2.7), respectively, given discount rate ξ.

The intuition behind the proposition is sketched out here. The proof, which closely follows the same line of reasoning as Proposition 3 of Laibson (1999), is in the appendix. Proposition 2 relies on the separability of utility with respect to the environment-specific addiction levels. That is, when the environment is A (or B), the addictive level associated with environment B (or A) is not updated and does not affect current utility.

The proposition reflects the intuition derived from the first-order conditions regarding discounting of future periods. Any given future period is discounted not only by the rate of time preference, but also by the probability that current consumption will not have an effect on that period. In essence, $\frac{\delta\mu}{1-\delta(1-\mu)} < \delta$ is the discount rate between the current period and the next period in which the environment is A. Likewise, $\frac{\delta(1-\mu)}{1-\delta\mu} < \delta$ is the discount rate between the current period and the next period in which the environment is B.

The optimal policy correspondences, conditional on the environment, can then be written as functions of the relevant addiction level only. The optimal consumption policy correspondences associated with the functional equation given by Equation (2.12) when the environment is A or B, respectively, are:

$$\begin{aligned}
\widetilde{\psi}^A(x) &= \Big\{x'|\widetilde{V}^{SE}(x|\frac{\delta\mu}{1-\delta(1-\mu)}) = u(y-\frac{p}{\beta}(x'-\alpha x), \frac{1}{\beta}(x'-\alpha x)) + \widetilde{v}(\frac{1}{\beta}(x'-\alpha x), x) \\
&\quad + \frac{\mu}{1-\delta(1-\mu)}\widetilde{V}^{SE}(x'|\frac{\delta\mu}{1-\delta(1-\mu)})\Big\} \\
\widetilde{\psi}^B(x) &= \Big\{x'|\widetilde{V}^{SE}(x|\frac{\delta(1-\mu)}{1-\delta\mu}) = u(y-\frac{p}{\beta}(x'-\alpha x), \frac{1}{\beta}(x'-\alpha x)) + \widetilde{v}(\frac{1}{\beta}(x'-\alpha x), x) \\
&\quad + \frac{1-\mu}{1-\delta\mu}\widetilde{V}^{SE}(x'|\frac{\delta(1-\mu)}{1-\delta\mu})\Big\}
\end{aligned}$$

and the optimal policy correspondences associated with equation (2.13), when the environment is A or B, are:

$$\begin{aligned}
\psi^A(x) &= \Big\{x'|V^{SE}(x|\frac{\delta\mu}{1-\delta(1-\mu)}) = u(y-\frac{p}{\beta}(x'-\alpha x), \frac{1}{\beta}(x'-\alpha x)) + v(\frac{1}{\beta}(x'-\alpha x), x) \\
&\quad + \frac{\mu}{1-\delta(1-\mu)}V^{SE}(x'|\frac{\delta\mu}{1-\delta(1-\mu)})\Big\} \\
\psi^B(x) &= \Big\{x'|V^{SE}(x|\frac{\delta(1-\mu)}{1-\delta\mu}) = u(y-\frac{p}{\beta}(x'-\alpha x), \frac{1}{\beta}(x'-\alpha x)) + v(\frac{1}{\beta}(x'-\alpha x), x) \\
&\quad + \frac{1-\mu}{1-\delta\mu}V^{SE}(x'|\frac{\delta(1-\mu)}{1-\delta\mu})\Big\}
\end{aligned}$$

Because the dual environment value functions are essentially linear combinations of the single environment value functions (with modified discount rates), we can also characterize these optimal stock evolutions as in proposition 1.

Proposition 3 *The optimal stock evolutions associated with the Bellman Equations given by Equations (2.12) and (2.13), $\widetilde{\psi}^A(x)$, $\widetilde{\psi}^B(x), \psi^A(x)$ and $\psi^B(x)$, can each be characterized as follows: (i) every optimal path is a monotonic sequence; (ii) any optimal path converges to a steady state; and (iii) there exists exactly one critical level between any two consecutive stable steady states.*

Using an example where the optimal consumption path has two steady stable states, it is easy to see how it is possible for an agent's consumption to converge to the addictive state in one state of the world, but not in the other.

Consider the example illustrated in Figure 2.3.[6] As in the SE Model example, the optimal consumption correspondences for Equation (2.12), $\widetilde{a}^{A*}(x) = \frac{1}{\beta}(\widetilde{\psi}^A(x) - \alpha x)$,

[6] $u(c,a) = \ln c + \ln a$, $v(a,x) = 5.6ax - 6x$ and $\widetilde{v}(a,x) = 5.6ax - 5.6x$, $y = p = 1$, $\alpha = 0.5$, $\beta = 1$, $\delta = 0.95$ and $\mu = 0.5$. $a^{A*}(x)$ and $a^{B*}(x)$ each have two stable steady states $\underline{x} = 0.23$ and $\overline{x} = 1.79$ and a critical value $x^* = 0.94$. $\widetilde{a}^{A*}(x)$ and $\widetilde{a}^{B*}(x)$ each have two stable steady states $\widetilde{\underline{x}} = 0.26$ and $\widetilde{\overline{x}} = 1.81$ and a critical value $\widetilde{x}^* = 0.83$.

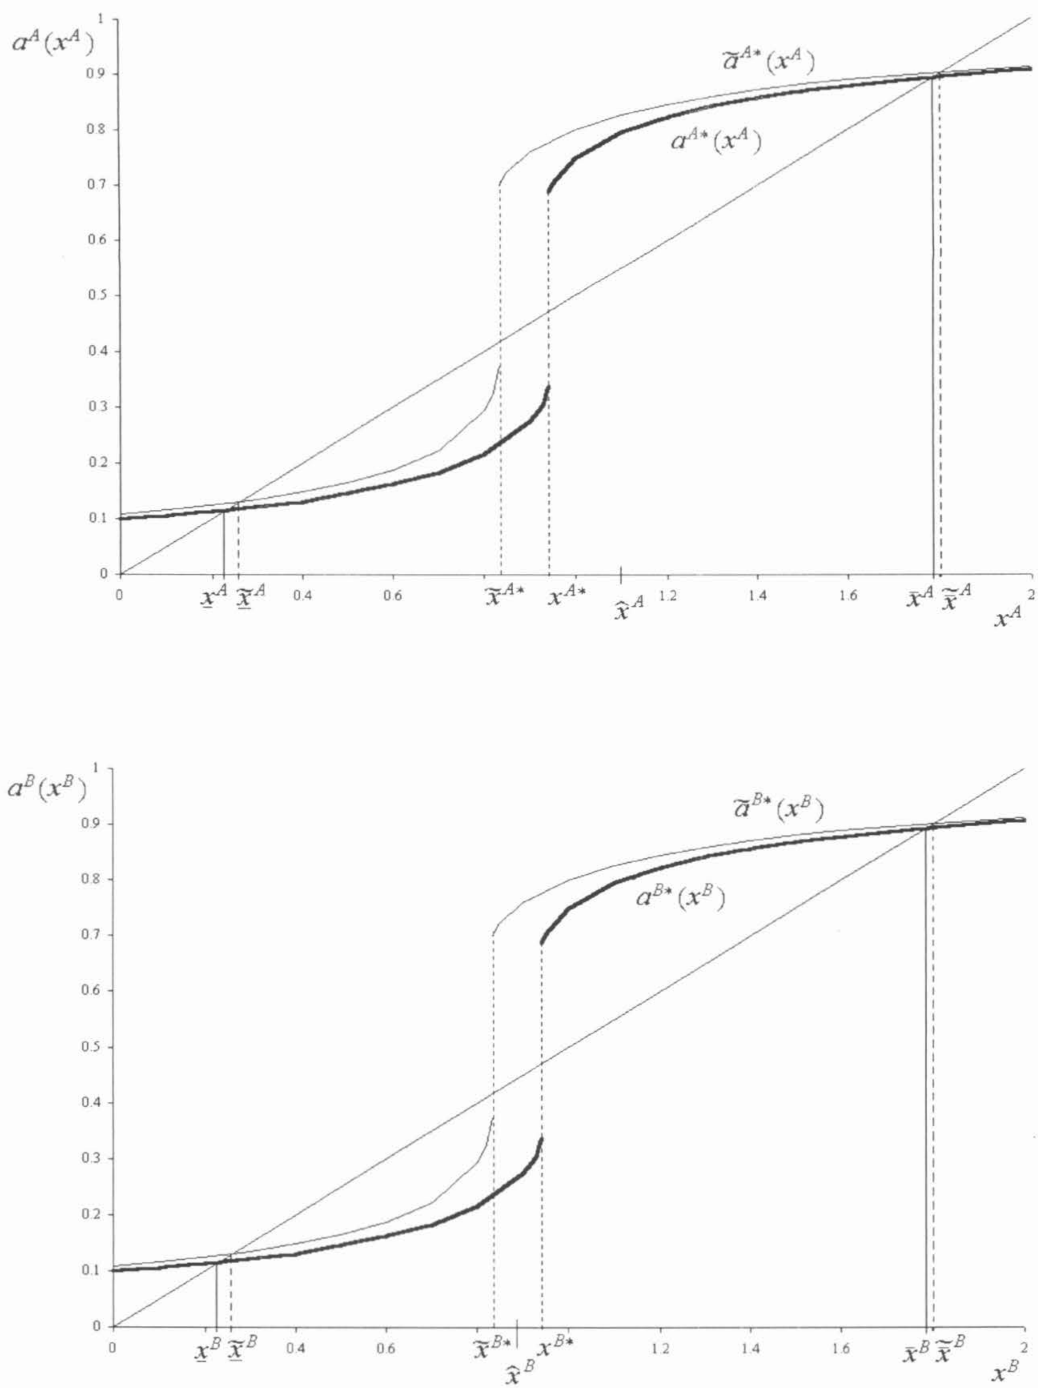

Figure 2.3: Dual environment example. The top panel represents Environment A and the bottom panel represents Environment B.

and $\widetilde{a}^{B*}(x) = \frac{1}{\beta}(\widetilde{\psi}^{B}(x) - \alpha x)$ are represented by the medium bold correspondence in their respective figures, and $a^{A*}(x) = \frac{1}{\beta}(\psi^{A}(x) - \alpha x)$ and $a^{B*}(x) = \frac{1}{\beta}(\psi^{B}(x) - \alpha x)$, the consumption policies for Equation (2.13), are represented by the heavy bold correspondence.

Assume that the threshold level in environment B, $\widehat{x}^{B}$, is less than the threshold level in environment A, $\widehat{x}^{A}$. In particular, suppose that $\widetilde{x}^{*} < \widehat{x}^{B} < x^{*}$ whereas $\widetilde{x}^{*} < x^{*} < \widehat{x}^{A}$, as illustrated in the Figure 2.3 above. In this case, once the addictive level that is activated in environment A reaches the threshold level, $\widehat{x}^{A}$, the agent will continue consuming the addictive good, and the addiction level will converge to the higher steady state $\overline{\overline{x}}^{A}$. On the other hand, when the agent realizes the true tolerance function in environment B, the agent will essentially quit her addiction in this state of the world.

In this example, there are four possible outcomes. One outcome is low consumption of the addictive good, regardless of the environment, which corresponds to the no-addiction steady state in the SE model. Another outcome is high consumption of the addictive good, regardless of the state of the world, which corresponds to the addiction steady state in the SE Model. The last two outcomes involve high consumption of the addictive good in one state of the world, but not the other.

The two outcomes that result in addiction in one environment but not the other are the outcomes that can not occur in the models such as those of B-M and O-Z, but they are essential to capturing the pattern of quitting and relapsing that is so prevalent. The agent behaves as if she has quit her addiction in Environment B

(a rehabilitation program, for example), but once she returns to Environment A (a neighborhood where she frequently consumed the addictive good), she continues to consume. Notice that there are other interpretations of these outcomes. For example, consider binge behavior. "Bingeing" is characterized by periods of high consumption alternating with periods of abstinence. Cocaine addiction is characterized by binge behavior (see for example, Gawin, 1991). It is not uncommon for cocaine addicts to function normally throughout the working day or week, and then to consume high doses of cocaine after working hours or on weekends. In terms of the model, in the working or office environment, the agent is at the lower of the two steady states, the one associated with no addiction. In the after-hours or weekend environment, however, the agent is at the addiction steady state.

These outcomes that correspond to high consumption in one environment but not the other demonstrate how environment can come to have a very real impact on behavior even though utility and all exogenous variables, other than threshold levels, are constant across environments. Discovery of the true consequences of substance use early on implies that it is optimal or "rational" for the agent to quit using the substance, and therefore, consume very little, or none, of the substance. In the other environment, realization of the true consequences when it is "too late" implies that the optimal or "rational" consumption path for the agent leads to addiction. However, notice that because preferences are constant across environments, if preferences are such that the agent chooses to quit in one environment, then if consumption in the other environment converges to the addiction steady state, the agent "regrets" this addiction. If the agent had known her true preferences, she would be at the lower steady state regardless of environment, where utility is higher than at the addiction steady state.

One interpretation of these outcomes might be that the agent chooses to quit in one environment, but, in the other environment, he chooses to continue his addiction. This interpretation seems contrary to observation of addicts who seek treatment. Apparently, they would like to quit regardless of environment. A more realistic interpretation of the behavior generated by the model is that the agent chooses to quit in one environment, but finds it too painful to quit in the other. However, if possible, he would seek to avoid the "addiction" environment, because his utility is higher in the environment in which he has quit. For example, many recovery programs advocate eliminating the "addiction" environment altogther, if possible (Frawley, 1988). Although the environment is exogenously determined in this model, an extension to allow for endogenous determination of the model could predict such behavior.

The agent in this model essentially behaves as if he has two personalities.[7] One personality corresponds to environment A, while the other corresponds to environment B. The consumption decisions that one personality makes do not affect the other, and therefore each personality has utility that is separate and independent of the other. If the agent's behavior approaches the outcome where his consumption of the addictive good is high in one environment but low in the other, then his two

[7] The dual personality that I describe here differs from the typical "divided self" in the literature where agents with self control problems are modelled as having multiple selves. One (or more) selves are myopic, and one (or more) selves are more forward looking. In this model, both selves are forward looking, but unconcerned with the utility of the other. See, for example, Winston (1980), Thaler and Shefrin (1981) for applications to addiction.

personalities can be interpreted as one that is operational when craving strikes, and another that is operational at all other times.

2.5 Learning and Tolerance

2.5.1 Gradual Learning

As pointed out in Section 3.1 of this chapter, the agent's assumed process of learning the true tolerance function is somewhat simplistic. Suppose instead that the agent gradually learns the true tolerance function, $v\left(a,x\right)$. As an additional modification to the original model, assume that the perceived tolerance function is not constant across environments. For example, the functional form may be constant across environments, but the parameters are not.

The agent initially believes that the tolerance function in environment j is $\widetilde{v}^{j0}(a,x)$, where the value $\widetilde{v}^{j0}(a,x_0^j) = v^j(a,x_0^j)$.

Every period in which the environment is j, the agent experiences the true tolerance or craving associated with his addiction level and updates his perceived tolerance function associated with environment j in such a way that the perceived tolerance function and the true tolerance function are equal for all previously realized (a,x) pairs. That is,

$$\widetilde{v}^{jt}(a,x) = v(a,x)$$

for all $(a,x) = (a_i^j, x_i^j)$, $\forall i \leq t$ such that $\omega_i = j$.

Lastly, assume that $\widetilde{v}^{jt}(a,x)$ satisfy all the assumptions on v. Figure 2.4 provides

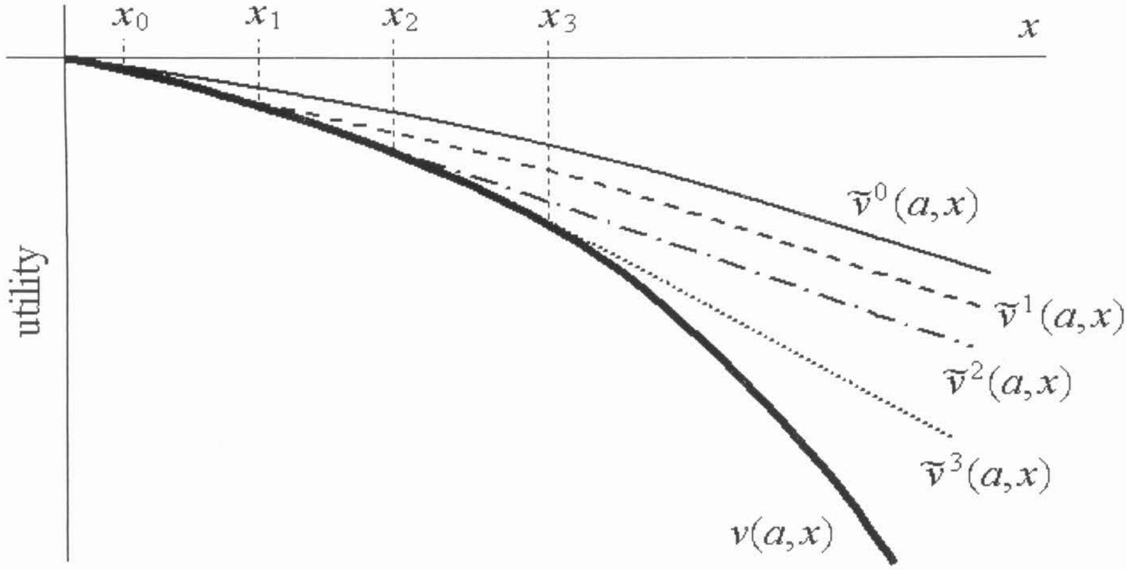

Figure 2.4: An example of gradual learning of the tolerance function

an example in which there is only one environment, and $a_1 = a_2 = a_3$.

Note that under this sort of learning, in each period, the agent behaves as if her perceived tolerance function at that time is the true tolerance function. At each time t, the agent believes that his lifetime utility maximization problem is:

$$\max_{\{a_t^A, a_t^B, c_t^A, c_t^B\}_{i=t}^{\infty}} E_t \sum_{i=t}^{\infty} \delta^i [\{\omega_i = A\} \left(u(c_i^A, a_i^A) + \widetilde{v}^{At}(a_i^A, x_i^A)\right) \qquad (2.14)$$
$$+ \quad \{\omega_i = B\} \left(u(c_i^B, a_i^B) + \widetilde{v}^{Bt}(a_i^B, x_i^B)\right)]$$

Therefore, the agent's problem at time t can be written as a stationary Bellman equation:

$$\begin{aligned}\widetilde{V}^t(x^A,x^B) &= \max_{a^A,a^B} \mu\left[u(y-pa^A,a^A)+\widetilde{v}^{At}(a^A,x^A)+\delta\widetilde{V}^t(\alpha x^A+\beta a^A,x^B)\right] \\ &+(1-\mu)\left[u(y-pa^B,a^B)+\widetilde{v}^{Bt}(a^B,x^B)+\delta\widetilde{V}^t(x^A,\alpha x^B+\beta a^B)\right]\end{aligned} \quad (2.15)$$

and the agent's behavior every period is qualitatively similar to that of the agent with the simple learning process used in the bulk of this paper. That is, the Theorem of the Maximum together with the assumptions on utility guarantee the existence and uniqueness of $\widetilde{V}^t(x^A,x^B)$. As before, because utility is separable with respect to environments, the value function can be written as the weighted sum of two separate, environment-specific value functions:

Proposition 4

$$\widetilde{V}^t(x^A,x^B)=\frac{\mu}{1-\delta(1-\mu)}\widetilde{W}^{At}(x^A|\widetilde{\delta}=\frac{\delta\mu}{1-\delta(1-\mu)})+\frac{1-\mu}{1-\delta\mu}\widetilde{W}^{Bt}(x^B|\widetilde{\delta}=\frac{\delta(1-\mu)}{1-\delta\mu})$$

where $\widetilde{W}^{jt}(x|\widetilde{\delta}=\xi)$ *is the solution to*

$$\widetilde{W}^{jt}(x)=\max_a\left[u(y-pa,a)+\widetilde{v}^{jt}(a,x)+\xi\widetilde{W}^{jt}(\alpha x+\beta a)\right] \quad (2.16)$$

for $j=A,B$.

Given the environment-specific value functions, $\widetilde{W}^{At}(x)$ and $\widetilde{W}^{Bt}(x)$, we can again characterize the optimal policy correspondence with the following Proposition:

Proposition 5 *The optimal stock evolutions associated with the Bellman Equations* *given by Equations (2.16), can each be characterized as follows: (i) every optimal path is a monotonic sequence; (ii) any optimal path converges to a steady state; and (iii) there exists exactly one critical level between any two consecutive stable steady states.*

In this setting, depending on the factors that determine the critical values, the agent may choose to quit (or relapse) in any period, not simply in the one period in which the threshold is reached, as in the simple learning case. Therefore, even though the agent's one-period behavior may be the same as in the simple learning case, her dynamic behavior may be substantially richer.

As a simple example in the single environment setting, consider Figure 2.5, in which the tolerance function takes the functional form $v(a,x) = -\gamma_x x + \gamma_{ax} ax$. Additionally, assume that the agent knows the functional form, but does not know γ_x or γ_{ax}. In particular, suppose the true tolerance function is $v(a,x) = -6x + 5.6ax$ but the agent's initial perceived tolerance function is $\widetilde{v}_0(a,x) = -5.5x + 5.5ax$. Suppose $x_0 = 0.85$. Given the agent's perceived tolerance function, his optimal consumption is $a_0 = 0.72$, which implies $x_1 = 1.14$. The agent observes the realized value of the actual tolerance function, and updates his perceived tolerance function accordingly. Suppose $\widetilde{v}_1(a,x) = -5.93x + 5.5ax$. Then his optimal consumption is $a_1 = 0.80$. At this point, after observing two realizations of the true tolerance process, the agent exactly identifies both parameters of the true tolerance function, and follows the optimal consumption path that results from the true problem.

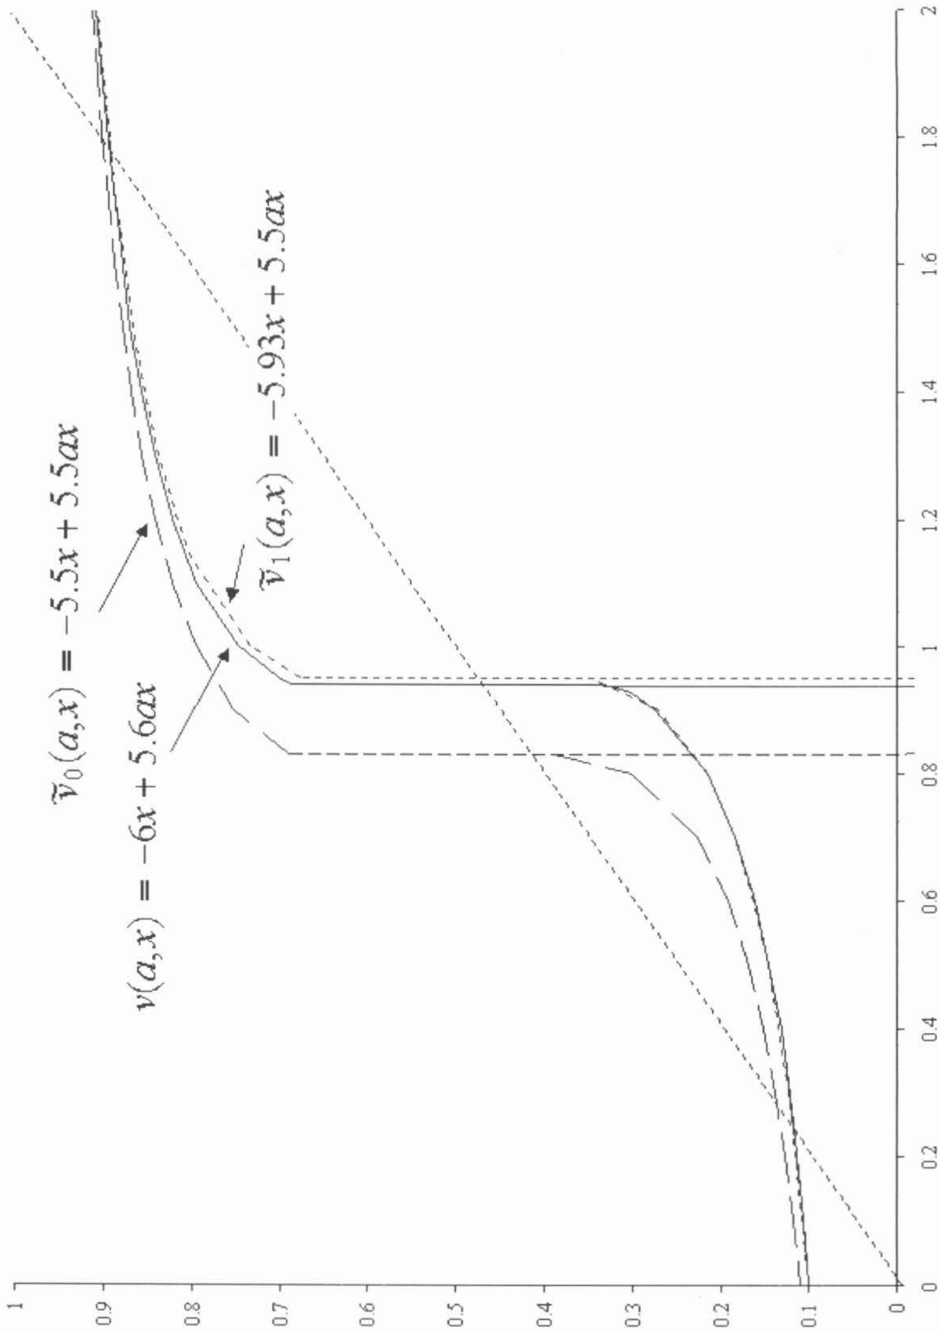

Figure 2.5: An example of gradual learning

2.5.2 A More Complete Model of Learning and Beliefs

As mentioned above, under this sort of learning, the agent behaves in each period as if her perceived tolerance function at that time is the true tolerance function. A richer model would include agents who realize that their perceived tolerance function is simply a belief about the true tolerance function and who know how their beliefs will be updated in the future.

Consider a more general model of learning, in which agents update their perceptions of the true tolerance function. Suppose that the agent knows the true functional form of the tolerance function, but does not know all the parameters of the true tolerance function. Let γ denote the vector of unknown parameters, and write the true tolerance function as $v\left(a, x, \gamma\right)$. As before, assume that the perceived tolerance function varies across environments. In environment j, the agent has some initial belief on γ, denoted by $\widetilde{\gamma}_0^j$, and updates $\widetilde{\gamma}_t^j$ according to some function g that is assumed to be twice continuously differentiable in both its arguments

$$\widetilde{\gamma}_{t+1}^j = g\left(\widetilde{\gamma}_t^j, \gamma^j\right)$$

The updated beliefs depend not only on last period's beliefs, but also on any knowledge having to do with the true parameter set. For example, at time t in environment j, the agent knows the value of $v(a_t^j, x_t^j, \gamma)$, and uses this information to update her beliefs.

At time t, if the environment is A, the agent's belief $\widetilde{\gamma}_t^B$ does not evolve: $\widetilde{\gamma}_{t+1}^B = \widetilde{\gamma}_t^B$.

Likewise, if the environment is B, the agent's belief $\widetilde{\gamma}_t^A$ does not evolve. The minimum number of periods in environment j for the agent to realize the true tolerance function is equal to the number of free parameters, assuming that he observes the value $v(a_t^j, x_t^j, \gamma^j)$ at time t. The time it takes for an agent to realize the true tolerance function is increasing in the degree of persistence of past beliefs. The perceived tolerance function is also written as a function of the perceived parameters: $v(a, x, \widetilde{\gamma})$.

As an example, assume as before that the tolerance function is linear in a and x. That is, suppose $v(a, x, \gamma) = \gamma_{ax} ax - \gamma_x x$. Now suppose that γ_{ax} is known to the agent, but γ_x is not. The agent has an initial belief on γ in environment A, denoted as $\widetilde{\gamma}_0^A$. Suppose that the agent's initial beliefs have some degree of persistence and every period in which he is in environment A, he updates his belief on γ according to this simple updating function:

$$\widetilde{\gamma}_{t+1}^A = \rho \widetilde{\gamma}_t^A + (1 - \rho)\,\gamma$$

In this case, if $\rho = 1$, the agent never updates his belief. He is constantly surprised by his realized utility. If $\rho = 0$, the agent updates his beliefs right away, and realizes the true tolerance function after the first period.

The functional equation for the general problem is given by

$$
\begin{aligned}
& V(x^A, x^B, \widetilde{\gamma}^A, \widetilde{\gamma}^B) = \\
& \quad \max_{a^A, a^B} \mu \left[u(y - pa^A, a^A) + v\left(a^A, x^A, \widetilde{\gamma}^A\right) + \delta V(\alpha x^A + \beta a^A, x^B, g\left(\widetilde{\gamma}^A, \gamma\right), \widetilde{\gamma}^B) \right] \\
& \quad + (1 - \mu) \left[u(y - pa^B, a^B) + v\left(a^B, x^B, \widetilde{\gamma}^B\right) + \delta V(x^A, \alpha x^B + \beta a^B, \widetilde{\gamma}^A, g\left(\widetilde{\gamma}^B, \gamma\right)) \right]
\end{aligned}
$$

As before, the separability of utility with respect to the addiction levels as well the beliefs about the true tolerance function allows the value function to be written as the sum of two environment-specific value functions.

Proposition 6

$$V(x^A, x^B, \widetilde{\gamma}^A, \widetilde{\gamma}^B) = \frac{\mu}{1-\delta(1-\mu)} W(x^A, \widetilde{\gamma}^A | \widetilde{\delta} = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu} W(x^B, \widetilde{\gamma}^B | \widetilde{\delta} = \frac{\delta(1-\mu)}{1-\delta\mu})$$

where $W(x, \widetilde{\gamma} | \widetilde{\delta} = \xi)$ *is the solution to*

$$W(x, \widetilde{\gamma}) = \max_a \left[u(y - pa, a) + v\left(a, x, \widetilde{\gamma}\right) + \xi W(\alpha x + \beta a, g\left(\widetilde{\gamma}, \gamma\right)) \right]$$

Once again, the assumptions on utility and the function $g(\cdot)$ and the Theorem of the Maximum ensure the existence, uniqueness, and differentiability of $W\left(x, \widetilde{\gamma}\right)$, as well as the existence of a non-empty upper semi-continuous policy correspondences:

$$\begin{aligned}
\lambda^A(x) &= \Big\{ x', \widetilde{\gamma}' | W(x, \widetilde{\gamma} | \frac{\delta\mu}{1-\delta(1-\mu)}) = u(y - \frac{p}{\beta}(x' - \alpha x), \frac{1}{\beta}(x' - \alpha x)) + v(\frac{1}{\beta}(x' - \alpha x), x, \widetilde{\gamma}) \\
&\quad + \frac{\mu}{1-\delta(1-\mu)} W(x', \widetilde{\gamma}' | \frac{\delta\mu}{1-\delta(1-\mu)}) \Big\} \\
\lambda^B(x) &= \Big\{ x', \widetilde{\gamma}' | W(x, \widetilde{\gamma} | \frac{\delta(1-\mu)}{1-\delta\mu}) = u(y - \frac{p}{\beta}(x' - \alpha x), \frac{1}{\beta}(x' - \alpha x)) + v(\frac{1}{\beta}(x' - \alpha x), x, \widetilde{\gamma}) \\
&\quad + \frac{1-\mu}{1-\delta\mu} W(x', \widetilde{\gamma}' | \frac{\delta(1-\mu)}{1-\delta\mu}) \Big\}
\end{aligned}$$

Unfortunately, without additional restrictions on the utility or tolerance functions, the optimal policy correspondences $\lambda^A(x)$, $\lambda^B(x)$ can not be further characterized.

2.6 Directions for Future Research

Clearly, there are many interesting extensions and applications of the model. A few are described below.

2.6.1 Spillover in Addiction Levels

In general, tolerance is environment-specific, but there may be some spillover of tolerance to other environments. For example, in the overdose studies of Siegel et al. (1982), experienced rats who were injected with heroin in an unfamiliar environment had a lower rate of overdose than rats who were completely inexperienced with heroin.

Consider the basic model, but with the following change to instantaneous utility. Utility depends on the addiction level associated with the present environment, as before. However, the addiction levels associated with the other non-present environments also have minor impact on the utility function.

Suppose that the state of the world at time t is $\omega_t = j$, $j \in \{A, B\}$. The addiction level associated with environment $k \neq j$ has a minor effect on utility. Before the agent realizes the true tolerance function:

$$\widetilde{U}(c_t^j, a_t^j, x_t^j, x_t^k) = u(c_t^j, a_t^j) + \widetilde{v}(a_t^j, x_t^j + \sigma x_t^k)$$

and after the agent realizes the true tolerance function:

$$U(c_t^j, a_t^j, x_t^j, x_t^k) = u(c_t^j, a_t^j) + v(a_t^j, x_t^j + \sigma x_t^k)$$

where $\sigma \in [0, 1]$.

For ease of exposition, I will work with only the actual problem to the individual. The results are the same for the problem that the individual mistakenly perceives that he must solve.

Again, begin by considering the first order conditions to the individual's problem. In environment A,

$$\begin{aligned} u_2(y - pa_t^A, a_t^A) + v_1(a_t^A, x_t^A + \sigma x_t^B) &= pu_1(y - pa_t^A, a_t^A) + \sum_{i=1}^{\infty} \mu \delta^i \alpha^{i-1} \beta v_2(a_{t+i}^A, x_{t+i}^A + \sigma x_{t+i}^B) \\ &+ \sum_{i=1}^{\infty} (1 - \mu)\, \delta^i \alpha^{i-1} \sigma \beta v_2(a_{t+i}^B, x_{t+i}^B + \sigma x_{t+i}^A) \qquad (2.17) \end{aligned}$$

and in environment B,

$$\begin{aligned} u_2(y - pa_t^B, a_t^B) + v_1(a_t^B, x_t^B + \sigma x_t^A) &= pu_1(y - pa_t^B, a_t^B) + \sum_{i=1}^{\infty} \mu \delta^i \alpha^{i-1} \beta v_2(a_{t+i}^B, x_{t+i}^B + \sigma x_{t+i}^A) \\ &+ \sum_{i=1}^{\infty} (1 - \mu) \delta^i \alpha^{i-1} \sigma \beta v_2(a_{t+i}^A, x_{t+i}^A + \sigma x_{t+i}^B) \qquad (2.18) \end{aligned}$$

The spillover in addiction levels has two effects that are not present in the case with no spillover. First, the addiction level associated with the non-present environment enters directly into the tolerance function. Therefore, as the amount or strength of the spillover (σ) increases, marginal utility from consumption of the addictive good is weakly increasing. Second, the agent realizes that current consumption will have an effect not only on the periods in which the environment is the same as the current environment, but also on the periods in which the environment is different. Hence, the "full cost" of consuming the addictive good is also increasing in σ.

Under this framework, the Bellman equation to the individual's problem is

$$S(x^A, x^B) =$$

$$\max_{a^A, a^B} \mu \left[u(y - pa^A, a^A) + v\left(a^A, x^A + \sigma x^B\right) + \delta S(\alpha x^A + \beta a^A, x^B)\right] \tag{2.19}$$

$$+ (1-\mu)\left[u(y - pa^B, a^B) + v\left(a^B, x^B + \sigma x^A\right) + \delta S(x^A, \alpha x^B + \beta a^B)\right]$$

When $\sigma = 0$, the actual problem that the individual must solve reduces to that given by Equation (2.2). First order conditions (2.17) and (2.18) reduce to (2.10) and (2.11) and Equation (2.19) reduces to Bellman Equation (2.13).

At the other extreme, when $\sigma = 1$, the problem is essentially the same problem as that given in the SE Model with Bellman equation given by (2.7).[8] In this case, optimal consumption is independent of the environment.

Without the separability of utility with respect to environments, the value function can not be written as the sum of two separate value functions as in Proposition 2. However, Equation (2.19) can be re-written as

$$S(x^A, x^B) = \mu\left[f(x^A, x^B)\right] + (1-\mu)\left[g(x^A, x^B)\right] \tag{2.20}$$

[8]Let $x = x^1 + x^0$. Then, because preferences are the same, regardless of environment, the only difference between this problem and that of the SE Model is in the updating of the consumption stock.

where

$$f(x^A, x^B) = \max_{a^A} \left[u(y - pa^A, a^A) + v\left(a^A, x^A + \sigma x^B\right) \right.$$
$$\left. + \delta \mu f(\alpha x^A + \beta a^A, x^B) + \delta (1 - \mu) g(\alpha x^A + \beta a^A, x^B) \right]$$

and

$$g(x^A, x^B) = \max_{a^B} \left[u(y - pa^B, a^B) + v\left(a^B, x^B + \sigma x^A\right) \right.$$
$$\left. + \delta \mu f(x^A, \alpha x^B + \beta a^B) + \delta (1 - \mu) g(x^A, \alpha x^B + \beta a^B) \right]$$

The problem to the individual when the environment is A is given by $f(x^A, x^B)$, and the problem to the individual when the environment is B is given by $g(x^A, x^B)$.

Again, the agent's consumption problem can be viewed as a problem of two separate personalities. Unlike the case with no spillover, however, the dual personalities in this case are intertwined. As σ increases, the separation between the two personalities decreases.

More work is needed in this area in order to characterize optimal behavior of an agent for whom there is spillover of addiction levels.

2.6.2 Endogenous Environments

Consider the case in which the probability of one of the environments is increasing in the addiction level associated with that environment. For example, consider an agent who can consume alcohol at home or at the neighborhood bar. Initially, he consumes more alcohol at the bar than at home. As he consumes more alcohol, his visits to the bar become more frequent. In this case, the probability of environment A, μ, is an increasing function of the addiction level associated with that environment: $\mu'(x^A) \geq 0$. Although the full cost of consuming the addictive good in environment A is greater when $\mu'(x^A) \geq 0$ than when $\mu'(x^A) = 0$, if the agent chooses to consume the good in environment A, the endogenous determination of environment may "speed up" the addiction process. As the agent's consumption increases, not only does the strength of future craving increase when he is in environment A, but also the probability that he will be in environment A also increases.

2.6.3 Peer Effects

The model with cues can serve as a framework to study peer effects on consumption of addictive goods. For example, among adolescents, the peer group may be a strong determinant of the decision to begin consuming an addictive good:

> The one feature that is consistent in every clinical case is the presence of peer-drug associations. The young person's pattern of drug use is matched, almost point by point by shared drug use with his or her "gang," best friend, and/or boyfriend/girlfriend. (Oetting and Beauvais 1988, p. 156).

The model developed in this paper presents a way to generate "neighborhoods" or "pockets" of addictive substance users without assuming an explicit preference for peer or social acceptance. For example, suppose one environment is the presence of a friend or group of friends. Then using the interaction among this group, it is possible to solve for conditions under which the entire group chooses to consume the addictive good when they are together. Alternatively, there are conditions under which none of the group chooses to consume the addictive good.

2.6.4 Advertising

Another application of this model is the study of how firms might be able to manipulate environments in order to affect demand. Typical advertisements convey very little information about such things as price and quality. One explanation is that firms are trying to induce craving for goods by re-creating environments that are associated with past consumption of these goods. For example, advertisements for food, drink, or alcohol may try to induce craving by simply showing the product or showing others enjoying their product. These phenomena may be captured by a generalization of the present model to one of goods that display some degree of habit formation.

2.6.5 Policy Implications

The model with environmental cues demonstrates the problem that cues produced by other agents may impose negative externalities on the addict or recovering addict. Recall that if preferences are such that an agent would choose to quit in one environment but not the other, she "regrets" her addiction in that environment. Her utility is higher in time periods in which she is in the environment in which she does not feel addicted than when she is in the environment in which she is addicted. Cues produced by others that place her in the environment in which she is addicted is therefore a negative externality for her.

From a policy perspective, it may be possible to increase welfare to agents who are either addicts or recovering addicts by limiting their exposure to environmental cues associated with consumption of addictive goods. For example, smoking bans in public places may serve not only to provide clean air to nonsmokers, but also to reduce the temptation of those who are trying to quit smoking.

Lastly, the model also implies that drug education can also be welfare-enhancing, because drug education can inform agents about their true tolerance function.

2.7 Conclusion

The relapse rate among abusers of addictive substances is strikingly high. This phenomenon has proven to be an intricate problem for addiction researchers from all fields. Previous economic research on addiction, however, has not fully utilized findings from other disciplines. Studies on how people misjudge the severity of future consequences of addictive substance use explain why people might begin to use an addictive substance that they eventually choose to quit. Research on conditioned responses offers an explanation for why addicts who decide to stop using addictive substances begin to consume again, even if, at that point, they no longer have misperceptions about the negative effects of consumption.

In this paper, I have shown that agents who misperceive the future consequences of using addictive substances may fall into this pattern of quitting and relapse. This occurs because preferences and environmental cues become intertwined in such a way that craving can be induced simply by the presence of environmental cues.

Chapter 3 An Empirical Test of Rational Addiction: Consumer Response to Price and Policy Changes

3.1 Introduction

Rational choice models of addiction describe perfectly rational, forward looking agents who may develop a harmful addiction. The Becker and Murphy (1998) framework is an infinite horizon continuous time problem where utility is a function of current consumption of addictive and non-addictive goods, as well as a stock of past consumption of the addictive good. The key to this model lies in the relaxation of the usual assumption of intertemporal separability. Consumption patterns consistent with addiction can result from forward-looking utility maximization with stable preferences; that is, the functional form of utility is invariant over time. In this model, agents fully anticipate the effects of their current consumption of a (harmfully) addictive good on future utility. The model predicts that current consumption of an addictive good is increasing in past consumption because past consumption increases current marginal utility of consumption. Furthermore, if the agent is rational and forward looking, then current consumption should also be increasing in expected future consumption.

In the previous chapter, I point out that the rational model omits two prevalent features of addiction–misperception by agents of the future consequences of consuming the addictive good and the pattern of quitting and relapse that is frequently seen in addicts. Drawing on evidence from addiction research in other fields, I develop a behavioral model of addiction that captures these two features of addiction. Related economic research includes the work of Laibson (1999), O'Donoghue and Rabin (1999), and Orphanides and Zervos (1995).

Thus far, Becker and Murphy's rational model has been the standard model of addiction in economics. There have been a few empirical tests of the rational addiction model that pertain to a variety of addictive substances and activities, such as cocaine (Grossman and Chaloupka, 1998), alcohol (Grossman, Chaloupka, and Sirtalan, 1998), and casino gambling (Nichols, 1999). A few papers use data on cigarette consumption to conduct empirical tests of the rational addiction model. Examples include Becker, Grossman and Murphy (1994), Chaloupka (1991) Keeler, Hu, Barnett and Manning (1993), and Gruber and Kőszegi (1999).

Becker, Grossman, and Murphy (BGM) use state level annual cigarette tax receipts from 1955 through 1985 to measure per capita consumption of cigarettes. They find that cigarettes are addictive. That is, current consumption is increasing in past consumption. Furthermore, addicts are forward looking, as opposed to myopic, in that current consumption is found to be increasing in future consumption. Lastly, they find that long-run responses to permanent price changes are almost twice as large as short run responses.

A serious problem with using cigarette and tobacco tax receipts to measure consumption is that, for most states, state-level tobacco taxes are paid by tobacco distributors, rather than tobacco consumers. For example, according to the California State Board of Equalization, the department that oversees cigarette and tobacco tax collections: "The [cigarette] tax and [cigarette and tobacco products] surtax are paid by distributors, who purchase tax stamps from banks and affix them to each package of cigarettes before distribution. Distributors can be reimbursed for these taxes by the businesses to whom they sell the cigarettes, and the businesses include the taxes as part of the retail selling price of the cigarettes" (Cigarette and Tobacco Products Tax Law, 1998). Furthermore, cigarette distributors do not necessarily hold the same number of packs of cigarettes and number of stamps in stock. Therefore, state-paid tobacco taxes more accurately reflect distributors' demand for cigarette and tobacco tax stamps, rather than consumer demand for cigarettes. Using state tax receipt data can therefore lead to mistaken inferences, as I discuss below.

Chaloupka (1991) tests the rational model using data from the second National Health and Nutrition Examination Survey, which includes 28,000 respondents and covers the time period 1976-1980. Unlike the aggregate data that BGM use, these data describe the consumption of individuals. Like BGM, Chaloupka finds that cigarettes are addictive and that individuals are not myopic. That is, he finds that both past and future consumption have positive effects on current consumption, although he finds long-run price elasticities that are about half those of BGM. Chaloupka also estimates separate cigarette demand equations for subsamples based on educational attainment and age. He finds that for less educated and younger individuals, the coefficient on future consumption is not significantly different from zero. Individuals in these groups behave more myopically than do more educated or older individuals.

To test the rational model, at least three consecutive periods of data are required. However, the survey data that Chaloupka uses include only two consecutive periods. Therefore, for future consumption, he uses what respondents report as current consumption. For current consumption, he uses reported one year lagged consumption. Lastly, for past consumption, he uses reported maximum average daily quantity for those who began smoking more than two years ago; otherwise past consumption is recorded as zero. In all likelihood, this measurement error is not independently and identically distributed across respondents. Therefore, the resulting estimates may be biased in an unanticipated manner.

Keeler, et al. (1993) include an analysis of the rational model of addiction in their study of taxation and regulation. They use cigarette tax receipts for the state of California as the measure of consumption. Unlike the previous studies, they acknowledge and attempt to correct for serial correlation in the rational model. As in the previous studies, they find a positive and significant coefficient on future consumption. However, they find a negative coefficient on lagged consumption, which is difficult to reconcile with the rational model. Their finding hints that the BGM implementation of the rational model is not very robust to different econometric specifications, as Gruber and Kőszegi (1999) explicitly discuss in their paper.

Gruber and Kőszegi have a thorough critique of the methods used by BGM and Chaloupka. In particular, they find that the BGM results are extremely sensitive to different specifications of the model. They replicate the BGM analysis using a similar dataset of state level tax receipt data, and they find similar results. However, when Gruber and Kőszegi attempt to control for any state-specific fixed effects, they find that the coefficient on future consumption is no longer significant.

To provide a better test of forward-looking behavior, Gruber and Kőszegi use the tax receipt data, as well as Vital Statistics Natality Data, to study the effect on consumption of announced tax increases that are not yet effective. This latter dataset describes the smoking behavior of expectant mothers. As predicted by the rational model, they find that consumption decreases during the period between enactment and implementation of tax hikes.

The main problem with the dataset that Gruber and Kőszegi use is that it is not representative of the population as a whole. Expectant mothers who smoke likely have lower discount rates than the average consumer. Furthermore, any concern that survey respondents may deny or downplay their consumption of such goods as cigarettes, alcohol, or illegal drugs due to social conformity should be heightened when the survey respondents are pregnant women.

The dataset that I use avoids the previously mentioned data problems. The data, compiled by Information Resources Incorporated from grocery store scanner data, describe weekly sales in 20 markets that span the states of California, Arizona, Colorado, Nevada, and Washington.

The previous tests of the rational addiction model have focused on the model's predictions concerning consumption responses to future and past price changes. This essay also allows an empirical test of the prediction that the response to an anticipated price change differs from the response to an unanticipated price change.

In particular, I consider the consumption effects of three institutional changes that occur during the time period 1996 through 1999. The first is the ban on smoking in bars and taverns in California as part of the state's comprehensive "Smoke-Free Workplace" law. I argue that this is an anticipated permanent shock to future consumption.

Secondly, as a result of the settlement that the five largest tobacco companies signed with 46 states in November 1998, these companies raised wholesale tobacco prices by 45 cents per pack, the largest cigarette price increase in history. Although analysts, and perhaps smokers, may have predicted a price increase contingent on settlement of the litigation, it was not clear when a settlement would occur and what the terms of the settlement would be. Tobacco companies announced the price increase the same day that they signed the settlement.

Lastly, in the November 1998 election, California voters approved a 50 cent tax increase on cigarettes. This tax increase was anticipated. The increase was not effective until January 1, 1999, but the official outcome of the election was announced in mid-November. The price increase due to the tobacco settlement occurred during the period between the approval of this tax and its implementation. Fortunately, for purposes of econometric identification, the tax increase applied to California only, whereas the tobacco settlement price increase affected the whole country. Therefore, I can study these two events separately.

The next section gives additional details on these three policy changes. Section 3 lays out the theoretical model of Becker and Murphy and discusses some of the econometric issues that arise in estimation of the model. Section 4 provides a description of the data.

Section 5 presents the empirical results. I begin by attempting to replicate the BGM results using the grocery store scanner data. Like BGM, I find that the coefficients on lagged and lead consumption are positive and significant. An attempt to correct for fixed trends in the panel data, however, reveals the sensitivity of the results to the econometric specification. While the results are not inconsistent with the rational model, they are difficult to interpret within the context of the rational framework.

As an additional modification, I assume that consumers forecast prices using lagged prices and other available information, rather than assume that they can perfectly predict future prices and consumption. While the results of this analysis are not very different from the previous results, the price forecast results raise questions about the validity of the methods used in the majority of the empirical studies of rational addiction.

Next, I consider the effects of the policy changes described in Section 2. I find that the ban on smoking in bars and restaurants has no effect on consumption. Again, the analysis of the smoking ban raises some issues with the econometric specification of the rational model. Lastly, I compare the effects of the unanticipated price increase due to the tobacco settlement and the anticipated Proposition 10 tax increase. Contrary to the rational model, I find that the consumption response to both price changes are similar.

Section 6 discusses some methodological issues, measurement error and serial correlation, that the results raise and Section 7 concludes.

3.2 Institutional Background

3.2.1 California Smoke-Free Workplace Law

This paper considers the effects of California's "Smoke-Free Workplace" Law. Assembly Bill 13, introduced to the Assembly in December, 1992, and chaptered in July, 1994, added to the Labor Code a section prohibiting smoking of tobacco in any enclosed spaces of a workplace. All workplaces, with some exceptions which include bars and taverns, were to comply immediately. [1] Bars and taverns were to begin compliance on January 1, 1997. However, assembly Bill 3037, chaptered in September, 1996, amended the Labor Code section to extend from January 1, 1997 to January 1, 1998 the date of compliance by bars and taverns.

The California ban on smoking in bars and taverns provides a natural test for the predictions of the rational model. This ban acts as a permanent shock to future consumption for smokers. Consider the rational, forward-looking addict for whom this shock is prohibitive enough that he would quit smoking once the ban is in effect.

[1] There are a few workplaces which are exempted from the smoke-free workplace laws such as private residences; employers with a total of five or fewer employees, with some additional conditions; or retail or wholesale tobacco shops and private smokers' lounges. Note that "retail or wholesale tobacco shop" is defined as "any business establishment the main purpose of which is the sale of tobacco products," (California Labor Code Section 6404.5) and therefore does not include bars or taverns which also sell tobacco products.

This individual should begin to decrease consumption after the passage of the law but prior to the date of compliance. The smoking ban is not a monetary price increase; therefore, there is no need to worry about individuals who might stockpile cigarettes before the implementation date, as could be a problem in an analysis of an anticipated tax or price increase.

3.2.2 California Proposition 10 Tax Increase

In the November 1998 election, California voters passed Proposition 10, a measure that would increase tobacco taxes by 50 cents on January 1, 1999, to finance early childhood development programs. Proposition 10 passed with 50.4% of the vote, making it one of the narrowest victories in California referenda history.

I argue that this tax increase to smokers is an anticipated price change. The final outcome for this proposition was announced on November 12, 1998, yet the tax hike would not go into effect until January 1, 1999. Furthermore, this proposition received extensive press coverage not only because the outcome was so close, but also because Proposition 10 was drafted by Hollywood actor, director and producer Rob Reiner, who also helped finance the campaign. There was also dramatic spending on the "No on 10" campaign: Tobacco companies alone spent over $30 million.

3.2.3 Tobacco Litigation Settlement Price Increase

On November 23, 1998, the five largest tobacco companies in the Unites States settled a lawsuit filed by 46 states, the District of Columbia, and five U.S. territories.

The settlement called for tobacco companies to pay \$206 billion dollars to reimburse states for providing health care to smokers. That same day, the three largest tobacco companies announced a 45 cent per pack price increase (Los Angeles Times, 1998).

When compared to the Proposition 10 tax increase to be implemented on January 1, 1999, this price increase seems much less likely to have been anticipated by consumers. Although analysts may have predicted a price increase if the tobacco companies and states were to settle, they did not predict such a large price increase, in part because this was the largest one time price increase in history.

In addition, the settlement occurred quite rapidly. The terms of the settlement were not formulated until November 14, nine days before the price hike, and the states signed the settlement November 20. Furthermore, some stores raised prices right away, while other stores waited until distributors passed the price increase on to them before raising prices (Howe, 1998).

3.3 Theoretical Model

In the rational model of Becker and Murphy (1988), period t utility depends on the current consumption of the addictive consumption good, a_t, current consumption of the non-addictive consumption good, c_t, and past consumption of the addictive good as summarized by x_t. In their empirical test of the rational model, Becker, Grossman, and Murphy (1994) add an error term, e_t, which is also referred to as the "impact of unmeasured life-cycle variables on utility" (p. 398). Instantaneous utility is given by $u(a_t, c_t, x_t, e_t)$, where the utility function $u(\cdot)$ is concave in all its arguments.

Let the price of the addictive good at time t be P_t, with c_t as numeraire. If the agent's discount rate is δ, the agent's maximization problem under perfect foresight is:

$$\max \sum_{t=1}^{\infty} \delta^{t-1} u(a_t, c_t, x_t, e_t)$$

subject to

$$\sum_{t=1}^{\infty} \delta^{t-1}(c_t + P_t a_t) = A^0$$

where x_0, initial consumption stock, is given exogenously and A^0, the present value of lifetime wealth, is assumed constant and exogenous.

Let λ be the Lagrange multiplier on the budget constraint. Then the first order conditions associated with the maximization problem are:

$$\begin{aligned} u_2(a_t, c_t, x_t, e_t) &= \lambda \\ u_1(a_t, c_t, x_t, e_t) + \sum_{j=1}^{\infty} \delta^j \alpha^{j-1} \beta u_3(a_{t+j}, c_{t+j}, x_{t+j}, e_{t+j}) &= \lambda P_t \end{aligned}$$

Becker, Grossman, and Murphy assume a quadratic utility function. Furthermore, they assume that past consumption can be summarized by last period's consumption. That is, $x_t = a_{t-1}$. Under these assumptions, we get the following demand function for the addictive good:

$$a_t = \theta_1 + \delta\theta_2 a_{t+1} + \theta_2 a_{t-1} + \theta_3 P_t + \theta_4 e_t + \delta\theta_5 e_{t+1} \tag{3.1}$$

If the agent can perfectly forecast her future consumption, as is implicitly assumed by BGM, then Equation (3.1) gives the agent's demand function. However, if the agent can not perfectly predict her future consumption, then the demand function can be rewritten as:

$$a_t = \theta_1 + \delta\theta_2 E_t\left[a_{t+1}|\Phi_t\right] + \theta_2 a_{t-1} + \theta_3 P_t + \theta_4 e_t + \delta\theta_5 E\left[e_{t+1}|\Phi_t\right] \tag{3.2}$$

where Φ_t is the information available to the agent at time t[2]. Such information can include the history of past prices, history of past consumption, and any price or tax change announcements.

According to this model, past consumption of the addictive good increases current consumption if $\theta_2 > 0$. That is, the addictive good is truly addictive if θ_2 is positive and significant. If the good is addictive and the agent is forward looking, current consumption will depend on future consumption; that is $\delta\theta_2 > 0$.

A serious problem for estimating either Equation (3.1) or Equation (3.2) is the endogeneity of lagged consumption and lead consumption (actual or expected). Current consumption (a_t) is a function of lagged consumption (a_{t-1}) and lead consumption (a_{t+1}), while lagged consumption (a_{t-1}) is a function of a_t and a_{t-2}. Likewise, a_{t+1} is a function of a_t and a_{t+1}. Therefore, these two right-hand side variables, a_{t-1} and a_{t+1}, are likely highly correlated with the error, e_t. Least squares estimation would lead to inconsistent estimates.

Previous empirical studies have attempted to find instruments for lead and lagged consumption. The most common instruments are lags and leads of prices. Lagged

[2]This simple relationship between Equations (3.1) and (3.2) arises because of the assumption of quadratic utility.

prices are correlated with lagged consumption, and lead prices are correlated with lead consumption. Equation (3.1) implies that, a_{t-1} and a_{t+1} fixed, current consumption a_t is independent of lagged and lead prices. Unfortunately, the validity of lagged and lead prices as instruments for lagged and lead consumption relies on strong assumptions on model specification and measurement errors.

Lastly, note that current period prices are assumed to be exogenous. It may be that this assumption is not too problematic during many time periods studied in which supply is very elastic at the market level. The 45 cent price increase due to the tobacco settlement is clearly an exogenous price shock during this sample period. However, evidence that retail prices rise in advance of an anticipated tax increase would indicate otherwise.

It could be argued that the 50 cent tax increase in California is an endogenous price change. That is, decreased demand in California would make it easier to pass a tax increase. However, as seen in Section 5 below, there is no apparent demand shift in California during the sample period. In fact, even though there is a downward trend in consumption, consumption in California is decreasing slower than consumption in the other states over the sample period.

Overall, prices do not appear to be driven by demand shocks. There is little weekly variation in prices until the sharp price increase in November, 1998, even though weekly consumption fluctuates. As discussed in section 5 below, prices over the sample period have a slight upward trend while there is a downward trend in consumption. However, the trend in prices is fairly constant across states, even though the trend in consumption varies greatly.

3.4 Data Description

Information Resources, Inc., Sales Data. The consumption and price data employed in this analysis are scanner data obtained from Information Resources Incorporated (IRI), a company that specializes in collecting scanner data from various grocery stores, drugstores, and convenience stores.

The unit of observation is the Designated Marketing Area (DMA) defined by Nielsen Media Research. DMAs are mutually exclusive and are defined as "all counties whose largest viewing share is given to [television] stations of that same market area" (Nielsen Media Research website http://www.nielsenmedia.com). DMAs cover all of the contiguous states, Hawaii and parts of Alaska.

The IRI data include weekly sales from January 1, 1996, through May 9, 1999, in 20 different markets that cover California, Arizona, Colorado, Nevada and Washington.[3] Eleven of these DMAs are California markets. IRI provides data on total revenues and total units sold in each DMA in each week. The average price of a pack of cigarettes within a DMA can then be calculated. Prices include state and federal cigarette and tobacco taxes, but do not include any additional state or local retail sales tax. As mentioned earlier, state cigarette and tobacco taxes are paid by distributors, and federal cigarette and tobacco taxes are paid by the cigarette manufacturer.

[3]There are actually 21 markets that cover the states of California, Arizona, Colorado, Nevada and Washington. One of the markets, the Grand Junction-Montrose market in Colorado, exhibits a sharp decrease in sales that starts in the beginning of August 1998. Because it is not clear whether this decline is due to a demand shift or a problem with the data collection, I have eliminated this market from the analysis.

Thus, the calculated price is the average price paid by consumers within a DMA net of (i.e., before) sales tax.

Yearly DMA population estimates through 1998 were obtained from the Polk Company, a company that specializes in collecting DMA-level data for marketing purposes. DMA populations for 1999 are estimated using linear projections based on population data from 1995-1998. In order to calculate per capita consumption of cigarettes purchased from the IRI grocery stores, I use a 52 week moving average of yearly population.

Board of Equalization Tax Receipt Data. In addition, I have collected California tobacco tax receipts from the California State Board of Equalization. These data allow for comparison between IRI sales data and state tax receipts.

In order to calculate the BOE per capita consumption variable, I use county population estimates from the Bureau of Economic Analysis.

Finally, per capita income data by county from 1996-1997 were obtained from the Bureau of Economic Analysis. Income for 1998-1999 are estimated using simple linear projections based on income data from 1983-1997. Per capita income used in the regressions are 52 week moving averages of yearly income.

3.5 Empirical Analysis

3.5.1 Descriptive Analysis

The descriptive statistics of the variables used in the analysis are summarized in Table 3.1. Per capita cigarette pack consumption in week t is denoted a_t. The average retail price of a cigarette pack in week t is denoted P_t.

Table 3.1: Descriptive statistics

variable	source	description	mean	std dev
a_t	IRI and Polk	weekly per capita consumption (packs)	0.261	0.192
a_t	BOE and BEA	CA monthly per capita consumption (packs)	4.311	0.679
P_t	IRI	price (January 1996 dollars)	1.985	0.441
Y_t	BEA	annual income (1000's of 1996 dollars)	20.333	3.290
Correlation between weekly per capita consumption and real price = -0.404				
Correlation between CA monthly per capita consumption and real price (in CA) = -0.573				

Figures 3.1-3.2 and Table 3.2 further describe the weekly IRI sales data. 3.1 plots the time series of per capita units of cigarette sales, where the data have been aggregated to the state level. The time series of average prices by state is plotted in Figure 3.2. It is clear that there is a general downward trend in consumption and an upward trend in prices.

Week 105 is the first week of 1998, when the ban on smoking in bars began. Examination of the California series indicates that the ban has no obvious effect on cigarette prices and per capita consumption. If anything, the downward trend in consumption appears to slow down after the implementation of the ban, a finding

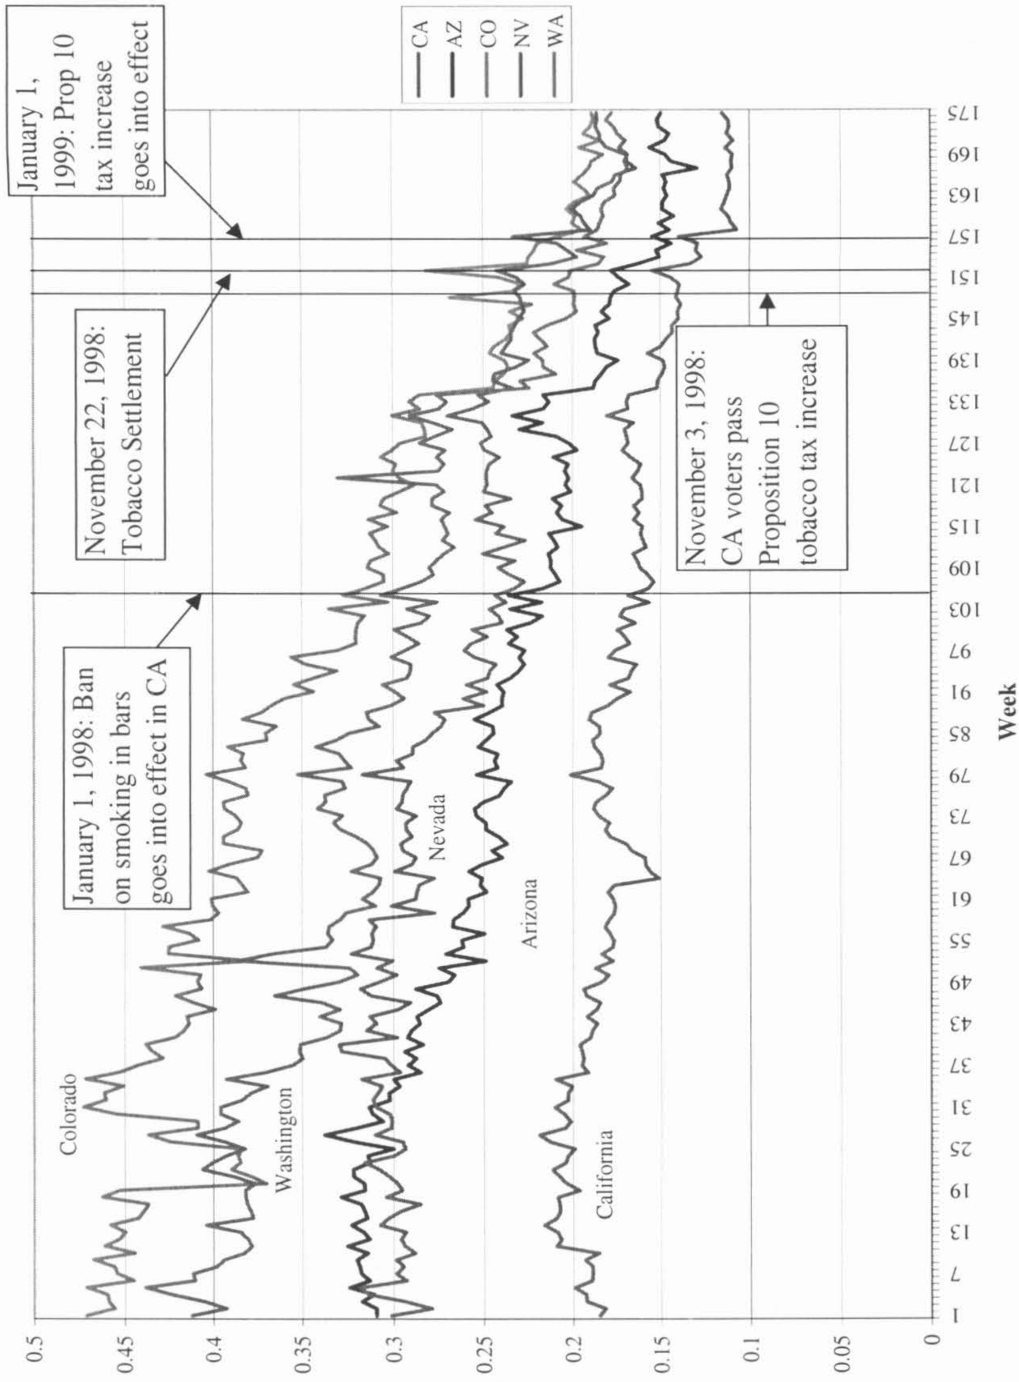

Figure 3.1: Per capita cigarette pack sales

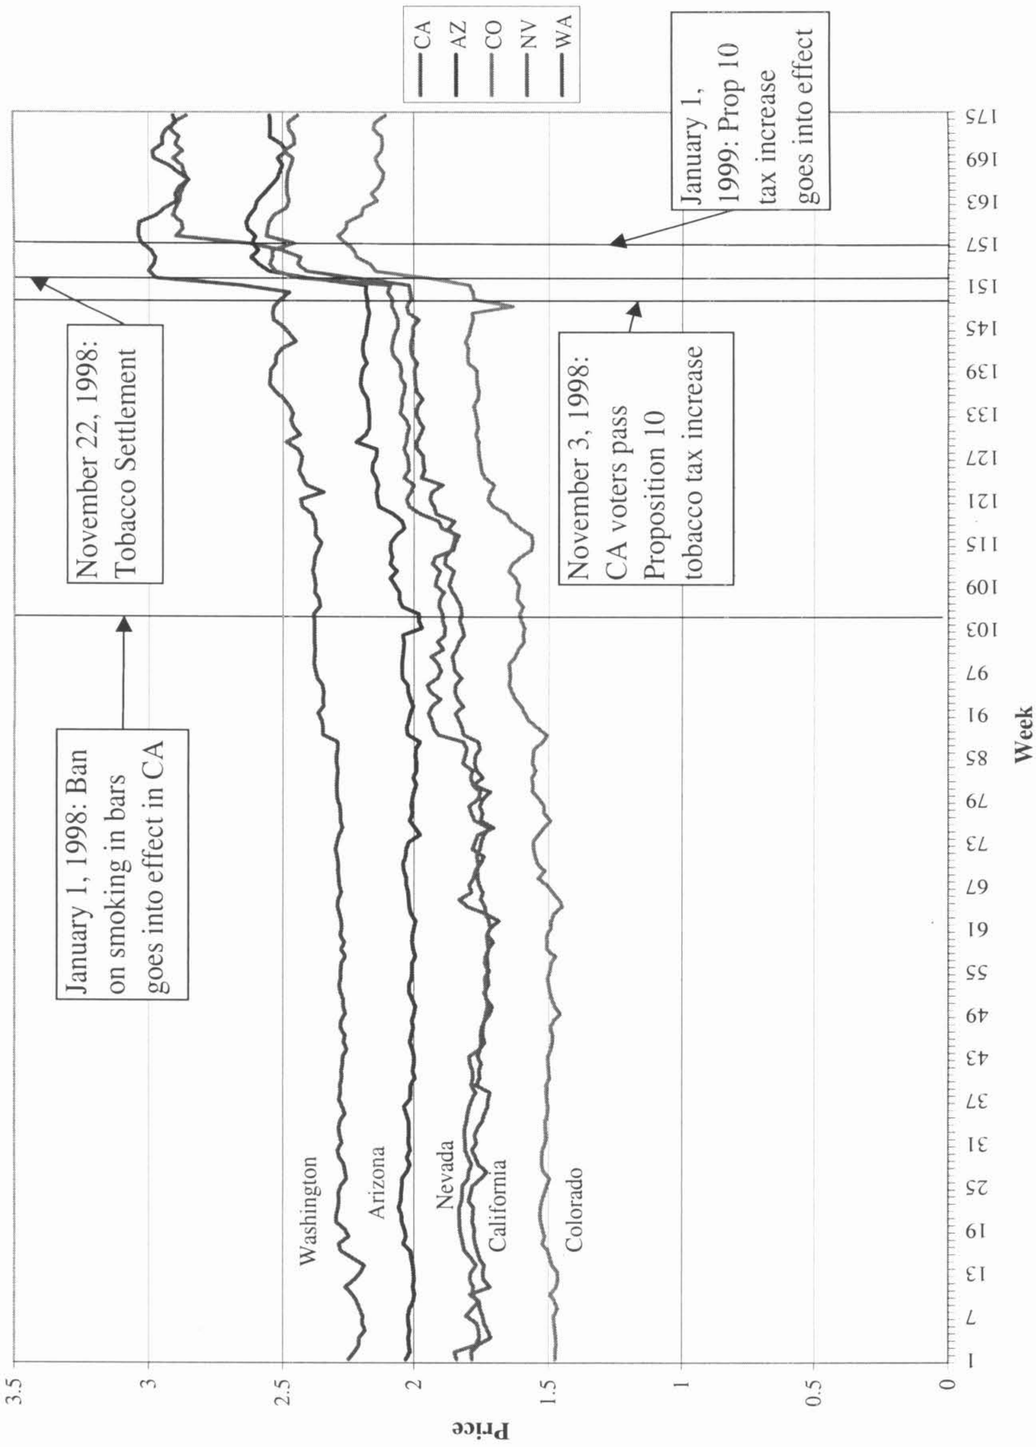

Figure 3.2: Average price per pack of cigarettes

that is consistent with forward-looking agents decreasing consumption in advance of the implementation. However, the lack of data on consumption prior to the enactment of the legislation makes inference difficult. The tobacco settlement was signed by the five largest tobacco companies on November 23, 1998, in Week 152. The 45 cent per pack price increase was announced the same day. Average prices in some states increased immediately (Nevada for example), while in other states, prices increased more slowly (Washington and Colorado, for example). Aside from the general time trend, there does not appear to be any significant effect on consumption. Week 157 is the first week of 1999, when the 50 cent Proposition 10 tax increase went into effect. Average prices started to increase prior to the new year, perhaps because grocery stores anticipated that consumers might stockpile cigarettes before the new tax increase. Indeed, there does appear to be a temporary increase in consumption (or sales) prior to week 157.

Table 3.2 presents average weekly per capita consumption by state over six-month intervals. Again,the downward trend in consumption is apparent.

Table 3.2: Weekly per capita consumption by state and six month interval

Time period	CA	AZ	CO	NV	WA	AZ,CO,NV,WA
01.01.96-06.30.96	0.200	0.317	0.437	0.299	0.394	0.362
07.01.96-12.31.96	0.195	0.291	0.429	0.308	0.360	0.347
01.01.97-06.30.97	0.176	0.252	0.395	0.296	0.326	0.317
07.01.97-12.31.97	0.177	0.238	0.353	0.264	0.305	0.290
01.01.98-06.30.98	0.163	0.209	0.303	0.243	0.279	0.259
07.01.98-12.31.98	0.147	0.183	0.246	0.233	0.217	0.220
01.01.99-05.09.99	0.113	0.147	0.194	0.186	0.177	0.176

IRI sales data in California and the Board of Equalization tax receipt data are compared in Figure 3.3. The IRI data are aggregated into monthly sales for easier comparison with the BOE data. The IRI and BOE data follow fairly closely, with the IRI sales about one-tenth of the BOE tax receipts, until November, 1998.[4] Once the weekly sales data are aggregated into monthly data, there appears to be no stockpiling of cigarettes before the January 1, 1999, tax increase. However, it appears as if cigarette distributors are stockpiling cigarettes (or tax stamps) from November through December 1998. In January 1999, tax receipts drop sharply.

3.5.2 Replication of Becker Grossman Murphy (1994)

Table 3.3 presents the results from attempts to replicate the BGM analysis using various data sources. Recall that BGM's analysis focuses on estimation of Equation (3.1). As a basis of comparison, one set of BGM two stage least squares estimates (2SLS) is reported in Column (i) of the top panel (corresponding to BGM's Table 3, Column (i) on page 406). To address the problem of endogeneity of lagged and lead consumption, they use lagged and lead prices as instruments. Additional regressors are full sets of dummy variables for state and year. The reported estimates correspond to the identified parameters of Equation (3.1).

There are some important distinctions between the data set that they employ and

[4]This proportion suggests that 90% of the cigarettes sold in California are sold outside of the grocery stores in the IRI data. Maintained assumptions in this analysis are that variation in grocery store purchases with prices (and other variables) is representative of the smoking population and that this proportion is constant across states and over time. However, measured levels of per capita consumption are affected by this low proportion because the population used to calculate per capita consumption is population of the entire market, not the population that shops at those grocery stores. Furthermore, as I discuss later, the data may not capture casual smokers.

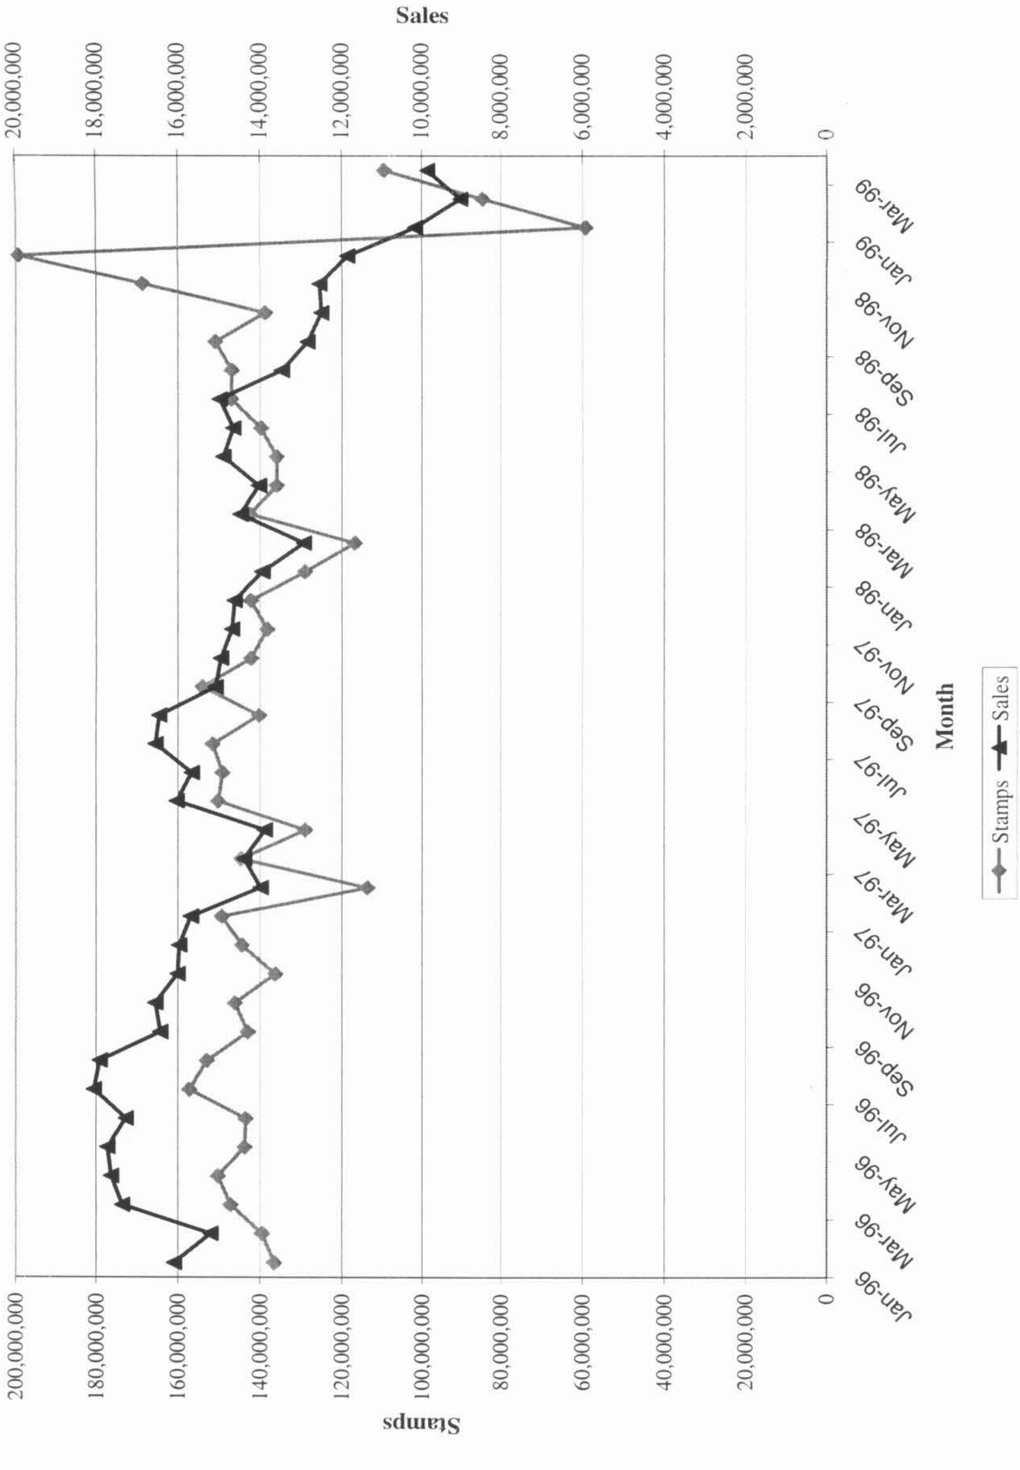

Figure 3.3: Monthly California Cigarette Tax Stamps vs. Monthly California Sales

the data sets that I use. The first is that BGM use *annual* per capita consumption as implied by the state tax receipts. I use *weekly* per capita cigarette consumption for the regressions using the IRI data, and *monthly* per capita consumption for the regression that uses the BOE tax data. Prices in my data set are expressed as January 1996 dollars, and weekly or monthly (for IRI and BOE, respectively) income is expressed as thousands of January 1996 dollars. BGM use hundreds of 1969 dollars for annual income and 1969 cents for prices. I use full sets of dummy variables for month, not year, for the IRI data analyses. For the BOE analysis, I use a linear time trend instead, because there are only 35 observations. Lastly, I do not replicate their use of the indices for importing and exporting across state lines and for long distance smuggling.[5]

The top panel of Table 3.3 reports 2SLS estimates. Column (ii) presents 2SLS estimates using the IRI data. Additional regressors are dummy variables for state. Estimates using the BOE tax receipt data over the time period January 1996-March 1999 are presented in Column (iii). Lastly, for another comparison with the BOE tax receipt results, Column (iv) presents 2SLS estimates using IRI data for the state of California only.

Standard errors are in parentheses below the estimates.[6] Lastly, implied short-run price elasticities and long-run price elasticities evaluated at the sample means are

[5] In constructing the long-distance smuggling index, BGM assume that states located over 1,000 miles away from Kentucky, Virginia, and North Carolina do not smuggle. All five states in my sample fall in this category. The short-distance smuggling index is a function of the difference between neighboring states' tobacco and cigarette taxes, which do not vary much, if at all, in my dataset.

[6] Although they are the correct standard errors for 2SLS estimates, they do not take into account the longitudinal nature of the panel data. The true standard errors are likely larger than those reported in the table.

reported.[7]

The lower panel of Table 3.3 presents results from the first stage regressions. The dependent variable is per capita consumption. Regressors include current, lagged and lead prices, as well as income. In the second stage, the fitted, rather than actual, values of lagged and lead per capita consumption are used.

Results based on IRI data

Compare the first two columns. Like BGM, I find that the coefficients on lagged and lead consumption are both positive and significant. According to the rational framework, the positive coefficient on lagged consumption implies that cigarettes are addictive and the positive coefficient on lead consumption implies that smokers are forward looking. The coefficient on lagged consumption is greater than the coefficient on lead consumption, as occurs if the discount rate is less than one.

To assess the relationship between consumption and price, focus on estimated elasticities, rather than the estimates of the coefficient on price. Otherwise, the difference in consumption units (annual vs. weekly vs. monthly) and price units

[7] The short-run price elasticities are calculated as follows:

$$\frac{da_t}{dP_t} = \frac{\theta_3}{\theta_2(1-\phi_1)(\phi_2)}$$

and the long-run price elasticities are calcaulated as follows:

$$\frac{da_\infty}{dP} = \frac{\theta_3}{\theta_2(1-\phi_1)(\phi_2-1)}$$

where ϕ_1 and ϕ_2 are given by:

$$\phi_1 = \frac{1-(1-4\theta_2^2\delta)^{\frac{1}{2}}}{2\theta_2}$$

$$\phi_2 = \frac{1+(1-4\theta_2^2\delta)^{\frac{1}{2}}}{2\theta_2}$$

See BGM Appendix A for derivation.

(1969 cents vs. 1996 dollars) confuses the comparison.

Comparison between the first two columns reveals that short-run price elasticity calculated using the IRI data is much greater than the elasticity that BGM find. In fact, consumers represented by the IRI data appear to be over three times as sensitive to price as the cigarette distributors represented by the tax receipt data that BGM use. One possible reason for this difference in response to permanent price changes may be the timing of purchases. The IRI data are weekly, whereas the BGM data are yearly. Another possible reason for the difference in elasticities may be measurement error in price in the BGM dataset. The prices used in their analysis are the retail prices paid by consumers, but the consumption variable is distributors' consumption of tax stamps.

Like BGM, I find that long-run price elasticities are greater than short-run price elasticities. They argue that this relationship lends support to the rational addiction model. However, the LeChatalier principle guarantees that all goods, not just addictive goods, are more elastic in the long run than in the short run. That is, because other consumption goods may be held fixed in the short run but not in the long run, short-run demand elasticities are lower than long-run demand elasticities. [8]

Lastly, BGM find that the relationship between cigarette consumption and income is positive and significant, whereas the coefficient on income in the IRI regression is negative and significant. The dummy variables for states control for any interstate differences in income. However, because the data are market-level data, the negative coefficient on income captures intrastate, market level differences in income and

[8] See Chapter 3 of Samuelson (1947).

consumption. The negative coefficient appears to capture the cross-sectional relationship between cigarette consumption and income. Of the total variation in the income data, 99.2% is due to variation between markets, whereas only 0.8% is due to variation within markets, over time. Therefore, the estimated income effect is unlikely to be caused by the downward trend of consumption as incomes have risen.

The BGM data span the years between 1955 and 1985, whereas the IRI data is very recent (1996-1999). The finding of a negative coefficient on income in the analysis that uses the IRI data may support the assertion of Keeler, et al. (1993) that during the past 30 years, cigarettes have moved from being a normal good to an inferior good.

Consider the first stage regression results in the second panel of Table 3.[9] The coefficients on current, lagged and lead prices all have the expected sign, according to the rational addiction model. According to Equation (3.1), lead and lagged prices have a negative effect on current consumption because these prices affect current consumption through lead and lagged consumption, which are not held fixed.

Results based on BOE data

The regression results based on the BOE data are striking. The coefficients on lagged and lead consumption are positive, but not significantly different from zero at conventional levels of significance. These results suggest that, for cigarette distributors, cigarettes are not an addictive good.

[9]The first stage regressions include not only prices, but also all other exogenous variables in the model. However, I only report the coefficients on price in Table 3.

Table 3.3: Replication of Becker Grossman Murphy (1994) 2SLS estimates

Dependent variable is per capita consumption.

Per capita consumption	BGM annual (i)	IRI data weekly (ii)	BOE data monthly (iii)	IRI data (CA only) weekly (iv)
a_{t-1}	0.418	0.470	0.546	0.303
	(0.047)	(0.059)	(0.458)	(0.104)
a_{t+1}	0.135	0.211	0.020	0.260
	(0.055)	(0.067)	(0.415)	(0.084)
P_t	-1.388	-0.117	-1.860	-0.198
	(0.155)	(0.024)	(1.219)	(0.044)
Y_t	0.837	-0.076	-1.587	-0.159
	(0.114)	(0.027)	(2.390)	(0.053)
SR price elasticity	-0.407	-1.314	-0.747	-3.490
LR price elasticity	-0.734	-2.791	-1.667	-5.222
$\overline{R}^2$	N/A	0.319	0.056	0.322
N	1415	3410	35	1881

Standard errors are in parentheses. Instruments include lagged and lead prices as well as the other regressors. For Column (i), prices are in 1969 cents and income is annual income in 1969 100's of dollars. Additional regressors are sets of dummy variables for state and year. In Columns (ii) and (iv), prices are in January 1996 dollars and income is weekly income in January 1996 1000's of dollars. Additional regressors include sets of dummy variables for month and state (Column (ii) only). In Column (iii) prices are in January 1996 dollars and income is monthly income in January 1996 1000's of dollars. An additional regressor in this regression is a time trend.

First stage regressions

Dependent variable is per capita consumption

	(i)	(ii)	(iii)	(iv)
P_t	N/A	-0.176	-5.866	-0.215
		(0.073)	(1.942)	(0.101)
P_{t-1}	N/A	-0.130	0.106	-0.093
		(0.057)	(1.395)	(0.083)
P_{t+1}	N/A	-0.014	4.355	-0.094
		(0.056)	(1.167)	(0.082)
$\overline{R}^2$	N/A	0.315	0.478	0.312

Price elasticities are much smaller than those based on the IRI California only data reported in Column (iv), but they are close to the price elasticities that BGM found (Column (i)). Cigarette distributors appear to be less sensitive to price changes than consumers. As mentioned above, one possible reason for the difference in price elasticity between the BOE and IRI California only results may be the timing of purchases. The BOE data is monthly, whereas the IRI data is weekly. Any intra-month response to price changes by cigarette distributors is not captured. As with the BGM data, another possible reason for the difference in elasticities may be measurement error in price in the BOE dataset. The prices used are the prices paid by consumers, whereas the consumption variable is distributors' consumption of tax stamps.

Consider the first stage regression results. The coefficient on current price is negative, as expected. The coefficient on lagged price is positive but small in magnitude and not significantly different from zero. Lastly, the coefficient on lead price is positive and significant. This result is most likely led by the last few months of the sample in which distributors appear to be stockpiling cigarettes in advance of the Proposition 10 tax hike (recall Figure 3.3). During this time, prices are rising due to the tobacco settlement, but sales are also increasing.

Results based on IRI data (California only)

Lastly, compare the parameter estimates for California only (Column (iv)) with the parameter estimates obtained when all fives states are included in the analysis. The estimated coefficients are all of the same sign. The coefficient on lagged consumption is smaller in magnitude and the coefficient on lead consumption is greater in magnitude than those reported in Column (ii). This implies a larger point estimate of the discount rate to Californians ($\delta = \frac{0.260}{0.303} = .86$) than the point estimate of the discount rate of the whole sample ($\delta = \frac{0.211}{0.470} = .45$).

California consumers are over twice as sensitive to price changes in the short run as consumers in the whole sample, and almost twice as sensitive to permanent price changes in the long run. That the California results are different from the general results suggests that perhaps allowing for state-specific coefficients, or at least state-specific time trends, is a necessary modification.

3.5.3 State-specific Time Trends

The BGM replications include dummy variables for months. These variables may account for general variation across time, but they can not capture state-specific or market specific time trends. The dummy variables for states may capture fixed effects, but not fixed trends. The graphs in Figure 3.1 indicate that each state has a general downward trend in consumption, but some states trend downward faster than others. For example, at the beginning of the time series, the difference between the highest and lowest per capita cigarette pack consumption (between California and Colorado) is 0.29. Both states have a downward trend in consumption, but Colorado's rate is faster than that of California. At the end of the time series, this difference is 0.07.

As in Gruber and Kőszegi, a regression using first differences of the independent and dependent variables may capture these fixed trends. However, because the IRI data pertain to weekly sales, there appears to be much negatively correlated week-to-week variation (see Figure 3.1) that could be easily explained by measurement errors in reporting or other timing issues.

As an alternative approach, I expand the preceding analysis by including state-specific linear time trends. Table 3.4 reports the results of a 2SLS regression using the IRI data with state-specific time trend variables. As in the analysis summarized in Table 3.3, instruments for lagged and lead consumption are lagged and lead prices. First stage regression results are presented in the lower panel of Table 3.4. The results presented in this table are analogous to those presented in Column (ii) of Table 3.3.

Consider first the estimated relationship between price and consumption. The point estimate of the coefficient on current price is similar to that in Column (ii) of Table 3.3 (-0.087 vs. -0.117) as are the estimated short-run price elasticity (-1.666 vs. -1.314) and long-run price elasticity (-2.770 vs. -2.791).

Again, the coefficients on lagged and lead consumption are positive and significant. Unlike the results of the 2SLS regression without state-specific time trends, the coefficient on lead consumption is greater than the coefficient on lagged consumption. Assuming that the econometric model is specified correctly, the finding of a larger coefficient on lagged than on lead consumption is difficult to reconcile. Recall that the coefficient on lead consumption divided by the coefficient on lagged consumption should yield the discount rate of consumers. The estimates of the coefficients in Table 3.4 imply that the point estimate of the discount rate is greater than one.

To examine whether the inclusion of state-specific time trends improves the fit of the model, I conduct a Wald test of joint significance of the state-specific time trends.

Table 3.4: 2SLS regression with state-specific time trend
Dependent variable is per capita consumption.

a_{t-1}	0.330
	(0.048)
a_{t+1}	0.431
	(0.055)
P_t	-0.087
	(0.019)
Y_t	-0.058
	(0.020)
SR price elasticity	-1.666
LR price elasticity	-2.770
$\overline{R}^2$	0.335
N	3410

Standard errors are in parentheses. Additional regressors are state-specific time trends and full sets of dummy variables for month and state. First stage regressors include lagged and lead prices as well as the other explanatory variables.

First stage regressions
Dependent variable is per capita consumption

P_t	-0.188
	(0.072)
P_{t-1}	-0.109
	(0.056)
P_{t+1}	-0.047
	(0.056)
$\overline{R}^2$	0.334

The F-statistic for the joint test, which is distributed $F[5, 3357]$, is 5.011 whereas the critical value for the 99th percentile is less than 1.69. Therefore, the test rejects the null hypothesis that the coefficients on the state-specific time trends are jointly zero at the 1% significance level.

3.5.4 Price Forecasts

Recall from Section 3 that the model estimated in the above regressions is given by Equation (3.1), in which current consumption is a function of actual future consumption. Without developing an explicit model of expectations, using actual future prices as an instrument for actual future consumption in the above regressions implies that agents perfectly forecast prices and consumption. Most likely, previous empirical studies have viewed actual future price and consumption as proxies for expectations generated under the assumption of rational expectations. Instead, it seems more reasonable to explicitly model expectations based on information available at the time that these expectations are formed. To account for this consideration, I estimate Equation (3.2) rather than Equation (3.1). That is, I include expected future consumption rather than actual future consumption as a right-hand side variable.

As instruments for lagged consumption and expected lead consumption, I use the one-period lag of price and one-period-ahead forecast of price. Price forecasts are estimated using current price and two lags of price:

$$E\left[P_{t+1}|\Phi_t\right] = \alpha_0 + \alpha_1 P_t + \alpha_2 P_{t-1} + \alpha_3 P_{t-2} + E\left[\varepsilon_{t+1}|\Phi_t\right] \tag{3.3}$$

Table 3.5 reports results from the price forecast regressions.

The results of the price forecast regressions reveal that current price predicts future price very closely. One-period lagged price adds a little more explanatory power, but the coefficients on two-period lagged price are not significantly different from zero.

Table 3.5: Price forecast regressions
Dependent variable is future price P_{t+1}.

	Multistate		CA only
	fixed trends	no fixed trends	
	(i)	(ii)	(iii)
P_t	0.913	0.915	0.854
	(0.017)	(0.017)	(0.023)
P_{t-1}	0.066	0.067	0.111
	(0.023)	(0.023)	(0.030)
P_{t-2}	0.016	0.013	0.031
	(0.017)	(0.017)	(0.023)
$\overline{R}^2$	0.987	0.987	0.984
N	3432	3432	1892

Standard errors are in parentheses. For Column (i) the forecast includes state-specific time trends whereas the forecast only includes a general time trend for Columns (ii) and (iii).

These declining coefficients over time are as expected. All together, in conjunction with the time trends, the independent variables predict over 98% of the variation in future price.

The high correlation between prices in adjacent time periods raises questions about the validity of lagged and lead prices as instruments for lagged and lead consumption in the basic model. That is, current consumption may not be independent of lagged and lead prices when lagged and lead consumption are held fixed. These results imply that perhaps two-period lead and lagged prices should be used as instruments.

Table 3.6 presents the results from 2SLS regressions where predicted, rather than actual, price is a first stage regressor. Columns (i) and (ii) use data from all five states. Column (i) allows for state-specific time trends, whereas Column (ii) does not. Lastly, Column (iii) uses the California data only. The lower panel presents first stage regression results.

The results in Column (i) are comparable to the results in Table 3.4. Results in Columns (ii) and (iii) are comparable to those in Table 3.3, Columns (ii) and (iv), respectively. Based on comparison with these preceding results, it is apparent that, for the analysis without state-specific time trends and the California only analysis, the estimated coefficients on price and income and the implied price elasticities are similar to those found before. When comparing the analyses that include state-specific time trends, we find that the coefficients on price and income are greater in magnitude when expected future price rather than actual future price is used as an instrument (-0.131 vs. -0.087 for price, and -0.090 vs. -0.058 for income). Short-run response to price changes is also more pronounced: short-run price elasticity of -2.014 when expected future price is an instrument as compared to a short-run price elasticity of -1.666 when actual future price is used as an instrument.

The estimates of the coefficients on lagged consumption are smaller in magnitude when expected rather than realized future prices are used as instruments. Lastly, when data from all states are used and state-specific time trends are included, the coefficient on future consumption is smaller in the 2SLS regression with expected future price as an instrument than in the regression with actual future price as an instrument. However, when state-specific time trends are not included or when only California data are used, the opposite occurs: the coefficient on future consumption is larger when expected future price is an instrument than when actual future price is used.

Consider now a comparison across columns of Table 3.6. Columns (i) and (ii)

Table 3.6: 2SLS Estimation of Equation 3.2
Dependent variable is per capita consumption.

	Multi-state		CA only
	fixed trends	no fixed trends	
	(i)	(ii)	(iii)
a_{t-1}	0.235	0.367	0.248
	(0.068)	(0.058)	(0.072)
$E[a_{t+1}]$	0.400	0.288	0.329
	(0.080)	(0.077)	(0.084)
P_t	-0.131	-0.122	-0.189
	(0.025)	(0.028)	(0.045)
Y_t	-0.090	-0.087	-0.155
	(0.027)	(0.030)	(0.051)
SR price elasticity	-2.014	-1.569	-3.747
LR price elasticity	-2.731	-2.691	-5.150
$\overline{R}^2$	0.336	0.320	0.325
N	3370	3370	1859

Standard errors are in parentheses. Expected price is forecast using 3 lags of price. Additional regressors in both columns are full sets of dummy variables for state and month. Additional regressors in the first column are state-specific time trends.

First stage regressions
Dependent variable is per capita consumption

	(i)	(ii)	(iii)
P_t	3.054	-4.618	-5.934
	(3.272)	(1.435)	(1.409)
P_{t-1}	0.178	-0.495	-0.969
	(0.290)	(0.136)	(0.235)
$E[P_{t+1}]$	-3.595	4.809	6.519
	(3.575)	(1.562)	(1.631)
$\overline{R}^2$	0.335	0.317	0.318

reveal the same relationships as were found when comparing Table 3.4 and Column (ii) of Table 3.3. Coefficients on income and price and similar, but the short run elasticity is greater in magnitude than the short run elasticity that results from the regression without a state-specific time trend. The higher estimated elasticity of short-run demand is due in part to the smaller coefficient on lagged consumption together with the larger coefficient on lead consumption. In fact, when state-specific time trends are included in the model, the coefficient on lagged consumption is smaller than the coefficient on lead consumption. Once again, a Wald test rejects at the 1% significance level the hypothesis that the state-specific time trends are jointly zero ($F[5, 3318] = 4.821$).

Comparing the adjusted R^2 from the second stage of the regression provides a cursory comparison between the two models represented by Equations (3.1) and (3.2). Consider the regression results using data from all five states. When expected price is used as an instrument, 0.336 of the variation in consumption is explained by variation in the independent variables that enter the second stage of the regression, including predicted values of lagged consumption and predicted values of expected lead consumption. When actual future price is used as an instrument, variation in the independent variables that enter the second stage of the regression explains 0.335 of the variation in consumption. The difference is probably not statistically significant, and this test is not a formal test of one model versus another because the models are not nested models.

To test whether one equation fits the data better than the other, I apply Davidson and MacKinnon's J test[10] to the second stage of the regression. Unfortunately, as can happen when comparing nonnested models, the test rejects one model in favor of the other, and then rejects the other in favor of the first. In comparing the models without state-specific time trends, the coefficient on the model described by Equation (3.1) is 1.324 with a standard error of 1.405, and therefore this model is rejected in favor of the model described by Equation (3.2). On the other hand, the coefficient on the Equation (3.2) model is -0.008 with standard error 1.347, and therefore this model is rejected for the Equation (3.1) model. The results for comparing the models with state-specific time trends are qualitatively similar.

Lastly, consider the first stage results in the lower panel of Table 3.6. Because expected future price is forecast using current price and two lags, there is obviously high correlation among these variables. From the first stage results, we see that this high correlation leads to anomalous estimates. In Column (i), none of the three coefficients is statistically different from zero. Only the coefficient on expected future price has the expected negative sign. In Columns (ii) and (iii), all the coefficients are significantly different from zero, but the coefficients on expected future price are positive.

[10]See Chapter 7 of Greene (1993).

3.5.5 Institutional Changes

Smoking Ban

Examination of Figure 3.1, together with Table 3.2, does not reflect any obvious effect of the ban on smoking in bars and taverns in California. In this section, I re-estimate the cigarette demand function using an indicator variable for the enactment of the smoking ban. It takes the value one if the state is California and the date is January 1, 1998 or later, and zero otherwise.

As reported in the top panel of Table 3.7, the coefficient on the smoking ban dummy variable is positive and significant, but small. In the first stage regression results, the coefficient on the smoking ban dummy variable is also positive and significant. This result seems surprising, given that the smoking ban should decrease smoking. One explanation may be that, following passage of the ban but prior to its implementation, cigarette consumption began declining in anticipation of its implementation, as predicted by the rational model. In fact, the rate of decline may decrease at the time the ban becomes effective, and therefore, the coefficient on the smoking ban is positive. However, without additional data on consumption before the passage of the legislation, this hypothesis is not testable.

Another important explanation is that the monthly dummy variables are picking up relevant variation across time. That is, this variation may be exactly what we would like the smoking ban dummy variable to capture. Therefore, it is possible that the smoking ban has no effect on cigarette consumption, but the positive coefficient on the smoking ban variable simply reflects a time variation that neither the general monthly variation nor the California-specific time trend is picking up.

An alternative approach is to exclude the monthly dummy variables from the analysis, and instead include various time trends and interactions between the smoking ban variable and these trends. As a basis of comparison, Table 3.8 presents results from replications of Table 3.3, Column (ii) that exclude monthly dummy variables. Column (i) of Table 3.8 includes a linear time trend instead, and Column (ii) includes adds a quadratic time trend. The results in Table 3.8 highlight the sensitivity of the model to the econometric specification. Comparison with Table 3.3, Column (ii), reveals that the regression with both linear and quadratic time trends yields similar coefficients on lagged and lead consumption. Note that, as in Table 3.3, Column (ii), the coefficient on lead consumption in Table 3.8, Column (ii) is highly significant. When the regression includes only a linear time trend, however, the coefficient on lead consumption is not significantly different from zero. The implication from this regression is that cigarette consumers are not forward looking. The coefficients on price and the short-run price elasticities are smaller, especially when both a linear and quadratic time trend are included in the regression.

In Column (i), the coefficient on the time trend is positive, but insignificant. This sign is surprising, given that there seems to be a steady downward trend in consumption (recall Figure 3.1). It appears that lagged and lead consumption are capturing the downward trend. The coefficients on these variables sum to less than one, implying a downward trend.

Table 3.7: Effect of smoking ban on consumption
Dependent variable is per capita consumption.

	fixed trends (i)	no fixed trends (ii)
a_{t-1}	0.328	0.408
	(0.048)	(0.054)
a_{t+1}	0.432	0.306
	(0.055)	(0.059)
P_t	-0.087	-0.107
	(0.019)	(0.022)
	[−1.635]	[−1.474]
Y_t	-0.058	-0.065
	(0.020)	(0.023)
Smoking Ban	-0.002	0.019
	(0.005)	(0.006)
$\overline{R}^2$	0.335	0.328
N	3410	3410

Standard errors are in parentheses. Additional regressors are full sets of dummy variables for month and state. Additional regressors in column (i) are state specific time trends.

First stage regressions
Dependent variable is per capita consumption

P_t	-0.188	-0.179
	(0.072)	(0.072)
P_{t-1}	-0.109	-0.115
	(0.056)	(0.056)
P_{t+1}	-0.047	-0.038
	(0.056)	(0.056)
Smoking ban	-0.003	0.082
	(0.021)	(0.011)
$\overline{R}^2$	0.334	0.325

Table 3.8: 2SLS Estimation of Equation 3.1 with time trend
Dependent variable is per capita consumption.

	(i)	(ii)
a_{t-1}	0.553	0.551
	(0.086)	(0.024)
a_{t+1}	0.101	0.386
	(0.139)	(0.034)
P_t	-0.085	-0.023
	(0.036)	(0.011)
Y_t	-0.132	-0.016
	(0.065)	(0.014)
t	0.398×10^{-4}	-0.209×10^{-3}
	(0.298×10^{-4})	(0.939×10^{-4})
t^2		0.135×10^{-5}
		(0.610×10^{-6})
SR price elasticity	-0.770	-0.570
LR price elasticity	-1.870	-2.778
$\overline{R}^2$	0.280	0.302
N	3410	3410

Standard errors are in parentheses. Additional regressors are full sets of dummy variables for state.

Table 3.9 presents results from 2SLS regression of Equation (3.1) that includes the dummy variable for the smoking ban but does not include monthly dummy variables. The regression associated with Column (i) includes a linear time trend. The regression in Column (ii) adds a quadratic time trend. The regression in Column (iii) includes a linear time trend as well as an interaction term between the smoking ban and the linear time trend. Finally, the regression in Column (iv) includes both trend variables, as well as interactions between the smoking ban and the trends.

Again, note the sensitivity of the model to changes in specification. The coefficient on lead consumption is not significantly different from zero in Column (i), but highly significant according to the other specifications. The coefficient on price is not significant at the 5% level in the regressions that include interaction terms (Columns (iii) and (iv)). The coefficient on price is negative in the regression that includes linear and quadratic terms, and interactions between these trends and the smoking ban dummy variable.

Table 3.9: Effect of smoking ban on consumption

Regression uses Time trends rather than Monthly dummy variables.
Dependent variable is per capita consumption.

	(i)	(ii)	(iii)	(iv)
a_{t-1}	0.566	0.542	0.450	0.504
	(0.077)	(0.025)	(0.023)	(0.020)
a_{t+1}	0.132	0.397	0.553	0.512
	(0.104)	(0.032)	(0.021)	(0.023)
P_t	-0.076	-0.022	-0.003	0.001
	(0.029)	(0.011)	(0.009)	(0.009)
Y_t	-0.112	-0.016	0.004	0.008
	(0.051)	(0.014)	(0.010)	(0.009)
Smoking ban	0.014	0.001	-0.010	0.115
	(0.007)	(0.002)	(0.015)	(0.054)
Smoking ban $\times$ t			0.688×10^{-4}	-0.002
			(0.131×10^{-3})	(0.001)
Smoking ban $\times$ t^2				0.637×10^{-5}
				(0.345×10^{-5})
t	-0.275×10^{-4}	-0.202×10^{-3}	0.110×10^{-4}	-0.356×10^{-4}
	(0.280×10^{-4})	(0.906×10^{-4})	(0.873×10^{-5})	(0.482×10^{-4})
t^2		0.127×10^{-5}		0.253×10^{-6}
		(0.559×10^{-6})		(0.293×10^{-6})
Overall effect of smoking ban at $t = 105$			-0.003	-0.025
$\overline{R}^2$	0.285	0.305	0.318	0.332
N	3410	3410	3410	3410

The results reported in Table 3.9 are mixed, but suggest that the smoking ban has no contemporaneous negative effect on consumption. The regression with a linear time trend only (Column (i)) yields a positive coefficient on the smoking ban dummy variable. Column (ii) reports a positive, but much smaller, coefficient on the smoking ban dummy variable from the regression that includes both linear and quadratic time trends.

On the other hand, the regression that includes a linear time trend and an interaction term between the smoking ban dummy variable and the trend (Column (iii)) yields a negative and insignificant coefficient on the dummy variable. Unexpectedly, the coefficients on the interaction term and the linear time trend are positive, indicating that the overall time trend is positive, and becomes steeper after the smoking ban goes into effect. As in the analysis without the smoking ban dummy variable, this anomalous result may be due in part to the presence of lag and lead consumption in the model. Lastly, noting that the smoking ban occurs at week $t = 105$, per capita consumption the first week after the ban is 0.003 packs less than it would be without the ban, assuming that the model in Column (iii) is the correct model. The positive coefficient on the interaction term implies that by October of 1998 (week 146), the per capita consumption is greater than it would be without the ban.

Lastly, consider the regression that includes linear and quadratic time trends, as well as interactions (Column (iv)). The coefficient on the smoking ban dummy variable is positive, but given the coefficients on the smoking ban dummy variable and the interaction terms, per capita consumption the first week after the ban is 0.025 packs less than it would be in the absence of the ban.

It would not be too surprising to find that the smoking ban has no short term effect on consumption for a number of reasons. To begin with, other studies of smoke-free workplace laws have found small effects, if any, on cigarette consumption. Evans, Farrelly and Montgomery (1999) find that smoke-free workplace bans decrease smoking participation among workers by 5%, and decrease consumption by about 10% among workers who are current smokers. Bar and restaurant workers who would be affected by the ban on smoking comprise about 3% of California's population according to estimates from the 1990 Census.

In a study of smoking among college students, Chaloupka and Wechsler (1997) find that workplace smoking bans have no effect on cigarette demand or smoking participation. They find that restaurant smoking bans have a small negative effect on smoking participation, but no effect on cigarette demand.

Casual observation suggests that there is a significant population of smokers who only smoke when they drink alcohol. The smoking ban may have the most significant effect on these casual smokers, but this effect may not be picked up by the grocery store data used in this analysis. Casual smokers may be less likely to purchase cigarettes at the grocery store during a regular grocery shopping trip. Rather, they seem more likely to buy cigarettes from a gas station, convenience store, or cigarette machine at a bar or restaurant while they are out for the night.

Lastly, there was (and still is) much uncertainty surrounding the ban. The effective date had already been postponed once. The original legislation (California Assembly Bill 13), signed by the governor on July 21, 1994, had set an implementation date of January 1, 1997. In September 1996, lawmakers approved Assembly Bill 3037 which postponed the implementation date until January 1, 1998. Through mid-1997, the Assembly debated an additional one year postponement of the ban. Even after implementation, California lawmakers are still discussing repealing the ban.

Proposition 10 Tax Increase

If consumers use all available information to forecast prices, the passage of Proposition 10 should lead them to incorporate the upcoming 50 cent tax increase into their future price forecast. Therefore, to forecast lead prices, I use current price and two lags of prices and a dichotomous variable for the pending tax hike. This dummy variable is equal to one in the two weeks before January 1, 1999, in any California market and zero for all other dates and all other markets. Note that this time period coincides with the tobacco litigation settlement. Because the Proposition 10 tax increase applies to California consumers only, whereas the unanticipated tobacco settlement price increase applies to consumers in all states, I can separately identify the effects of these two price increases. That is, the price forecast equation includes an indicator variable, denoted as Dec98, that is a dummy variable on the last two weeks of 1998 to control for the effects of the tobacco litigation settlement price increase. The variable denoted as Prop10 is the interaction between the indicator on the last two weeks of 1998 and a California state dummy variable.

The price forecast equation is:

$$E\left[P_{t+1}|\Phi_t\right] = \alpha_0+\alpha_1 P_t+\alpha_2 P_{t-1}+\alpha_3 P_{t-2}+\alpha_4 DEC98+\alpha_5 PROP10+E\left[e_{t+1}|\Phi_t\right] \quad (3.4)$$

Table 3.10 reports the price forecast regression results.

Table 3.10: Price forecast includes Prop 10 indicator
Dependent variable is future price.

	fixed trends (i)	no fixed trends (ii)
P_t	0.909	0.911
	(0.168)	(0.017)
P_{t-1}	0.072	0.072
	(0.023)	(0.023)
P_{t-2}	0.013	0.012
	(0.017)	(0.017)
Dec98	-0.003	-0.006
	(0.012)	(0.012)
Prop10	0.102	0.106
	(0.016)	(0.016)
$\overline{R}^2$	0.987	0.987
N	3432	3432

Standard errors are in parentheses. Dec98 is the dummy variable on the last two weeks of 1998. Prop10 is the indicator variable on the last two weeks of 1998, if the state is California. For Column (i) the forecast also includes state-specific time trends and for Column (ii), the forecast also includes a general time trend.

The results in Table 3.10 are comparable to those reported in Table 3.5, Columns (i) and (ii). As in the price forecasts that do not include the Prop10 or Dec98 indicators, these results indicate that current price predicts future price very closely. The coefficients on one-period lagged price are small, but significant, and the coefficients on two-period lagged prices are not significantly different from zero. The coefficient on the indicator for the last two weeks of December is negative, but not significantly different from zero. This is not surprising given that, holding the upcoming Proposition 10 tax increase in California fixed, together with holding current price and two lags of price fixed, the last two weeks of December 1998 should not affect future price. Lastly, the coefficient on the Prop10 variable is positive and significant, as expected. The impending tax hike does affect future prices.

Table 11 reports the results from 2SLS regression using expected lead prices rather than actual lead prices. Comparison between these results and those reported in Table 3.6, Columns (i) and (ii), yields insight as to how the inclusion of the Prop10 variable in the price forecasts affects the 2SLS regression results.

Compare the results from the model that includes state-specific time trends (Table 3.11, Column (i) vs. Table 3.6, Column (i)). The estimates are not much affected by the inclusion of the Proposition 10 variable. In fact, the differences are most likely not statistically significant. When the impending tax hike is included in the price forecast, the coefficients on lagged and lead consumption are larger (0.257 vs. 0.235 for the coefficient on a_{t-1}, and 0.419 vs. 0.400 for a_{t+1}) and the implied discount rate is also slightly larger (0.614 vs. 0.588). The coefficient on price and the estimated short run price elasticities are smaller, and the coefficient on income is also slightly smaller.

Consider the results from the model that does not include state-specific time trends. The coefficient on lagged consumption is larger (0.428 vs. 0.367) and the coefficient on expected lead consumption is smaller (0.204 vs. 0.288) which implies a smaller discount rate (0.477 vs. 0.785). The coefficients on price and income are similar.

Comparison across columns reveals the previously identified relationship when state-specific time trends are included in the regression: the coefficient on lead consumption is greater than the coefficient on lagged consumption. Once again, a Wald test indicates that the hypothesis that state-specific time trends are jointly zero can be rejected[11].

Table 3.11: 2SLS Estimation of Equation 3.2
Expected price forecast includes Prop10 indicator.
Dependent variable is per capita consumption.

	fixed trends (i)	no fixed trends (ii)
a_{t-1}	0.257	0.428
	(0.057)	(0.064)
$E\left[a_{t+1}\right]$	0.419	0.204
	(0.069)	(0.074)
P_t	-0.117	-0.134
	(0.024)	(0.027)
	$[-1.973]$	$[-1.480]$
Y_t	-0.079	-0.089
	(0.025)	(0.031)
$\overline{R}^2$	0.336	0.320
N	3370	3370

Standard errors are in parentheses. Additional regressors in both columns are full sets of dummy variables for state and month. Additional regressors in the first column are state-specific time trends.

First stage regressions
Dependent variable is per capita consumption.

	fixed trends (i)	no fixed trends (ii)
P_t	-0.047	-0.635
	(0.362)	(0.345)
P_{t-1}	-0.089	-0.157
	(0.066)	(0.066)
$E\left[P_{t+1}\right]$	-0.209	0.475
	(0.395)	(0.380)
$\overline{R}^2$	0.335	0.315

As in the first stage regression results reported in Table 3.6, the first stage re-

[11] $F[5,3318] = 4.678$

gression results reported in Table 3.11 are not as expected, due in part to the high correlation between $E[P_{t+1}]$ and P_t. In the first stage regression results for the model that includes state-specific time trends, the signs on the coefficients are all negative, as expected, but are not precisely estimated. In the results for the model that does not include state-specific time trends, only the coefficient on lagged price is negative.

To compare the model that includes the Proposition 10 indicator with the one that does not, I apply Davidson and McKinnon's J test to the second stage, because neither model nests the other. Unfortunately, using the Davidson and McKinnon J test to compare the various models yields inconclusive results. For example, in comparing the models that include state-specific time trends, the coefficient on the model without the Prop10 variable is -0.242 with a standard error of 2.93, and therefore, this model is rejected in favor of the model that includes the Prop10 variable in the price forecast. However, the coefficient on the model that includes Prop10 is -0.057 with a standard error of 0.166, and therefore, this model is rejected in favor of the model without the Prop10 variable in the price forecast. Likewise, comparing the model discussed in this section with the model that uses actual future consumption as a right-hand side variable yields the same inconclusive results.

Tobacco Settlement Price Increase and Proposition 10 Tax Increase

It is difficult to assess the effects of the tobacco settlement price increase and the Proposition 10 tax increase, because they occur during the same time period, even though the former affects all five states in the data, whereas the tax increase affects California only. However, the tobacco settlement price increase is unanticipated, whereas the Proposition 10 tax increase is anticipated. The rational model predicts that consumers will act in advance of an anticipated tax or price hike, because they anticipate that their consumption after the tax hike will decrease.

As a cursory glance at the effects of the tobacco settlement price increase and the Proposition 10 tax increase in California, consider again Figure 3.1. Ignoring the slight hoarding effects the week of the tax and price increases, the levels of consumption appear to decrease below the original levels before the tobacco settlement price increase. In California, after the Proposition 10 tax increase, the consumption level again decreases. Holding all else equal, the rational model predicts not only that Californians should decrease consumption in reaction to the tobacco settlement, but during the same time period, they should further decrease consumption in anticipation of the Proposition 10 tax increase. In the five weeks before the California election, consumption averages 0.189 packs per capita. In the eight weeks after the election, but before the tax increase, consumption averages 0.186 packs per person. Following the tax hike, consumption averages 0.147 pack per person. Without holding other effects such as price or time fixed, it appears that consumption has decreased slightly in anticipation of the imminent tax increase, as the rational model predicts. However, consider the consumption pattern of the other four states which do not have the 50 cent tax increase. Before November 3, 1998, average per capita consumption is 0.208. Between November 3 and January 1, 1999, average per capita consumption is 0.199 and after January 1, it is 0.173. In Colorado, Arizona, Nevada, and Washington, in which consumers are reacting to the tobacco settlement price increase, consumption decreases by about 4%. On the other hand, in California, consumption decreases by less than 2%. Comparison of these figures suggests that, contrary to the prediction of the rational model, cigarette consumers in California did not decrease consumption in anticipation of the Proposition 10 tax increase.

Of course, this is not a rigorous test of the rational model, because it ignores all other factors that affect consumption. An alternative approach is to compare the rational model with a restricted version of the model in which the coefficient on future consumption is constrained to be zero. The rational model predicts that the reaction will be different depending on whether the price increase is anticipated or unanticipated. A model in which addicts are not forward-looking, however, predicts the same reaction.

I use the full sample to estimate two models. The first is the rational addiction model with a linear time trend and state-specific dummy variables. The second is a restricted rational addiction model where the coefficient on lead consumption is constrained to be zero. Previous authors refer to this model as "myopic addiction." In the unrestricted model, I assume that agents forecast prices using the model in the previous section. Using the estimated coefficients, I predict per capita consumption for the period between January, 1999 and May, 1999. The estimation results are in Table 12.[12]

[12] A better test of the models may be to estimate the two models using only the first 143 weeks of data (through the end of September, 1998). Using these estimates to forecast consumption for the next 32 weeks (i.e., out of sample), I could then compare the forecasts with actual per capita consumption. However, in the first 143 weeks of the dataset, there are no price increases that are nearly as large as the ones following the tobacco settlement and the implementation of the Proposition 10 tax increase. Without the price changes, it appears that the estimates of the models are driven by noise. For example, using only the 143 weeks of data to estimate the unrestricted model

Table 3.12: 2SLS Estimation of Equation 3.2 with time trend
Dependent variable is per capita consumption.

	Unrestricted model (i)	Restricted model ($a_{t+1}=0$) (ii)
a_{t-1}	0.578	0.488
	(0.031)	(0.033)
$E[a_{t+1}]$	0.311	
	(0.049)	
P_t	-0.032	-0.124
	(0.014)	(0.008)
Y_t	-0.038	-0.196
	(0.024)	(0.015)
t	0.0021	-0.0001
	(0.0113)	(0.0111)
$\overline{R}^2$	0.281	0.280
N	3370	3374

Additional regressors are dummy variables for state.

Note, from Table 3.12, that the estimated coefficient on price is almost three times as large in the restricted model than in the unrestricted model. The coefficient on lagged consumption is smaller, and the coefficient on income is almost four times as large in the restricted model than in the unrestricted model.

Using the estimates in Table 3.12, I predict per capita consumption for October, 1998, through May, 1999. For the right-hand side variable of lead consumption in the unrestricted model, I use the reduced form predictions for lead consumption. The sum of squared deviations between actual per capita consumption and consumption predicted by the unrestricted model is 2.48. The sum of squared deviations between actual consumption and consumption predicted by the restricted model is 3.37.

gives insignificant estimates of the coefficients on lagged and lead consumption. In fact, the coefficient on lagged consumption is negative. Furthermore, the coefficient on price is over 5 times as large as the coefficient that results from estimation using the full sample. Using these estimates to predict consumption for October, 1998, through May, 1999, yields predictions of negative consumption in both models.

Figure 3.4 plots the average of realized per capita consumption for markets in California, as well as predicted values for both the restricted and unrestricted models. Figure 3.5 plots the average of realized per capita consumption for markets in the other states, along with the predicted values.

Examination of both figures reveals that the models' predictions are very close to one another. Consider Figure 3.4 first. Both models perform well until the tobacco settlement price increase. At that time, the models overreact to the price increase. At the January 1, 1999 tax hike, both models again predict a stronger reaction than actually occurred.

Figure 3.5 reveals that the models predict per capita consumption much better in the other states than in California. Before the tobacco settlement, the predicted values of consumption are higher than actual consumption. As with the California data, the models predict stronger responses to the price increase than actually occur. Therefore, after the price increase, the predictions from both models follow actual consumption closely.

The unrestricted model does slightly better than the restricted model. The predictions of the restricted and the unrestricted models are very similar, even though the unrestricted model should predict a different reaction to an unanticipated price change than an anticipated price change, whereas it is irrelevant to the restricted model whether a price change is anticipated or unanticipated. However, the predictions are more similar in the other states, in which all price changes are unanticipated, as the rational model would predict. One possibility is that both price changes are

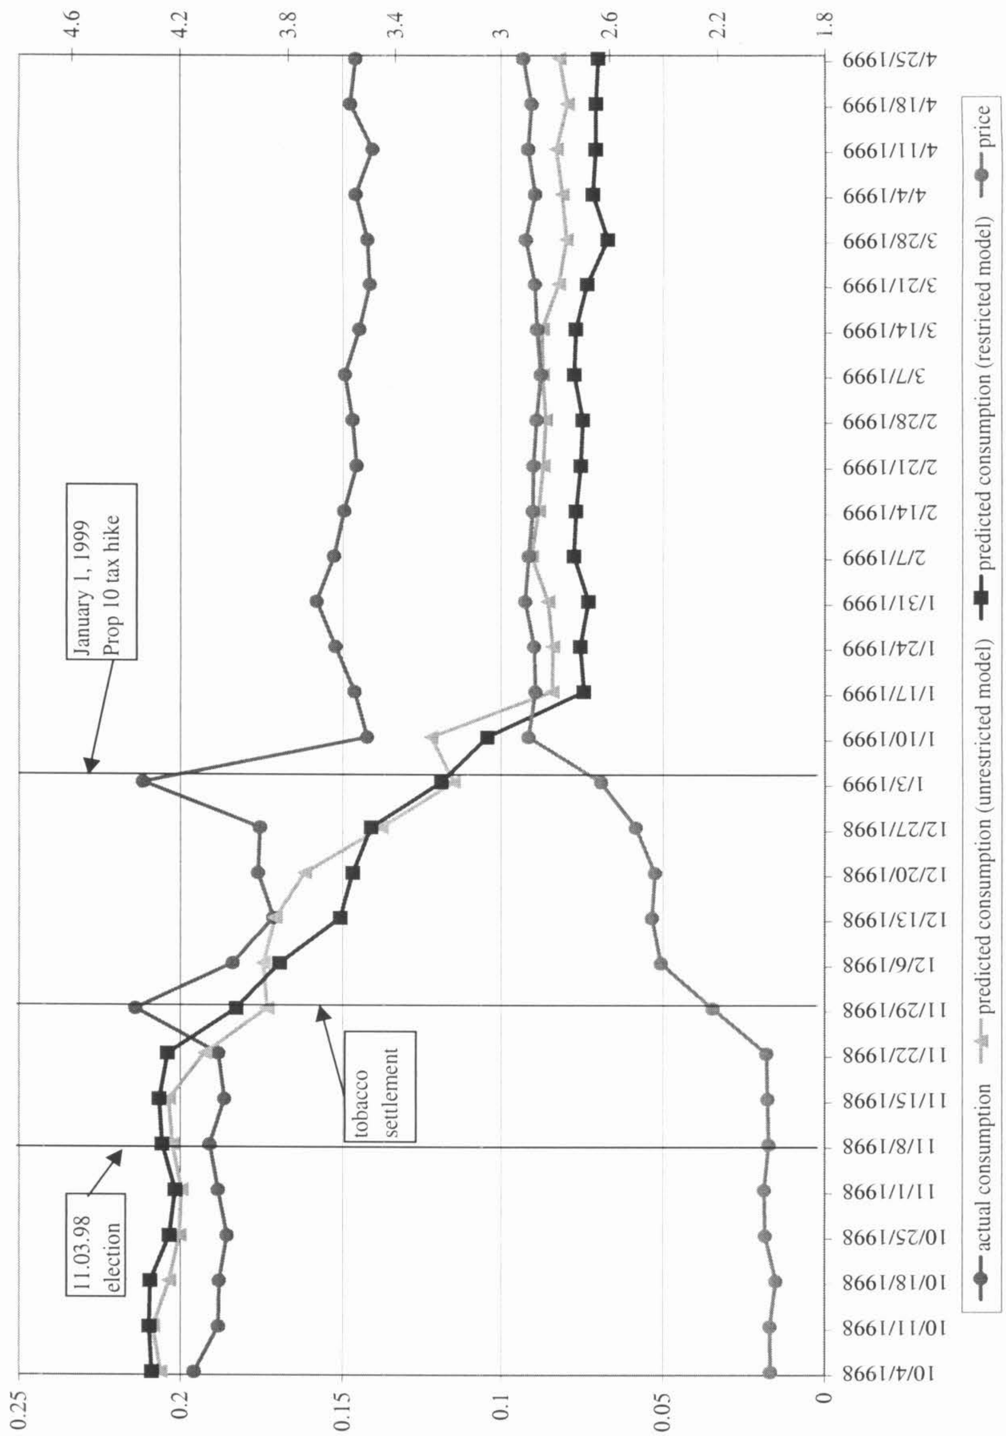

Figure 3.4: Actual and Predicted Per Capita Consumption Averaged over California Markets

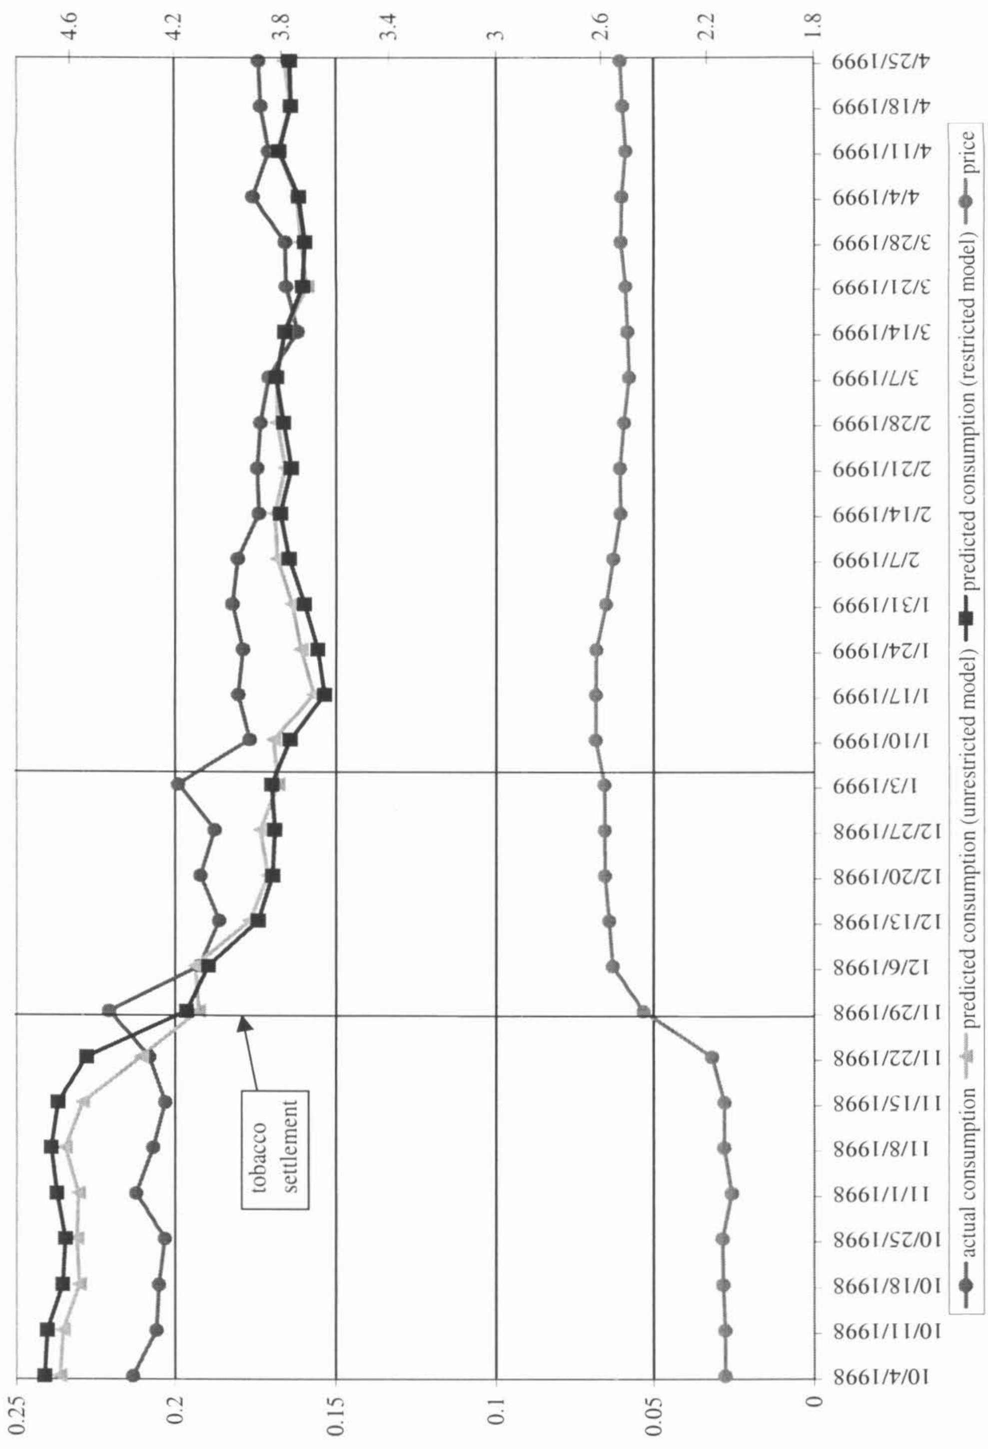

Figure 3.5: Actual and Predicted Per Capita Consumption Averaged over AZ, CO NV, and WA Markets

unanticipated to the consumers: California consumers did not anticipate the Proposition 10 tax hike, despite the election and publicity surrounding the proposition. Another possibility is that the stores began raising prices prior to the tax hike, and the models predict that the consumers are merely responding to the price increases.

Lastly note that neither the unrestricted nor the restricted model predicts the hoarding of cigarettes that occurs, most noticeably in the last week of December. Even though the data used in this analysis measure sales, the model of rational addiction (as well as the restricted model) describes consumption behavior, rather than purchasing behavior, and therefore, does not predict hoarding.

In fact, when the analysis is applied using aggregated monthly data, rather than weekly data, so that the effects of hoarding are smoothed over, predicted consumption more closely follows actual consumption. The mean squared deviation between actual per capita consumption and consumption predicted by the unrestricted model is 0.0035 and the mean squared deviation between actual and predicted consumption for the restricted model is 0.0014, which are less than the mean squared deviations from the original analysis in which weekly data was used.

Hoarding behavior is not inconsistent with the rational addiction model. If the anticipated tax hike will not cause an agent to quit smoking, then subject to budget and storage constraints, it is forward-looking and utility-maximizing to hoard cigarettes before the tax increase. It is not clear what the optimal level of hoarding is, but the consumers represented in this dataset do not appear to be hoarding very much. There is a small spike in sales in the week of the tobacco settlement price hike, the week ending November 29, and there is larger spike in sales the last week of December. A finding that consumers do not hoard enough would lend support to studies that have found that consumers engage in purchase quantity rationing when purchasing goods that are vices, such as cigarettes (see Wertenbroch, 1998).

3.6 A Methodological Note on Measurement Error and Serial Correlation

The positive and significant coefficient on future (or expected future) consumption may imply that consumers are indeed forward looking and rational. However, there are some serious issues that this test of rational addiction raises. One issue is the question of how serial correlation might affect the estimates. Recall from Figures 3.1 and 3.2 that cigarette sales follow a general upward trend, while prices follow a general downward trend. It is possible that the coefficient on future consumption is picking up the serial correlation that is not fully accounted for by the time trend and time variation variables.

A related issue is that the positive and significant coefficient may simply be a statistical artifact due to measurement error. Grether and Maddala (1973) demonstrate how measurement error in the independent variable can lead to a non-zero coefficient on lags of the independent or dependent variables, even when the true model contains no lags. Following the same type of analysis, we can show that measurement error in the dependent variable can also result in biased coefficients on lags or leads of dependent variables.

Suppose, for simplicity, that consumption of addictive goods depends only on current price, which is exogenously given. For ease of exposition, let the variables be expressed as deviations from their means. Then the true model is given by:

$$a_t = \gamma P_t + e_t \tag{3.5}$$

where, as before, a_t and P_t are consumption and price, respectively, at time t. However, instead of observing consumption, we observe expenditures, state tax receipts, or perhaps survey responses. That is, we observe

$$z_t = a_t - \eta_t$$

where η_t is measurement error. Assume

$$E[e_t] = E[\eta_t] = 0$$

and for all j,

$$E[e_t P_{t-j}] = E[\eta_t P_{t-j}] = E[\eta_t a_{t-j}] = E[\eta_t e_{t-j}] = 0$$

Lastly, to simplify the computation, assume

$$E[P_t] = 0$$

Now, suppose the econometric model is misspecified as

$$a_t = \beta a_{t+1} + \gamma P_t + e_t \tag{3.6}$$

Then the estimated model yields

$$\widehat{z}_t = \widehat{\beta} z_{t+1} + \widehat{\gamma} P_t$$

In the limit, the least squares estimate of β is given by:

$$\widehat{\beta} = \frac{\sigma_{zz}(1)\sigma_P^2 - \sigma_{Pz}\sigma_{Pz}(1)}{\sigma_z^2\sigma_P^2 - \sigma_{Pz}^2(1)} \tag{3.7}$$

where

$$\sigma_{xy}(j) = cov(x_t y_{t+j})$$

Using sample variances and covariances of sales and prices from the IRI data, the bias is estimated to be 0.9912, assuming that the coefficient on lead consumption is zero.

Now suppose that the true model is given by Equation (3.6) rather than Equation (3.5), but consumption is again measured with error. In the limit, we can show that

$$\widehat{\beta} = \beta \frac{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2} - \frac{\sigma_P^2\sigma_{\eta\eta}(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2} \tag{3.8}$$

Derivation of Equations (3.7) and (3.8) can be found in the Appendix.

If there is no serial correlation in the measurement error, then the second term drops out and the estimate of β is biased toward zero because the denominator of the first term is always greater than the numerator. Notice that this bias toward zero is similar to the downward bias that results from the classical problem of measurement error in an independent variable. The sign of the second term depends on the serial correlation of the measurement error. If the errors are positively correlated, then the estimate of β is even smaller, assuming that β is positive. On the other hand, if the errors are negatively correlated, the sign of the second term is negative, and it is inconclusive as to whether or not $\widehat{\beta}$ is biased downward.

This measurement error bias may have affected the analysis of Keeler, et al. (1993) in which they estimate the rational model using California state cigarette and tobacco tax data. In their analysis, rather than one-period lagged and lead consumption, they use previous 12 month and subsequent 12 month moving averages, respectively. They argue: "Use of only one-month lead and lag values in [Equation (3.1)] generated nonsense results quite anithetical to the rational-addiction hypothesis, specifically insignificant and/or negative values of both lead and lagged quantity demanded." (p.14)

3.7 Conclusion

Previous empirical studies of addictive goods typically use either state cigarette tax receipt data or survey data. Analysis of grocery store scanner data is a promising approach to the empirical study of cigarette addiction. Unlike tax receipt data, grocery store sales reflect sales to consumers rather than sales by distributors to wholesalers or retailers. Furthermore, grocery store sales data have the advantage over survey data that there is no concern about adjusting answers due to social conformity.

The results I find are mixed with respect to the rational addiction model, and there are some important issues raised. The results are very sensitive to econometric specification, and some results are difficult to reconcile within the rational framework. For example, in the regressions that include a linear time trend rather than monthly variation (Tables 3.8 and 3.9), the coefficient on lead consumption is not significantly different from zero. Furthermore, the specifications of the model with state-specific time trends fit the data better than those without. However, when state-specific time trends are included in the model, the coefficients on lagged consumption are smaller than the coefficients on lead consumption, yielding point estimates of discount rates that exceed unity.

I believe that the model represented by Equation (3.2), in which current consumption depends on expected future consumption rather than actual future consumption, as in the model represented by Equation (3.1), is the more realistic model. It is noteworthy that, although model (3.2) uses less information (in particular, actual future consumption is not included in the model), model (3.2) does equally well, if not better, at explaining variation in the dependent variable than model (3.1). Furthermore, the data reveal strong downward trends in consumption that differ across states, and I reject the hypothesis that the state-specific time trends are jointly zero in all regressions in which they are included. Therefore, the results I focus on are those given in Table 3.11, Column (i) based on model (3.2) with state-specific time trends. Support for the rational addiction model is found in that the coefficients on past and expected future consumption are positive and significant. However, as mentioned above, the results also cast doubt on the rational addiction model because of the negative point estimate of the rate of time preference.

Furthermore, the analyses that use price forecasts raise questions about the validity of lagged and lead prices as instruments. However, the majority of empirical studies rely on the validity of these instruments.

In analyzing the effects of policy changes, I find that the ban on smoking in bars and restaurants does not have a contemporaneous negative effect on consumption of cigarettes. In comparing the effects of an anticipated tax change with the effects of an unanticipated price change, I find that an unrestricted rational addiction model performs slightly better than a model in which the coefficient on future consumption is restricted to be zero.

This analysis may not adequately account for the correlation structure of these panel data. For example, it does appear that the errors are serially correlated. Furthermore, the price forecast equation estimates also suggest strong correlation between prices realized in adjacent periods. It should be recognized, however, that this analysis, as well as previous empirical studies, have found that the econometric implementation of the rational model is not very robust to different specifications. Therefore, although investigating the degree of serial correlation in the data may be fruitful, correcting for serial correlation may not offer more insight.

A larger project would involve conducting a similar econometric analysis using similar data on sales of goods that are believed to be non-addictive. Because many of the issues raised are problems with the econometric specification of the model, rather than the theoretical model itself, using consumption of other non-addictive goods as a comparison could help differentiate results that are artifacts of the econometric specification from results that lend support to the rational model.

Chapter 4 Conclusion

This dissertation presents a theoretical and empirical analysis of consumption of addictive goods.

The pattern of quitting and relapse that is prevalent among substance abusers has proved to be a difficult problem for addiction researchers from all fields. Previous economic research on addiction has not fully utilized findings from other disciplines. Research on people's misjudgment of the severity of future consequences of substance use explain why people might begin to use an addictive substance that they eventually choose to quit. Research on conditioned responses sheds light on why addicts who decide to stop using addictive substances begin to consume again, even though they no longer have misperceptions about the negative effects of consumption.

I develop a model of addiction of addictive goods that departs from conventional models of consumption in two ways: first, by introducing craving that can be induced by the presence of environmental cues such as locations, persons, or drug paraphernalia, and second, by allowing for the possibility that individuals misperceive the severity of the future consequences from consuming addictive substances. Whereas for addicts in "rational addiction" models addiction is simply a utility-maximizing decision, addicts in my model may experience regret and could improve their welfare by controlling or eliminating environmental cues that generate cravings. Agents in my model may exhibit consumption patterns that resemble a pattern of quitting and relapse.

The third chapter presents an empirical analysis of the Becker and Murphy model of rational addiction using data on grocery store sales of cigarettes. Analysis of grocery store scanner data on cigarette sales is a promising approach to the empirical study of cigarette addiction. Unlike tax receipt data, grocery store sales reflect sales to consumers rather than sales by distributors to wholesalers or retailers. Sales data also have an advantage over survey data in that there is no concern about adjusting answers due to social conformity.

The results I find are mixed with respect to the rational model, but there are some important issues raised. First of all, the results are very sensitive to econometric specification, and some results are difficult to reconcile within the rational framework. Secondly, the analyses that use price forecasts raise questions about the validity of lagged and lead prices as instruments, an approach adopted in previous empirical studies.

I analyze the effects of three policy changes: the implementation of the California Smoke-Free Workplace ban on smoking in bars and restaurants; a nationwide price increase due to a settlement between tobacco companies and state governments; and a 50 cent tax increase in California. I find that the ban on smoking in bars and restaurants does not have a contemporaneous negative effect on consumption of cigarettes. I compare the effects of an anticipated tax change with those of an unanticipated price change. The rational addiction model predicts that, because agents are forward-looking and because of the lack of intertemporal seperability with respect to addictive goods, agents' behavior will change in advance of an anticipated price change. Indeed, I find that an unrestricted rational addiction model performs slightly better than a model in which the coefficient on future consumption is restricted to be zero.

A future project would focus on distinguishing between the rational model of addiction and various behavioral models of addiction, including my own. Because the various models have been designed to explain why individuals would become addicts using a decision-theoretic framework, many of these models yield similar observable outcomes. It is important to develop sharp predictions that are empirically testable in order to distinguish between the various models. For example, the models might differ in their predictions as to whether people will seek treatment and the success of treatment programs. Alternatively, the models might predict different responses to price changes or policy changes.

Bibliography

[1] Becker, Gary S.; Grossman, Michael and Murphy, Kevin M. "An Empirical Analysis of Cigarette Addiction." *American Economic Review* 84, June 1994, 396-418.

[2] Becker, Gary S. and Murphy, Kevin M. "A Theory of Rational Addiction." *Journal of Political Economy* 96, August 1988, 675-700.

[3] Boyer, Marcel. "A Habit Forming Optimal Growth Model." *International Economic Review* 19, October 1978, 585-609.

[4] California State Board of Equalization. "Cigarette and Tobacco Products Tax Law." Publication 93, July 1999.

[5] Chaloupka, Frank. "Rational Addictive Behavior and Cigarette Smoking." *Journal of Political Economy* 99, August 1991, 722-42.

[6] Chaloupka, Frank J. and Wechsler, Henry. "Price, Tobacco Control Policies and Smoking among Young Adults." *Journal of Health Economics* 16, June 1997, 359-73.

[7] Champion R. A. and Bell, David S. "Attitudes toward Drug Use: Trends and Correlations with Actual Use." *The International Journal of the Addictions* 15, 1980, 551-67.

[8] Eikelboom, Roelof and Stewart, Jane. "Conditioning of Drug-Induced Physiological Responses." *Psychological Review* 89, 1982, 507-28.

[9] Evans, William N.; Farrelly, Matthew C. and Montgomery, Edward. "Do Workplace Smoking Bans Reduce Smoking?" *American Economic Review* 89, September 1999, 728-47.

[10] Frawley, P. Joseph. "Neurobehavioral Model of Addiction: Addiction as a Primary Disease." in *Visions of Addiction: Major Contemporary Perspectives on Addiction and Alcoholism.* ed. by Peele, Stanton, 1988, Lexington: Lexington Books, 25-43.

[11] Gawin, Frank H. "Cocaine Addiction: Psychology and Neurophysiology." *Science* 251, March 1991, 1580-6.

[12] Goldstein, Avram and Kalant, Harold. "Drug Policy: Striking the Right Balance." *Science* 249, September 1990, 1513-21.

[13] Greene, William H. *Econometric Analysis.* 2nd ed. New York: Macmillan. 1993.

[14] Grether, David and Maddala, G.S. "Errors in Variables and Serially Correlated Disturbances in Distributed Lag Models." *Econometrica* 41, March 1973, 255-62.

[15] Grossman, Michael and Chaloupka, Frank J. "The Demand for Cocaine by Young Adults: A Rational Addiction Approach." *Journal of Health Economics* 17, August 1998, 427-74.

[16] Grossman, Michael; Chaloupka, Frank J. and Sirtalan, Ismail. "An Empirical Analysis of Alcohol Addiction: Results from the Monitoring the Future Panels." *Economic Inquiry* 36, January 1998, 39-48.

[17] Gruber, Jonathan and Kőszegi, Botond. "Rational and Irrational Addiction: Theory and Evidence." Working paper, 1999.

[18] Herrnstein, Richard J. and Prelec, Drazen. "A Theory of Addiction." in *Choice over Time.* ed. by Loewenstein, George and Elster, Jon, 1992, New York: Russell Sage Foundation, 331-60.

[19] Howe, Kenneth. "Price Rise Puts Heat On Smokers Run on Cigarette Sales and Internet Vendors." *San Francisco Chronicle*, November 27, 1998.

[20] Jones, Andrew M. "Adjustment Costs, Withdrawal Effects, and Cigarette Addiction." *Journal of Health Economics* 18, 1999, 125-37.

[21] Keeler, Theodore E.; Hu, Teh-Wei; Barnett, Paul G., and Manning Willard G. "Taxation, Regulation, and Addiction: a Demand Function for Cigarette Based on Time-series Evidence." *Journal of Health Economics* 12, April 1993, 1-18.

[22] Koob, George F. and Le Moal, Michel. "Drug Abuse: Hedonic Homeostatic Dysregulation." *Science* 278, October 1997, 52-8.

[23] Laibson, David. "A Cue-Theory of Consumption." Working paper, 1999.

[24] Leshner, Alan I. "Addiction is a Brain Disease, and it Matters." *Science* 278, October 1997, 45-7.

[25] Los Angeles Times. "Settlement Spurs Record Cigarette Price Hike." November 24, 1998.

[26] Loewenstein, George. "A Visceral Account of Addiction." in *Getting Hooked: Rationality and Addiction.* ed. by Elster, Jon and Skog, Ole-Jørgen, 1999, Cambridge, England: Cambridge University Press, 235-264.

[27] Loewenstein, George; O'Donoghue, Ted and Rabin, Matthew. "Projection Bias in Predicting Future Utility." Working paper, 1999.

[28] Marlatt, G. Alan. "Craving Notes." *British Journal of Addiction* 82, 1987, 42-3.

[29] Nichols, Mark W. "Casino Gambling and Addictive Behavior: An Empirical Analysis." Working paper, 1999.

[30] O'Brien, Charles P. "A Range of Research-Based Pharmacotherapies for Addiction." *Science* 278, October 1997, 66-70.

[31] O'Brien, Charles P. and McLellan, A. Thomas. "Myths about the Treatment of Addiction." *Lancet* 347, January 1996, 237-240.

[32] O'Donoghue, Ted and Rabin, Matthew. "Addiction and Self-Control." Working paper, 1999.

[33] Oetting, E.R. and Beauvais, Fred. "Common Elements in Youth Drug Abuse: Peer Clusters and Other Psychosocial Factors." in *Visions of Addiction: Major Contemporary Perspectives on Addiction and Alcoholism.* ed. by Peele, Stanton, 1988, Lexington: Lexington Books, 141-161.

[34] Orphanides, Athanasios and Zervos, David. "Optimal Consumption Dynamics with Non-Concave Habit-Forming Utility." *Economics Letters* 44, 1994, 67-72.

[35] Orphanides, Athanasios and Zervos, David. "Rational Addiction with Learning and Regret." *Journal of Political Economy 103, 1995,* 739-758.

[36] Pollak, Robert A. "Habit Formation and Dynamic Demand Functions." *Journal of Political Economy* 78, July-August 1970, 745-63.

[37] Pomerleau, Ovide and Pomerleau, Cynthia. "A Biobehavioral View of Substance Abuse and Addiction." in *Visions of Addiction: Major Contemporary Perspectives on Addiction and Alcoholism.* ed. by Peele, Stanton, 1988, Lexington: Lexington Books 117-139.

[38] Robins, Lee N.; Davis, Darlene H. and Goodwin, Donald W. "Drug Use by U.S. Army Enlisted Men in Vietnam: a Follow-up on their Return Home." *American Journal of Epidemiology* 99, April 1974, 235-49.

[39] Robins, Lee N. "Vietnam Veterans' Rapid Recovery from Heroin Addiction: a Fluke or Normal Expectation?" *Addiction* 88, 1993, 1041-54.

[40] Rohsenow, Damaris J. "Drinking Habits and Expectancies about Alcohol's Effects for Self Versus Others." *Journal of Consulting and Clinical Psychology* 51, 1983, 752-6.

[41] Ryder, Harl E. and Heal, Geoffrey M. "Optimal Growth with Intertemporally Dependent Preferences." *Review of Economic Studies* 40, 1973, 1-33.

[42] Samuelson, Paul A. *Foundations of Economic Analysis*, 1958, Cambridge: Harvard University Press.

[43] Schelling, T.C. "Egonomics, or the Art of Self-Management." *American Economic Review 68*, May 1978, 290-4.

[44] Siegel, Shepard; Hinson, Riley E.; Krank, Marvin D. and Jane McCully. "Heroin 'Overdose' Death: Contribution of Drug Associated Environmental Cues." *Science* 216, 1982, 436-7.

[45] Siegel, Shepard; Krank, Marvin D. and Riley E. Hinson. "Anticipation of pharmacological and nonpharmacological events: Classical conditioning and addictive behavior." in *Visions of Addiction: Major Contemporary Perspectives on Addiction and Alcoholism.* ed. by Peele, Stanton, 1988, Lexington: Lexington Books, 85-116.

[46] Stigler, George J. and Gary S. Becker. "De Gustibus Non Est Disputandum." *American Economic Review* 67, March 1977, 76-90.

[47] Suranovic, Steven M.; Goldfarb, Robert S. and Leonard, Thomas C. "An economic theory of cigarette addiction." *Journal of Health Economics* 18, 1999, 1-29.

[48] Thaler, Richard H. and H. M. Shefrin. "An Economic Theory of Self-Control." *Journal of Political Economy* 89, 1981, 392-406.

[49] Wertenbroch, Klaus. "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice." *Marketing Science* 17, 1998, 317-37.

[50] Winston, Gordon C. "Addiction and Backsliding: A Theory of Compulsive Consumption." *Journal of Economic Behavior and Organization* 1, 1980, 295-324.

Appendix A Proof Appendix

For ease of exposition, let $w(a) = u(y - pa, a)$ and let $U(a) = w(a) + v(a, x)$ and $\widetilde{U}(a) = w(a) + \widetilde{v}(a, x)$.

Proof of Proposition 1: By our assumptions on $u(\cdot,\cdot)$, $v(\cdot,\cdot)$ and $\widetilde{v}(\cdot,\cdot)$, the partial derivative of current utility (with the perceived tolerance function) with respect to the consumption stock, which is given by

$$\frac{\partial}{\partial x} \quad \left[w(\frac{1}{\beta}(x' - \alpha x)) + \widetilde{v}(\frac{1}{\beta}(x' - \alpha x), x)\right] =$$
$$-\frac{\alpha}{\beta}w'(\frac{1}{\beta}(x' - \alpha x)) - \frac{\alpha}{\beta}\widetilde{v}_1(\frac{1}{\beta}(x' - \alpha x), x) + \widetilde{v}_2(\frac{1}{\beta}(x' - \alpha x), x)$$

and the partial derivative of current utility (with the true tolerance function) with respect to the consumption stock, which is given by

$$\frac{\partial}{\partial x} \quad \left[w(\frac{1}{\beta}(x' - \alpha x)) + v(\frac{1}{\beta}(x' - \alpha x), x)\right] =$$
$$-\frac{\alpha}{\beta}w'(\frac{1}{\beta}(x' - \alpha x)) - \frac{\alpha}{\beta}v_1(\frac{1}{\beta}(x' - \alpha x), x) + v_2(\frac{1}{\beta}(x' - \alpha x), x)$$

are both strictly increasing x'. The assumptions on the utility and tolerance functions satisfy Assumptions 1-4 of Orphanides and Zervos (1994), and therefore their Lemmas 2 and 3 and Proposition 1 apply. ∎

Proof of Proposition 2: I will prove

$$V(x^A, x^B) = \frac{\mu}{1 - \delta(1 - \mu)} V^{SE}(x^A | \widetilde{\delta} = \frac{\delta\mu}{1 - \delta(1 - \mu)}) + \frac{1 - \mu}{1 - \delta\mu} V^{SE}(x^B | \widetilde{\delta} = \frac{\delta(1 - \mu)}{1 - \delta\mu})$$

The proof for

$$\widetilde{V}(x^A, x^B) = \frac{\mu}{1-\delta(1-\mu)}\widetilde{V}^{SE}(x^A|\widetilde{\delta} = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu}\widetilde{V}^{SE}(x^B|\widetilde{\delta} = \frac{\delta(1-\mu)}{1-\delta\mu})$$

is the same. Begin by proving the following Lemma

Lemma 7 *If $V^*()$ is the solution to*

$$V^*(x) = \max_a \mu\left[w\left(a\right) + v(a,x) + \delta V^*(\alpha x + \beta a)\right] + \left(1-\mu\right)V^*(x)$$

then

$$V^*(x) = \frac{\mu}{1-\delta(1-\mu)}V^{SE}(x|\delta = \frac{\delta\mu}{1-\delta(1-\mu)})$$

Proof:

$$\begin{aligned}
V^*(x) &= \frac{\mu}{1-\delta(1-\mu)}V^{SE}(x|\frac{\delta\mu}{1-\delta(1-\mu)}) \\
&= \max_{a\in[0,\frac{y}{p}]}\frac{\mu}{1-\delta(1-\mu)}\left[U(a,x) + \frac{\delta\mu}{1-\delta(1-\mu)}V^{SE}(\alpha x + \beta a|\frac{\delta\mu}{1-\delta(1-\mu)})\right] \\
&= \max_{a\in[0,\frac{y}{p}]}\frac{\mu}{1-\delta(1-\mu)}\left[U(a,x) + \delta V^*(\alpha x + \beta a)\right] \\
&= \max_{a\in[0,\frac{y}{p}]}\mu\left[U(a,x) + \delta V^*(\alpha x + \beta a)\right] + (1-\mu)\delta V^*(x)
\end{aligned}$$

∎

Now, let

$$V(x^A, x^B) = \frac{\mu}{1-\delta(1-\mu)}V^{SE}(x^A|\delta = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu}V^{SE}(x^B|\delta = \frac{\delta(1-\mu)}{1-\delta\mu})$$

and show that

$$\begin{aligned}
V(x^A, x^B) &= \max_{a^A,a^B}\mu\left[U(a^A, x^A) + \delta V(\alpha x^A + \beta a^A, x^B)\right] \\
&\quad + (1-\mu)\left[U(a^B, x^B) + \delta V(x^A, \alpha x^B + \beta a^B)\right]
\end{aligned}$$

$$\begin{aligned}
V(x^A, x^B) &= \frac{\mu}{1-\delta(1-\mu)}V^{SE}(x^A|\delta = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu}V^{SE}(x^B|\delta = \frac{\delta(1-\mu)}{1-\delta\mu}) \\
&= V^*(x^A|\widetilde{\mu} = \mu) + V^*(x^B|\widetilde{\mu} = 1-\mu)
\end{aligned}$$

$$
\begin{aligned}
&= \max_{a^A \in [0, \frac{y}{p}]} \mu \left[U(a^A, x^A) + \delta V^*(\alpha x^A + \beta a^A | \mu) + \delta V^*(x^B | 1 - \mu) \right] \\
&\quad + \max_{a^B \in [0, \frac{y}{p}]} (1 - \mu) \left[U(a^B, x^B) + \delta V^*(\alpha x^B + \beta a^B | 1 - \mu) + \delta V^*(x^A | \mu) \right] \\
&= \max_{a^A \in [0, \frac{y}{p}]} \mu \left[U(a^A, x^A) + \delta V(\alpha x^A + \beta a^A, x^B) \right] \\
&\quad + \max_{a^B \in [0, \frac{y}{p}]} (1 - \mu) \left[U(a^B, x^B) + \delta V(x^A, \alpha x^B + \beta a^B) \right]
\end{aligned}
$$

∎

Proof of Proposition 3: By our assumptions on $u(\cdot,\cdot)$, $v(\cdot,\cdot)$ and $\widetilde{v}(\cdot,\cdot)$, marginal utility (with the perceived tolerance function) with respect to the consumption stock, which is given by

$$
\begin{aligned}
\frac{\partial}{\partial x} \quad & \left[w(\frac{1}{\beta}(x' - \alpha x)) + \widetilde{v}(\frac{1}{\beta}(x' - \alpha x), x) \right] = \\
& -\frac{\alpha}{\beta} w'(\frac{1}{\beta}(x' - \alpha x)) - \frac{\alpha}{\beta} \widetilde{v}_1(\frac{1}{\beta}(x' - \alpha x), x) + \widetilde{v}_2(\frac{1}{\beta}(x' - \alpha x), x)
\end{aligned}
$$

and marginal utility (with the true tolerance function) with respect to the consumption stock, which is given by

$$
\begin{aligned}
\frac{\partial}{\partial x} \quad & \left[w(\frac{1}{\beta}(x' - \alpha x)) + v(\frac{1}{\beta}(x' - \alpha x), x) \right] = \\
& -\frac{\alpha}{\beta} w'(\frac{1}{\beta}(x' - \alpha x)) - \frac{\alpha}{\beta} v_1(\frac{1}{\beta}(x' - \alpha x), x) + v_2(\frac{1}{\beta}(x' - \alpha x), x)
\end{aligned}
$$

are both strictly increasing x'. The assumptions on the utility and tolerance functions satisfy Assumptions 1-4 of Orphanides and Zervos (1994), and therefore their Lemmas 2 and 3 and Proposition 1 apply to the policy correspondences $\widetilde{\psi}^A(x)$, $\widetilde{\psi}^B(x)$, $\psi^A(x)$ and $\psi^B(x)$. ∎

Proof of Proposition 4: Begin by proving the following Lemma

Lemma 8 *If $V^*()$ is the solution to*

$$V^*(x) = \max_a \mu \left[\widetilde{U}^{jt}(a,x) + \delta V^*(\alpha x + \beta a) \right] + (1-\mu) V^*(x)$$

then

$$V^*(x) = \frac{\mu}{1-\delta(1-\mu)} \widetilde{W}^{jt} \left(x | \delta = \frac{\delta\mu}{1-\delta(1-\mu)} \right)$$

for $j = A, B$.

Proof:

$$\begin{aligned}
V^*(x) &= \frac{\mu}{1-\delta(1-\mu)} \widetilde{W}^{jt} \left(x | \delta = \frac{\delta\mu}{1-\delta(1-\mu)} \right) \\
&= \max_{a\in[0,\frac{y}{p}]} \frac{\mu}{1-\delta(1-\mu)} \left[\widetilde{U}^{jt}(a,x) + \frac{\delta\mu}{1-\delta(1-\mu)} \widetilde{W}^{jt}(\alpha x + \beta a | \frac{\delta\mu}{1-\delta(1-\mu)}) \right] \\
&= \max_{a\in[0,\frac{y}{p}]} \frac{\mu}{1-\delta(1-\mu)} \left[\widetilde{U}^{jt}(a,x) + \delta V^*(\alpha x + \beta a) \right] \\
&= \max_{a\in[0,\frac{y}{p}]} \mu \left[\widetilde{U}^{jt}(a,x) + \delta V^*(\alpha x + \beta a) \right] + (1-\mu)\delta V^*(x)
\end{aligned}$$

∎

Now, let

$$\widetilde{V}^t(x^A,x^B) = \frac{\mu}{1-\delta(1-\mu)} \widetilde{W}^{At}(x^A | \delta = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu} \widetilde{W}^{Bt}(x^B | \delta = \frac{\delta(1-\mu)}{1-\delta\mu})$$

and show that

$$\begin{aligned}
\widetilde{V}^t(x^A,x^B) &= \max_{a^A,a^B} \mu \left[\widetilde{U}^{At}(a^A,x^A) + \delta \widetilde{V}^t(\alpha x^A + \beta a^A, x^B) \right] \\
&\quad + (1-\mu) \left[\widetilde{U}^{Bt}(a^B,x^B) + \delta \widetilde{V}^t(x^A, \alpha x^A + \beta a^B) \right]
\end{aligned}$$

$$\begin{aligned}
V(x^A,x^B) &= \frac{\mu}{1-\delta(1-\mu)} \widetilde{W}^{At}(x^A | \delta = \frac{\delta\mu}{1-\delta(1-\mu)}) + \frac{1-\mu}{1-\delta\mu} \widetilde{W}^{Bt}(x^B | \delta = \frac{\delta(1-\mu)}{1-\delta\mu}) \\
&= V^*(x^A | \widetilde{\mu} = \mu) + V^*(x^B | \widetilde{\mu} = 1-\mu) \\
&= \max_{a^A\in[0,\frac{y}{p}]} \mu \left[\widetilde{U}^{At}(a^A,x^A) + \delta V^*(\alpha x^A + \beta a^A | \mu) + \delta V^*(x^B | 1-\mu) \right] \\
&\quad + \max_{a^B\in[0,\frac{y}{p}]} (1-\mu) \left[\widetilde{U}^{Bt}(a^B,x^B) + \delta V^*(\alpha x^B + \beta a^B | 1-\mu) + \delta V^*(x^A | \mu) \right]
\end{aligned}$$

$$= \max_{a^A \in [0,\frac{y}{p}]} \mu \left[\widetilde{U}^{At}(a^A, x^A) + \delta V(\alpha x^A + \beta a^A, x^B)\right]$$
$$+ \max_{a^B \in [0,\frac{y}{p}]} (1-\mu) \left[\widetilde{U}^{Bt}(a^B, x^B) + \delta V(x^A, \alpha x^B + \beta a^B)\right]$$

∎

Proof of Proposition 5: By our assumptions on $u(\cdot,\cdot)$ and $\widetilde{v}^{jt}(\cdot,\cdot)$, perceived marginal utility in environment A with respect to the consumption stock, which is given by

$$\frac{\partial}{\partial x} \quad \left[w(\frac{1}{\beta}(x'-\alpha x)) + \widetilde{v}^{At}(\frac{1}{\beta}(x'-\alpha x), x)\right] =$$
$$-\frac{\alpha}{\beta}w'(\frac{1}{\beta}(x'-\alpha x)) - \frac{\alpha}{\beta}\widetilde{v}_1^{At}(\frac{1}{\beta}(x'-\alpha x), x) + \widetilde{v}_2^{At}(\frac{1}{\beta}(x'-\alpha x), x)$$

and perceived marginal utility in environment B with respect to the consumption stock, which is given by

$$\frac{\partial}{\partial x} \quad \left[w(\frac{1}{\beta}(x'-\alpha x)) + \widetilde{v}^{Bt}(\frac{1}{\beta}(x'-\alpha x), x)\right] =$$
$$-\frac{\alpha}{\beta}w'(\frac{1}{\beta}(x'-\alpha x)) - \frac{\alpha}{\beta}\widetilde{v}_1^{Bt}(\frac{1}{\beta}(x'-\alpha x), x) + \widetilde{v}_2^{Bt}(\frac{1}{\beta}(x'-\alpha x), x)$$

are both strictly increasing x' for all t. The assumptions on the utility and tolerance functions satisfy Assumptions 1-4 of Orphanides and Zervos (1994), and therefore their Lemmas 2 and 3 and Proposition 1 apply. ∎

Proof of Proposition 6: Begin by proving the following Lemma

Lemma 9 *If $V^*()$ is the solution to*

$$V^*(x,\widetilde{\gamma}) = \max_a \mu \left[U(a,x,\widetilde{\gamma}) + \delta V^*(\alpha x + \beta a, g(\widetilde{\gamma},\gamma))\right] + (1-\mu) V^*(x,\widetilde{\gamma})$$

then

$$V^*(x,\widetilde{\gamma}) = \frac{\mu}{1-\delta(1-\mu)} W\left(x,\widetilde{\gamma}|\delta = \frac{\delta\mu}{1-\delta(1-\mu)}\right)$$

Proof:

$$
\begin{aligned}
&V^*(x,\widetilde{\gamma})\\
=\;& \frac{\mu}{1-\delta(1-\mu)} W\left(x,\widetilde{\gamma}|\frac{\delta\mu}{1-\delta(1-\mu)}\right)\\
=\;& \max_{a\in[0,\frac{y}{p}]} \frac{\mu}{1-\delta(1-\mu)}\left[U(a,x,\widetilde{\gamma})+\frac{\delta\mu}{1-\delta(1-\mu)}W(\alpha x+\beta a, g\left(\widetilde{\gamma},\gamma\right)|\frac{\delta\mu}{1-\delta(1-\mu)})\right]\\
=\;& \max_{a\in[0,\frac{y}{p}]} \frac{\mu}{1-\delta(1-\mu)}\left[U(a,x,\widetilde{\gamma})+\delta V^*(\alpha x+\beta a, g\left(\widetilde{\gamma},\gamma\right))\right]\\
=\;& \max_{a\in[0,\frac{y}{p}]} \mu\left[U(a,x,\widetilde{\gamma})+\delta V^*(\alpha x+\beta a, g\left(\widetilde{\gamma},\gamma\right))\right]+(1-\mu)\delta V^*(x)
\end{aligned}
$$

∎

Now, let

$$
V(x^A,x^B,\widetilde{\gamma}^A,\widetilde{\gamma}^B)\frac{\mu}{1-\delta(1-\mu)}W(x^A,\widetilde{\gamma}^A|\delta=\frac{\delta\mu}{1-\delta(1-\mu)})+\frac{1-\mu}{1-\delta\mu}W(x^B,\widetilde{\gamma}^B|\delta=\frac{\delta(1-\mu)}{1-\delta\mu})
$$

and show that

$$
\begin{aligned}
V(x^A,x^B,\widetilde{\gamma}^A,\widetilde{\gamma}^B) \;=\;& \max_{a^A,a^B}\mu\left[U(a^A,x^A,\widetilde{\gamma}^A)+\delta V(\alpha x^A+\beta a^A,x^B,g(\widetilde{\gamma}^A,\gamma),\widetilde{\gamma}^B)\right]\\
&+(1-\mu)\left[U(a^B,x^B,\widetilde{\gamma}^B)+\delta V(x^A,\alpha x^B+\beta a^B,\widetilde{\gamma}^A,g(\widetilde{\gamma}^B,\gamma))\right]
\end{aligned}
$$

$$
\begin{aligned}
&V(x^A,x^B,\widetilde{\gamma}^A,\widetilde{\gamma}^B)\\
=\;& \frac{\mu}{1-\delta(1-\mu)}W(x^A,\widetilde{\gamma}^A|\frac{\delta\mu}{1-\delta(1-\mu)})+\frac{1-\mu}{1-\delta\mu}W(x^B,\widetilde{\gamma}^B|\frac{\delta(1-\mu)}{1-\delta\mu})\\
=\;& V^*(x^A,\widetilde{\gamma}^A|\widetilde{\mu}=\mu)+V^*(x^B,\widetilde{\gamma}^B|\widetilde{\mu}=1-\mu)\\
=\;& \max_{a^A\in[0,\frac{y}{p}]}\mu[U(a^A,x^A,\widetilde{\gamma}^A)+\delta V^*(\alpha x^A+\beta a^A,g(\widetilde{\gamma}^A,\gamma)|\mu)+\delta V^*(x^B,\widetilde{\gamma}^B|1-\mu)]\\
&+\max_{a^B\in[0,\frac{y}{p}]}(1-\mu)\left[U(a^B,x^B,\widetilde{\gamma}^B)+\delta V^*(\alpha x^B+\beta a^B,g(\widetilde{\gamma}^B,\gamma)|1-\mu)+\delta V^*(x^A,\widetilde{\gamma}^A|\mu)\right]
\end{aligned}
$$

$$= \max_{a^A \in [0, \frac{y}{p}]} \mu \left[U(a^A, x^A, \widetilde{\gamma}^A) + \delta V(\alpha x^A + \beta a^A, x^B, g(\widetilde{\gamma}^A, \gamma), \widetilde{\gamma}^B) \right]$$

$$+ \max_{a^B \in [0, \frac{y}{p}]} (1 - \mu) \left[U(a^B, x^B, \widetilde{\gamma}^B) + \delta V(x^A, \alpha x^B + \beta a^B, \widetilde{\gamma}^A, g(\widetilde{\gamma}^B, \gamma)) \right]$$

$\blacksquare$

Appendix B Derivation of Equations (3.7) and (3.8)

I'll begin with the derivation of equation (3.7)

$$\widehat{\beta} = \frac{\sigma_{zz}(1)\sigma_P^2 - \sigma_{Pz}\sigma_{Pz}(1)}{\sigma_z^2\sigma_P^2 - \sigma_{Pz}^2(1)}$$

Recall that the model that is estimated is

$$z_t = \widehat{\beta} z_{t+1} + \widehat{\gamma} P_t$$

In the limit, the least squares estimates are given by:

$$\begin{aligned}\sum_t \left(-z_t + \widehat{\beta} z_{t+1} + \widehat{\gamma} P_t\right) z_{t+1} &= 0 \\ \sum_t \left(-z_t + \widehat{\beta} z_{t+1} + \widehat{\gamma} P_t\right) P_t &= 0\end{aligned}$$

or

$$\begin{aligned}\sum_t z_t z_{t+1} &= \widehat{\beta} \sum z_{t+1}^2 + \widehat{\gamma} \sum P_t z_{t+1} \\ \sum_t z_t P_t &= \widehat{\beta} \sum z_{t+1} P_t + \widehat{\gamma} \sum P_t^2\end{aligned}$$

Therefore, $\widehat{\beta}$ is given by

$$\widehat{\beta} = \frac{\sum z_t z_{t+1} \sum P_t^2 - \sum z_t P_t \sum P_t z_{t+1}}{\sum z_{t+1}^2 \sum P_t^2 - \left(\sum z_{t+1} P_t\right)^2}$$

Taking limits, we arrive at:

$$\widehat{\beta} = \frac{\sigma_{zz}(1)\sigma_P^2 - \sigma_{Pz}\sigma_{Pz}(1)}{\sigma_z^2\sigma_P^2 - \sigma_{Pz}^2(1)}$$

Now, to derive equation (3.8), suppose that the true coefficient on a_{t+1} is not 0, that is $\beta \neq 0$.

β is given by

$$\begin{aligned} \sum a_t a_{t+1} - \beta \sum a_{t+1}^2 - \gamma \sum a_{t+1} P_t &= 0 \\ \sum a_t P_t - \beta \sum a_{t+1} P_t - \gamma \sum P_t^2 &= 0 \end{aligned}$$

or

$$\beta = \frac{\sum a_t a_{t+1} \sum P_t^2 - \sum a_t P_t \sum a_{t+1} P_t}{\sum a_{t+1}^2 \sum P_t^2 - \left(\sum a_{t+1} P_t\right)^2}$$

so that in the limit:

$$\beta = \frac{\sigma_{aa}(1)\sigma_P^2 - \sigma_{aP}\sigma_{Pa}(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1)}$$

Now consider $\widehat{\beta}$

$$\begin{aligned} \widehat{\beta} &= \frac{\sigma_{zz}(1)\sigma_P^2 - \sigma_{Pz}\sigma_{Pz}(1)}{\sigma_z^2\sigma_P^2 - \sigma_{Pz}^2(1)} \\ &= \frac{\sigma_P^2\left(\sigma_{aa}(1) - \sigma_{\eta\eta}(1)\right) - \sigma_{Pa}\sigma_{Pa}(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2} \\ &= \frac{\sigma_P^2\sigma_{aa}(1) - \sigma_{Pa}\sigma_{Pa}(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2} - \frac{\sigma_P^2\sigma_{\eta\eta}(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2} \\ &= \frac{\beta}{1 + \frac{\sigma_\eta^2\sigma_P^2}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1)}} - \frac{\sigma_P^2\sigma_{\eta\eta}(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2} \end{aligned}$$

which gives us equation (3.8):

$$\widehat{\beta} = \beta\frac{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2} - \frac{\sigma_P^2\sigma_{\eta\eta}(1)}{\sigma_a^2\sigma_P^2 - \sigma_{Pa}^2(1) + \sigma_\eta^2\sigma_P^2}$$