Addiction and Cue-Triggered Decision Processes

By B. Douglas Bernheim and Antonio Rangel*

We propose a model of addiction based on three premises: (i) use among addicts is frequently a mistake; (ii) experience sensitizes an individual to environmental cues that trigger mistaken usage; (iii) addicts understand and manage their susceptibilities. We argue that these premises find support in evidence from psychology, neuroscience, and clinical practice. The model is tractable and generates a plausible mapping between behavior and the characteristics of the user, substance, and environment. It accounts for a number of important patterns associated with addiction, gives rise to a clear welfare standard, and has novel implications for policy. (*JEL* D01, D11, H20, H21, H23, H31, I12, I18, K32)

According to clinical definitions, substance addiction occurs when, after significant exposure, users find themselves engaging in compulsive, repeated, and unwanted use despite clearly harmful consequences, and often despite a strong desire to quit unconditionally (see, e.g., the American Psychiatric Association's *Diagnostic and Statistical Manual of Mental Disorders,* known as DSM-IV). There is widespread agreement that certain substances have addictive properties,[1] and there is some debate as to whether formal definitions of addiction should be expanded to include other substances (such as fats and sugars) and activities (such as shopping, shoplifting, sex, television viewing, and internet use).

The consumption of addictive substances raises important social issues affecting members of all socioeconomic strata.[2] Tens of millions of Americans use addictive substances. Nearly 25 million adults have a history of alcohol dependence, and more than five million qualify as "hard-core" chronic drug users. Estimates for 1999 place total U.S. expenditures on tobacco products, alcoholic beverages, cocaine, heroin, marijuana, and methamphetamines at more than $150 billion, with still more spent on caffeine and addictive prescription drugs. Estimated social costs (health care, impaired productivity, crime, and so forth) total more than

[1] Eliot Gardner and James David (1999) provide the following list of 11 addictive substances: alcohol, barbiturates, amphetamines, cocaine, caffeine and related methylxanthine stimulants, cannabis, hallucinogenics, nicotine, opioids, dissociative anasthetics, and volatile solvents.

[2] The statistics in this paragraph were obtained from the following sources: Center for Disease Control (1993), National Institute on Drug Abuse (1998), National Institute on Alcohol Abuse and Alcoholism (2001), Office of National Drug Control Policy (2001a, b), and U.S. Census Bureau (2001). There is, of course, disagreement as to many of the reported figures.

$300 billion per year. On average over 500,000 deaths each year are attributed directly to cigarettes and alcohol.

Public policies regarding addictive substances run the gamut from laissez-faire to taxation, subsidization (e.g., of rehabilitation programs), regulated dispensation, criminalization, product liability, and public health campaigns. Each alternative policy approach has passionate advocates and detractors. Economic analysis can potentially inform this debate, but it requires a sound theory of addiction.

This paper presents a new theory of addiction based on three central premises: first, use among addicts is frequently a mistake; second, experience with an addictive substance sensitizes an individual to environmental cues that trigger mistaken usage; third, addicts understand their susceptibility to cue-triggered mistakes and attempt to manage the process with some degree of sophistication. We argue that these premises find strong support in evidence from psychology, neuroscience, and clinical practice. In particular, research has shown that addictive substances systematically interfere with the proper operation of an important class of processes which the brain uses to forecast near-term hedonic rewards (pleasure), and this leads to strong, misguided, cue-conditioned impulses that often defeat higher cognitive control.

We provide a parsimonious representation of this phenomenon in an otherwise standard model of intertemporal decision-making. Specifically, we allow for the possibility that, upon exposure to environmental cues, the individual may enter a "hot" decision-making mode in which he always consumes the substance irrespective of underlying preferences, and we assume that sensitivity to cues is related to past experiences. The individual may also operate in a "cold" mode, wherein he considers all alternatives and contemplates all consequences, including the effects of current choices on the likelihood of entering the hot mode in the future.[3]

As a matter of formal mathematics, our model involves a small departure from the standard framework. Behavior corresponds to the solution of a dynamic programming problem with stochastic state-dependent mistakes. Our approach therefore harmonizes economic theory with evidence on the biological foundations of addiction without sacrificing analytic tractability. We underscore this point by providing results that illuminate the relationships between behavior and the characteristics of the user, substance, and environment. For example, we find that, when one substance is more addictive than another, then *ceteris paribus* the more addictive substance is associated with less consumption among relatively new users, but with more consumption (both intentional and accidental) among highly experienced users.

The theory can account for a number of important patterns associated with addiction. It also gives rise to a clear welfare standard and has novel implications for public policy. Our policy analysis focuses on consumer welfare and therefore ignores supply-side effects and externalities. It emphasizes the role of policy in averting mistakes and in either ameliorating or magnifying significant, uninsurable monetary risks indirectly caused by exposure to stochastic environmental cues. We show that a beneficial policy intervention potentially exists if and only if there are circumstances in which users unsuccessfully attempt to abstain. In that case, the optimal policy depends on usage patterns. In a natural benchmark case, it is optimal to subsidize an addictive substance when the likelihood of use rises with the level of past experience. In contrast, provided the substance is sufficiently inexpensive, it is optimal to tax the substance when the likelihood of use declines with the level of past experience. Under weak conditions, a small subsidy for rehabilitation is beneficial, and a small tax is harmful. When substance taxation is optimal, under some conditions criminalization can perform even better. Programs that make addictive substances available on a prescription basis have potentially large benefits. Restrictions on advertising and public consumption, and

---

[3] Our analysis is related to work by George Loewenstein (1996, 1999), who considers simple models in which an individual can operate either in a hot or cold decision-making mode. Notably, Loewenstein assumes that behavior in the hot mode reflects the application of a "false" utility function, rather than a breakdown of the processes by which

a utility function is maximized. He also argues, contrary to the findings of this paper, that imperfect self-understanding is necessary for addiction-like behaviors.

statutes requiring counter-cues on packaging, are also potentially beneficial.

The remainder of the paper is organized as follows. Section I describes some important behavioral patterns associated with addiction that require explanation. Section II lays out and justifies, with particular reference to evidence from psychology and neuroscience, the central premises of our theory. Section III presents the formal model. Section IV explores the model's positive implications, including its ability to generate observed behavioral patterns. Section V concerns policy analysis. Section VI clarifies the relationships between our theory of addiction and others that appear in the literature, including the standard model of rational addiction (Gary Becker and Kevin Murphy, 1988), various extensions of this model (Athanasios Orphanides and David Zervos, 1995; Angela Hung, 2000; David Laibson, 2001), and a number of behavioral alternatives (Loewenstein, 1996, 1999; Ted O'Donoghue and Matthew Rabin, 1999, 2000; Jonathan Gruber and Botond Koszegi, 2001; Loewenstein et al., Faruk Gul and Wolfgang Pesendorfer, 2001a, b). Section VII concludes and discusses directions for future research. The appendices provide additional technical details and proofs; in some cases we sketch proofs to conserve space.

I. Patterns of Addictive Behavior

What makes addiction a distinctive phenomenon? From the extensive body of research on addiction in neuroscience, psychology, and clinical practice, we have distilled five important behavioral patterns requiring explanation.

1. *Unsuccessful Attempts to Quit.*—Addicts often express a desire to stop using a substance permanently and unconditionally but are unable to follow through. Short-term abstention is common while long-term recidivism rates are high. For example, during 2000, 70 percent of current smokers expressed a desire to quit completely, and 41 percent stopped smoking for at least one day in an attempt to quit, but only 4.7 percent successfully abstained for more than three months (see J. E. Harris, 1993; Y. I. Hser et al., 1993; C. O'Brien, 1997; A. Goldstein, 2001; A. Trosclair et al., 2002). This pattern is particularly striking because regular users initially experience painful withdrawal symptoms when they attempt to quit, and these symptoms decline over time with successful abstention. Thus, recidivism often occurs after users have borne the most significant costs of quitting, sometimes following years of determined abstention.

2. *Cue-Triggered Recidivism.*—Recidivism rates are especially high when addicts are exposed to cues related to past drug consumption. Long-term usage is considerably lower among those who experience significant changes of environment (see O'Brien, 1975, 1997; Goldstein and H. Kalant, 1990; Hser et al., 1993, 2001; Goldstein, 2001).[4] Treatment programs often advise recovering addicts to move to new locations and to avoid the places where previous consumption took place. Stress and "priming" (exposure to a small taste of the substance) have also been shown to trigger recidivism (see Goldstein, 2001; Terry Robinson and Kent Berridge, 2003).

3. *Self-Described Mistakes.*—Addicts often describe past use as a mistake in a very strong sense: they think that they would have been better off *in the past as well as the present* had they acted differently. They recognize that they are likely to make similar errors in the future, and that this will undermine their desire to abstain. When they succumb to cravings, they sometimes characterize choices as mistakes *even while in the act of consumption.*[5] It is instructive that the 12-step program of Alcoholic Anonymous begins: "We admit we are powerless over alcohol—that our lives have become unmanageable."

---

[4] L. Robins (1974) and Robins et al. (1974) found that Vietnam veterans who were addicted to heroin and/or opium at the end of the war experienced much lower relapse rates than other young male addicts during the same period. A plausible explanation is that veterans encountered fewer environmental triggers (familiar circumstances associated with drug use) upon returning to the United States.

[5] Goldstein (2001, p. 249) describes this phenomenon as follows: the addict had been "suddenly overwhelmed by an irresistible craving, and he had rushed out of his house to find some heroin. ... it was as though he were driven by some external force he was powerless to resist, *even though he knew while it was happening that it was a disastrous course of action for him*" (italics added).

4. *Self-Control through Precommitment.*—Recovering users often manage their tendency to make mistakes by voluntarily removing or degrading future options. They voluntarily admit themselves into "lock-up" rehabilitation facilities, often not to avoid cravings, but precisely because they expect to experience cravings and wish to control their actions. They also consume medications that either generate unpleasant side effects, or reduce pleasurable sensations, if the substance is subsequently consumed.[6] Severe addicts sometimes enlist others to assist with physical confinement to assure abstinence through the withdrawal process.

5. *Self-Control through Behavioral and Cognitive Therapy.*—Recovering addicts attempt to minimize the probability of relapse through behavioral and cognitive therapies. Successful behavioral therapies teach cue-avoidance, often by encouraging the adoption of new lifestyles and the development of new interests. Successful cognitive therapies teach cue-management, which entails refocusing attention on alternative consequences and objectives, often with the assistance of a mentor or trusted friend or through a meditative activity such as prayer. Notably, these therapeutic strategies affect addicts' choices *without providing new information.*[7]

While consumption patterns for addictive substances are distinctive in some respects, it is important to bear in mind that they are ordinary in other respects. A number of studies have shown that aggregate drug use responds both to prices and to information about the effects of addictive substances. For example, an aggressive U.S. public health campaign is widely credited with reduction in smoking rates. There is also evidence that users engage in sophisticated forward-looking deliberation, reducing current consumption in response to anticipated price increases.[8]

It is important to remember that consumption patterns for the typical addictive substance vary considerably from person to person.[9] Some people never use it. Some use it in a controlled way, either periodically or for a short time period. Some experience occasional episodes where they appear to "lose control" (binge) but suffer no significant ongoing impairment and have no desire to quit permanently. Some fit the DSM-IV definition of addiction. In the rest of the paper the term *addict* is reserved for the third and fourth groups, whereas the term *user* is applied to everyone.

II. Central Premises

The theory developed in this paper is based on three premises: (i) use among addicts is frequently a mistake—that is, a pathological divergence between choice and preference; (ii) experience with an addictive substance sensitizes an individual to environmental cues that trigger mistaken usage; and (iii) addicts understand their susceptibility to cue-triggered mistakes and attempt to manage the process with some degree of sophistication. The third premise is consistent with observed behavioral patterns involving cue-avoidance and/or precommitments, and should be relatively uncontroversial. In contrast, the notion that choices and preferences can diverge is contrary to the standard doctrine of revealed preference and therefore requires thorough justification.

There are plainly circumstances in which it makes no sense to infer preferences from choices. For example, American visitors to the

---

[6] Disulfiram interferes with the liver's ability to metabolize alcohol; as a result, ingestion of alcohol produces a highly unpleasant physical reaction for a period of time. Methadone, an agonist, activates the same opioid receptors as heroin, and thus produces a mild high, but has a slow onset and a long-lasting effect, and it reduces the high produced by heroin. Naltrexone, an antagonist, blocks specific brain receptors and thereby diminishes the high produced by opioids. All of these treatments reduce the frequency of relapse (see O'Brien, 1997; Goldstein, 2001).

[7] Goldstein (2001, p. 149) reports that there is a shared impression among the professional community that 12-step programs such as AA "are effective for many (if not most) alcohol addicts." However, given the nature of these programs, objective performance tests are not available. The AA treatment philosophy is based on "keeping it simple by putting the focus on not drinking, on attending meetings, and on reaching out to other alcoholics."

[8] See F. Chaloupka and K. Warner (2001), Gruber and Koszegi (2001), and R. MacCoun and P. Reuter (2001) for a review of the evidence.

[9] Even for a substance such as cocaine, which is considered highly addictive, only 15–16 percent of people become addicted within 10 years of first use (F. A. Wagner and J. C. Anthony, 2002).

United Kingdom suffer numerous injuries and fatalities because they often look only to the left before stepping into streets, even though they know traffic approaches from the right. One cannot reasonably attribute this to the pleasure of looking left or to masochistic preferences. The pedestrian's objectives—to cross the street safely—are clear, and the decision is plainly a mistake. The source of this systematic error is traceable to features of the human brain. Habituated, semi-automatic responses beneficially increase the speed of decision-making in some circumstances but lead to systematic mistakes in others.

Recent research on the neuroscience of addiction has identified specific features of the brain that appear to produce systematic errors with respect to decisions involving the consumption of addictive substances. The key process involves a mechanism (henceforth called the "hedonic forecasting mechanism" or HFM) that is responsible for associating environmental cues with forecasts of short-term hedonic (pleasure/pain) responses.[10] Normally, the HFM learns through feedback from the hedonic system: with experience, it associates a situation and action with an anticipatory biochemical response, the magnitude of which reflects the intensity of expected pleasure. Addictive substances interfere with the normal operation of the HFM by acting directly (i.e., independent of the pleasure experienced) on the learning process that teaches the HFM to generate the anticipatory response. With repeated use of a substance, cues associated with past consumption cause the HFM to forecast grossly exaggerated pleasure responses, creating a powerful (and disproportionate) impulse to use. When this happens, a portion of the user's decision processes functions as if it has systematically skewed information, which leads to mistakes in decision-making.

Next we describe some of the key evidence that leads to these conclusions. We organize our discussion around four points.

1. Brain Processes Include a Hedonic Forecasting Mechanism (HFM) Which, with Experience, Produces a Biochemical Response to

Situations and Opportunities, the Magnitude of Which Constitutes a Forecast of Near-Term Pleasure.—Neuroscientists have long recognized that the mesolimbic dopamine system (MDS) is a basic component of human decision processes.[11] A large body of recent research indicates that the MDS functions, at least in part, as an HFM. In a series of experiments, subjects (often monkeys) are presented with a cue that is associated with a reward delivered a few seconds later (see Wolfram Schultz et al., 1997; Schultz, 1998, 2000). Initially, the MDS fires in response to the delivery of the reward and not in response to the cue. However, as time passes, the MDS fires with the presentation of the cue and not with the delivery of the reward. Moreover, the level of cue-triggered MDS activity is proportional to the size of the eventual reward. If, after a number of trials, the experimenter increases the magnitude of the reward, the MDS fires twice: with the presentation of the cue (at a level proportional to the original anticipated reward), and with the delivery of the reward (at a level reflecting the difference between the anticipated and actual rewards). After repeated trials with the new reward, the MDS fires more intensely upon presentation of the cue and, once again, does not respond to the delivery of the reward. Thus, with experience, the MDS generates a cue-conditioned dopamine response that anticipates the magnitude of the eventual reward.

2. Activation of the HFM Does Not Necessarily Create Hedonic Sensation, and Hedonic Sensation Can be Experienced without HFM Activation.—Since the MDS produces a dopamine response prior to an anticipated experience and no response during the experience, it is natural to conjecture that this mechanism is neither a source nor a manifestation of pleasure.

---

[10] The phrase "hedonic forecasting mechanism" summarizes the role of this process in economic terms; this terminology is not used in the existing behavioral neuroscience literature.

[11] The MDS originates in the ventral tegmental area, near the base of the brain, and sends projections to multiple regions of the frontal lobe, especially to the nucleus accumbens. The MDS also connects with the amygdala, basal forebrain, and other areas of the prefrontal lobe. These connections are believed to serve as an interface between the MDS and attentional, learning, and cognitive processes (Robinson and Berridge, 2003).

Indeed, the human brain appears to contain a separate hedonic system that is responsible for producing sensations of "well-being."[12] In a series of papers, neuroscientists Kent Berridge and Terry Robinson have argued that two separate processes are at work in decision making: a "wanting" process, which encompasses the impulse created by a positive MDS forecast, and a "liking" process, which refers to a hedonic response (see Robinson and Berridge, 1993, 2000, 2003; Berridge, 1996, 1999; Berridge and Robinson, 1998, 2003).[13] Their hypothesis emerges from numerous experimental studies, including the following. Using measures of "liking" based on rats' facial expressions when responding to sweet and sour tastes, several experiments have shown that neither the direct activation of the MDS nor its suppression affects liking (S. Pecina et al., 1997; H. J. Kaczmarek and S. W. Kiefer, 2000; C. L. Wyvell and Berridge, 2000). Others have demonstrated that the "liking" system functions well even with massive lesions to the MDS (see Berridge and Robinson, 1998). Direct activation of the MDS through microinjections of amphetamine in the nucleus accumbens (NAc) increases wanting but fails to increase liking (Wyvell and Berridge, 2000). Finally, blocking the MDS with dopamine antagonists does not have an impact on the level of pleasure obtained from using a drug reported by amphetamine and nicotine users (L. H. Brauer et al., 1997, 2001; S. R. Wachtel et al., 2002).

3. *HFM-Generated Forecasts Influence Choices.*—A series of classic experiments by J. Olds and P. Milner (1954) demonstrated that rats learn to return to locations where they have received direct electrical stimulation to the MDS. When provided with opportunities to self-administer by pressing a lever, the rats rapidly became addicted, giving themselves approximately 5,000–10,000 "hits" during each one-hour daily session, ignoring food, water, and opportunities to mate. These rats are willing to endure painful electric shocks to reach the lever (see Gardner and David, 1999 for a summary of these experiments). Complementary evidence shows that rats given drugs that block dopamine receptors, thereby impeding the appropriate operation of the MDS, eventually stop feeding (Berridge, 1999).

Notably, the MDS activates "seeking behaviors" as well as immediate consumption choices. That is, it learns to make associations not just between consumption opportunities and hedonic payoffs, but also between environmental cues and activities that tend to produce these consumption opportunities. For example, the sight of food may create a powerful impulse to eat, while an odor may create a powerful impulse to seek food. The size of the set of environmental cues that trigger an associated seeking behavior increases with the strength of the hedonic forecast (see Robinson and Berridge, 1993, 2000, 2003; Berridge and Robinson, 1998, 2003).

While the MDS plays a key role in determining choices, it is not the only process at work. In an organism with a sufficiently developed frontal cortex, higher cognitive mechanisms can override HFM-generated impulses. Though the specific mechanisms are not yet fully understood, structures in the frontal cortex appear to activate competing "cognitive incentives" (Berridge and Robinson, 2003), for example, by identifying alternative courses of action or projecting the future consequences of choices. The outcome depends on the intensity of the HFM forecast and on the ability of the frontal cortex to engage the necessary cognitive operations.[14] Thus, a more attractive HFM-generated forecast makes cognitive override less likely. In addition, the MDS seems to affect which stimuli the brain attends to, which cognitive operations it

[12] The existing evidence suggests that the hedonic system is modulated by a distributed network, separate from the structures involved in the HFM, that includes GABAergic neurons in the shell of the NAc, the ventral palladium, and the brainstem parachial nucleus (see Berridge and Robinson, 2003).

[13] For decades, neuroscientists and psychologists have used the term "reward" to describe both liking and wanting. In most experimental settings, the distinction is immaterial since outcomes that are liked are also wanted, and vice versa. However, as we will see, this distinction is critical to understanding why repeated exposure to drugs leads to mistaken usage.

[14] The activation of the cognitive operations required for cognitive control depends on neocortical structures such as the insula and the orbitofrontal cortex (see e.g., J. D. Cohen and K. I. Blum, 2002; D. C. Krawczyk, 2002; E. T. Rolls, 2002; Masataka Watanabe et al., 2002).

activates (what it thinks about), and which memories it preserves, and this may make it more difficult to engage the cognitive operations required to override the HFM.[15]

We emphasize that the HFM and higher cognitive processes are not two different sets of "preferences" or "selves" competing for control of decisions. Hedonic experiences are generated separately, and an individual maximizes the quality of these experiences by appropriately deploying both forecasting processes to anticipate outcomes. The HFM's main advantage is that it can produce rapid decisions with generally beneficial near-term outcomes, provided the environment is stable. It cannot, however, anticipate sufficiently delayed consequences, and when the environment changes, it can neither ignore irrelevant past experiences nor adjust forecasts prior to acquiring further experience. The competing cognitive forecasting system addresses these shortcomings (albeit imperfectly) but is comparatively slow. Balanced competition between these two processes apparently emerged as evolution's best compromise.

4. *Addictive Substances Act Directly on the HFM, Disrupting Its Ability to Construct Accurate Hedonic Forecasts and Exaggerating the Anticipated Hedonic Benefits of Consumption.*—Although addictive substances differ considerably in their chemical and psychological properties, there is a large and growing consensus in neuroscience that they share an ability to activate the firing of dopamine into the NAc with much greater intensity and persistence than other substances. They do this either by activating the MDS directly, or by activating other networks that have a similar effect on the NAc (see Ingrid Wickelgren, 1997; Steven Hyman and Robert Malenka, 2001; E. J. Nestler, 2001; Robinson and Berridge, 2003; Nestler and Robert Malenka, 2004).[16]

For nonaddictive substances, the MDS learns to assign a hedonic forecast that bears some normal relation to the subsequent hedonic experience. For addictive substances, consumption activates dopamine firing directly, so the MDS learns to assign a hedonic forecast that is out of proportion to the subsequent hedonic experience. This not only creates a strong (and misleading) impulse to seek and use the substance, but also undermines the potential for cognitive override.[17] Cognitive override still occurs, but in a limited range of circumstances.[18,19]

Our central premises have two implications that are worth emphasizing because they are at odds with some of the alternative models of addiction discussed in Section VI. First, the processes that produce systematic mistakes are triggered by stochastic environmental cues and are not always operative. Second, cue-triggered

---

hallucinogens are properly classified as addictive substances (see Goldstein, 2001, Ch. 14). Notably, laboratory animals and humans learn to self-administer the same set of substances, with the possible exception of hallucinogenics (Gardner and David, 1999, pp. 97–98).

[17] A stronger MDS-generated impulse is more likely to overcome competing cognitive incentives of any given magnitude. In addition, the MDS-generated impulse may make it more difficult to engage the cognitive operations required to override the HFM. For example, recovering addicts may pay too much attention to drugs, activate and maintain thoughts about the drug too easily, and retain particularly vivid memories of the high. Consistent with this, S. Vorel et al. (2001) have shown that the stimulation of memory centers can trigger strong cravings and recidivism among rats that have previously self-administered cocaine (J. P. Berke et al., 2001, and C. Holden, 2001a, b, provide nontechnical discussions).

[18] The importance of cognitive override is evident from comparisons of rats and humans. When rats are allowed to self-administer cocaine, after a short period of exposure they begin to ignore hunger, reproductive urges, and all other drives, consuming the substance until they die (R. Pickens and W. C. Harris, 1968; E. Gardner and David, 1999). In contrast, even severely addicted humans sometimes resist cravings and abstain for long periods of time. The difference is that rats rely solely on the HFM.

[19] Several studies (see J. I. Bolla et al., 1998; J. D. Jentsch and J. R. Taylor, 1999; T. W. Robbins and B. J. Everitt, 1999; A. Behara and H. Damasio, 2002; Behara, S. Dolan, and A. Hindes, 2002) have shown that addicts share psychological disorders with patients who have damaged frontal lobes affecting functions related to cognitive control. In addition, some of these studies have argued that drug use is partly responsible for this impairment. Thus, use may increase the likelihood of subsequent use by crippling cognitive control mechanisms.

---

[15] Notably, more educated individuals are far more likely to quit smoking successfully, even though education bears little relation either to the desire to quit or to the frequency with which smokers attempt to quit (Trosclair et al., 2002).

[16] Of the addictive substances listed in footnote 1, only hallucinogenics (or psychedelics) do not appear to produce intense stimulation of the MDS. Instead, they act on a "subtype of serotonin receptor which is widely distributed in areas of the brain that process sensory inputs" (Goldstein, 2001 p. 231). There is some disagreement as to whether

mistakes are specific to narrow domains. That is, they adhere to particular activities in particular circumstances and do not reflect a general bias toward immediate gratification. Since poor cognitive control increases the likelihood of becoming addicted, it should not be surprising that the typical addict exhibits other self-control problems. However, it does not follow that a general deficit in cognitive control is necessary for addiction.

In emphasizing the effects of addictive substances on decision process, we do not mean to discount the significance of their hedonic effects. The typical user is initially drawn to an addictive substance because it produces a hedonic "high." Over time, regular use leads to hedonic and physical tolerance. That is, the drug loses its ability to produce a high unless the user abstains for a while,[20] and any attempt to discontinue the drug may have unpleasant side effects (withdrawal). Cue-conditioned "cravings" may have hedonic implications as well as non-hedonic causes (i.e., HFM-generated impulses). All of these effects are clearly important, and with one exception discussed in the next section, our model subsumes them. However, there is an emerging consensus in neuroscience and psychology that decision-process effects, rather than hedonic effects, provide the key to understanding addictive behavior (see Roy Wise, 1989; G. Di Chiara, 1999; A. E. Kelley, 1999; Robbins and Everitt, 1999; Robinson and Berridge, 2000; Hyman and Malenka, 2001; Berridge and Robinson, 2003; Nestler and Malenka, 2004).

III. The Model

We consider a decision-maker (DM) who can operate in either of two modes: a "cold" mode in which he selects his most preferred alternative (by imposing cognitive control), and a dysfunctional "hot" mode in which decisions and

preferences may diverge (because he responds to distorted HFM-generated forecasts). He lives for an infinite number of discrete periods. In each period, he makes two decisions in succession. First he selects a "lifestyle" activity (a); then he allocates resources between an addictive substance ($x \in \{0, 1\}$) and a nonaddictive substance ($e \geq 0$). He enters each period in the cold mode and chooses his lifestyle activity rationally. This choice, along with his history of use and other environmental factors, determines the probability with which he encounters cues that trigger the hot mode. If triggered, he always uses the substance, even if this is not his best choice. If he is not triggered, he rationally decides whether to indulge or abstain.

The intensity (or volume) of substance-related cues encountered, $c(a, \omega)$, depends on the activity a and an exogenous state of nature, ω, drawn randomly from a state space Ω according to some probability measure μ. The function $M(c, s, a, \omega)$ denotes the attractiveness assigned to the drug by the HFM-generated forecast; this depends on the intensity of cues, the chosen activity, the state of nature, and a variable s summarizing the DM's history of use (his *addictive state*). The impulse from this forecast defeats cognitive control and places the individual in the "hot" mode when its strength exceeds some threshold, M^T.

There are $S + 1$ addictive states labeled $s = 0, 1, \dots, S$. Usage in state s leads to state $\min\{S, s + 1\}$ in the next period. No use leads to state $\max\{1, s - 1\}$ from state $s \geq 1$, and to state 0 from state 0. Note that it is impossible to reach state 0 from any state $s \geq 1$. The state $s = 0$ represents a "virgin state" in which the DM has had no contact with the substance. Since people become sensitized to cues through repeated use, we assume $M(c, s', \omega, a) < M(c, s'', a, \omega)$ for $s' < s''$, with $M(c, 0, a, \omega) < M^T$.

The lifestyle activity a is chosen from the set $\{E, A, R\}$. Activity E ("exposure") entails a high likelihood that the DM will encounter a large number of substance-related cues. Examples include attending parties at which the substance is readily available. Activity A ("avoidance") is less intrinsically enjoyable than E, but exposes the DM to fewer substance-related cues [$c(E, \omega) > c(A, \omega)$] and potentially reduces sensitivity to cues [$M(c, s, A, \omega) \leq M(c, s, E, \omega)$]. Examples include staying at home to read or

attending AA meetings. Activity R ("rehabilitation") entails a commitment to clinical treatment at a residential center during the current period. It is even less intrinsically enjoyable than A, it may further reduce exposure and sensitivity to substance-related cues [$c(A, \omega) \geq c(R, \omega)$ and $M(c, s, R, \omega) \leq M(c, s, A, \omega)$], and most importantly, it guarantees abstention ($x = 0$) during the current period.

Let $T(s, a) \equiv \{\omega \in \Omega | M(c(a, \omega), s, a, \omega) \geq M^T\}$. The DM enters the hot state if and only if $\omega \in T(s, a)$. Let $p_s^a \equiv \mu(T(s, a))$ denote the probability of entering the hot mode in addictive state s after selecting lifestyle activity a. Our assumptions on the functions c and M imply the following.

ASSUMPTION 1: $p_{s+1}^a \geq p_s^a$, $p_0^a = 0$, and $p_s^E \geq p_s^A \geq p_s^R$.

In state s, the DM receives an immediate hedonic payoff, $w_s(e, x, a)$ (recall that e denotes his consumption of nonaddictive goods). The dependence of the payoff function on the addictive state incorporates the effect of past usage on current well-being (tolerance, deterioration of health, and so forth). When evaluating the desirability of any possible set of current and future outcomes, the DM discounts future hedonic payoffs at a constant rate δ.

Notice that we do not allow the hedonic payoff to depend on the state of nature, ω. This is in contrast to the more conventional assumption that cravings reflect cue-triggered taste shocks (Laibson, 2001). As indicated in the previous section, we recognize that cravings have hedonic implications. We abstract from this possibility to focus more narrowly on the novel aspects of our theory, which involve cue-triggered mistakes. Allowing for dependence of w_s on ω is straightforward, but our model can account for the key features of addictive behavior without this extension.

With w_s independent of ω, rehabilitation serves only as a precommitment to abstain.[21] Since the

DM's hedonic payoff from abstention is the same regardless of whether he is hot or cold, he never enters rehabilitation with the object of reducing the likelihood of cravings.[22] As a result, the probabilities p_s^R are irrelevant parameters.

In state s the DM has access to resources y_s ("income"). In many cases, it is natural to assume that y_s declines with s due to deteriorating health, reduced productivity (e.g., through absenteeism), and increased out-of-pocket medical expenses. The price of the addictive substance is q, the cost of rehabilitation is r_s (which potentially depends on the addictive state), and the price of the nonaddictive substance is normalized to unity. For simplicity, we assume that the DM cannot borrow or save.

The following notation simplifies our discussion. Let $u_s^a \equiv w_s(y_s, 0, a)$ and $b_s^a \equiv w_s(y_s - q, 1, a) - u_s^a$ for $a \in \{E, A\}$; and let $u_s^R = w_s(y_s - r_s, 0, R)$. Intuitively, u_s^a represents the baseline payoff associated with successful abstention in state s and activity a, and b_s^a represents the marginal instantaneous benefit from use the individual receives in state s after taking activity a. Thus, $u_s^a + b_s^a$ is the payoff for usage. Let $\mathbf{p}_s = (p_s^E, p_s^A, p_s^R)$, $\mathbf{u}_s = (u_s^E, u_s^A, u_s^R)$, $\mathbf{b}_s = (b_s^E, b_s^A)$, $\boldsymbol{\theta}_s = (\mathbf{p}_s, \mathbf{u}_s, \mathbf{b}_s)$, and $\boldsymbol{\theta} = (\boldsymbol{\theta}_0, \dots, \boldsymbol{\theta}_S)$. The vector $\boldsymbol{\theta}$ specifies all pertinent "derivative" parameters. It reflects the properties of the substance, the method of administration, the characteristics of the individual user, and the public policy environment. We make the following assumption (the latter part of which is in keeping with our earlier discussion).

ASSUMPTION 2: *The payoff function w_s is increasing, unbounded, strictly concave, and twice differentiable with bounded second derivative in the variable e (consumption of the nonaddictive good). Moreover, $u_s^E > u_s^A \geq u_s^R$, and $u_s^E + b_s^E > u_s^A + b_s^A$.*

For each state s, the DM follows one of five contingent plans: engage in activity E and then

---

[21] Though we assume that the DM can commit to rehabilitation only one period at a time, this is without loss of generality since he starts each period in the cold mode. In practice, rehabilitation programs may also teach self-management skills and desensitize addicts to cues. One can model these possibilities by assuming that p_s^a (for a given

state or states) declines subsequent to rehabilitation or therapy. Since the evidence suggests that these treatments are not completely effective (Goldstein, 2001, p. 188), the forces described here would still come into play after treatment.

[22] When w_s depends on ω and $p_s^A > p_s^R$, rehabilitation can serve as a strategy for avoiding cues that trigger reductions in hedonic payoffs (through cravings).

use the substance when in the cold mode [$(a, x) = (E, 1)$], engage in E and refrain from use when in the cold mode [$(a, x) = (E, 0)$, henceforth "half-hearted abstention"], engage in A and use when in the cold mode [$(a, x) = (A, 1)$], engage in A and refrain from use when in the cold mode [$(a, x) = (A, 0)$, henceforth "concerted abstention"], or enter rehabilitation [$(a, x) = (R, 0)$]. From Assumption 2, it follows that $(E, 1)$ always dominates $(A, 1)$, so there are in practice only four pertinent choices.

The cold-mode DM is sophisticated in the sense that he correctly anticipates his future choices in either decision mode, and he understands the process triggering the hot mode. Accordingly, his choices in the cold mode correspond to the solution of a simple dynamic stochastic programming problem with a value function $V_s(\theta)$ (evaluated as of the beginning of a period) satisfying

$$(1) \quad V_s(\boldsymbol{\theta}) = \max_{(a,x)\in\{(E,1),(E,0),(A,0),(R,0)\}} u_s^a + \sigma_s^{a,x}b_s^a$$

$$+ \delta[(1 - \sigma_s^{a,x})V_{\max\{1,s-1\}}(\boldsymbol{\theta}) + \sigma_s^{a,x}V_{\min\{S,s+1\}}(\boldsymbol{\theta})]$$

for $s \geq 1$,[23] where $\sigma_s^{a,x}$ represents the probability of consuming the substance in state s with contingent plan (a, x) (so $\sigma_s^{E,1} = 1$, $\sigma_s^{E,0} = p_s^E$, $\sigma_s^{A,0} = p_s^A$, and $\sigma_s^{R,0} = 0$). Existence, uniqueness, and continuity of $V_s(\theta)$ in θ follow from standard arguments.

We close this section with several remarks. First, though simple and stylized, our model adheres closely to the three key premises described in Section II. Specifically, use among addicts is potentially a mistake; experience with an addictive substance sensitizes the user to environmental cues that subsequently trigger mistaken use; and the awareness of this possibility leads users to manage their susceptibilities.

Second, our model reduces to the standard "rational addiction" framework when $p_s^a = 0$ for all s and a. Thus, the novelty of our approach involves the introduction of stochastic shocks (occurring with probability $p_s^a > 0$) that potentially cause decisions to diverge from preferences. This possibility is a central feature of our

model since, without it, the DM would never choose to avoid cues or enter rehabilitation (with $p_s^E = 0$, $(E, 0)$ dominates both $(A, 0)$ and $(R, 0)$). For the same reason, a naive DM who incorrectly believes he does not suffer from a self-control problem (that is, who acts as if $p_s^a = 0$) will never choose cue-avoidance or rehabilitation.

Third, even though our model allows for the possibility that choices and preferences may diverge, with careful use of appropriate data it should still be possible to recover preferences and other critical parameters (such as hot-mode probabilities) empirically. Since we assume that preferences and choices are sometimes aligned, the most obvious approach involves the selective application of the revealed preference principle. The empirical challenge is to identify instances of alignment. One cannot make this determination using only information on choices. We contend, however, that other evidence, such as the research results summarized in Section II, justifies treating the central assumptions of our model as maintained hypotheses. This means that we can use choice data involving precommitments and cue-avoidance to infer hot-mode probabilities and the utility costs of unintended use (recall the discussion in the preceding paragraph). Furthermore, measures of physiological arousal and/or self-reported affective states could be used to differentiate "cold" choices from "hot" choices in experimental settings. For a more general discussion of preference measurement when choices and preferences systematically diverge, see Bernheim and Rangel (2005).

Fourth, unlike other economic theories of addiction, ours does not *necessarily* assume that present use increases the marginal benefit of future use ($b_{s+1} \geq b_s$). We show that, contrary to some claims in the literature, it is possible to explain the central features of addiction without invoking intertemporal preference complementarities (provided the probability of cue-triggered mistakes increases with s). This is important because intertemporal complementarities do not appear to drive some distinctive addictive behaviors,[24] and these behaviors are

---

[23] The associated expression for $s = 0$ is virtually identical, except that $V_0(\boldsymbol{\theta})$ replaces $V_{\max\{1,s-1\}}(\boldsymbol{\theta})$.

[24] The phenomenon of withdrawal is often interpreted as the key manifestation of intertemporal comple-

observed in contexts where such complementarities are probably not present (e.g., compulsive shopping and kleptomania).

Fifth, though one could incorporate the realistic possibility that some individuals are partially myopic with respect to the likelihood and effects of becoming addicted, we assume that the DM is sophisticated in the cold mode. If, as we argue, counterproductive addictive behaviors can arise even with sophisticated decision-makers, efforts to eradicate addiction solely through education and information are misguided.

IV. Positive Analysis

A. *Comparative Dynamics*

Our comparative dynamic results concern the intensity with which the DM voluntarily uses the addictive substance. We study two notions of intensity. We say that the *disposition to use* is greatest for $(E, 1)$, followed in order by $(E, 0)$, $(A, 0)$, and $(R, 0)$. Thus, for example, the disposition to use increases when the DM's choice shifts from $(A, 0)$ to $(E, 1)$. We judge the intensity of *intentional use* by asking whether the DM plans to consume the substance. Thus, intentional use is highest for $(E, 1)$, and equivalent for $(E, 0)$, $(A, 0)$, and $(R, 0)$.[25] An increase in intentional use implies an increase in the disposition to use, but not vice versa. Both definitions permit us to compare the intensity of voluntary use both within states and across states.

We study comparative dynamics with re-

spect to the elements of the parameter vector $\boldsymbol{\theta}$. Since some of these are simple functions of prices and income (q, r_s, and y_s), comparative dynamics with respect to the latter variables follow immediately. We are particularly interested in the effects of the parameters p_s^E and p_s^A, since these are directly tied to the novel aspects of our framework (stochastic events that create pathological discrepancies between preferences and choice). We are also interested in the effects of u_s^A, b_s^A, and u_s^R, since these parameters are relevant only if the novel components of our model are operational ($p_s^E > 0$).

1. *Changes in Individual Parameters.*—In practice, we are rarely interested in phenomena that affect only one state-specific parameter. However, examining these effects in isolation lays the groundwork for subsequent results involving changes in groups of parameters.

PROPOSITION 1: *(i) The disposition to use in state j is:*

(i-a) *weakly increasing in b_k^a and u_k^a, and weakly decreasing in p_k^a, for $k > j$;*
(i-b) *weakly decreasing in b_k^a and u_k^a, and weakly increasing in p_k^a, for $k < j$;*
(i-c) *weakly decreasing in p_j^E and u_j^R and weakly increasing in b_j^E.*

(ii) Intentional use in state j is invariant with respect to p_j^E, p_j^A, u_j^A, b_j^A, and u_j^R.

Parts (i-a) and (i-b) establish the intuitive property that beneficial changes in parameters for more (less) advanced states of addiction increase (decrease) the disposition to use in the current state. Thus, an increase in the likelihood or severity of a cue-triggered mistake in state *s* induces the DM to make choices that reduce the likelihood of reaching state *s*. Part (i-c) is also intuitive: the disposition to use in the current state rises with the benefits to current use and falls with both the desirability of rehabilitation and the likelihood that the exposure activity triggers the hot mode (since this increases the attractiveness of concerted abstention and rehabilitation relative to half-hearted abstention). Part (ii) is perhaps less transparent. A change in parameters can affect intentional use only if it

---

mentarities. Notably, W. E. McAuliffe (1982) showed that only 27.5 percent of heroin addicts experienced cue-triggered withdrawal symptoms, and only 5 percent of these felt these symptoms were responsible for recidivism.

[25] For some parameter values, the DM may be indifferent between two (but never more than two) choices in any particular addictive state. When this occurs, the set of optimal choices is always $\{(E, 1), (E, 0)\}$, $\{(E, 0), (A, 0)\}$, $\{(E, 0), (R, 0)\}$, or $\{(A, 0), (R, 0)\}$. We say that a change in parameters from $\boldsymbol{\theta}'$ to $\boldsymbol{\theta}''$ weakly increases the disposition to use (intentional use) if it leads to a weak increase in both the minimum and maximum disposition to use (intentional use) among optimal choices, and strictly increases the disposition to use (intentional use) if either the minimum or the maximum strictly increases and neither declines.

tips the balance between $(E, 1)$ and $(E, 0)$.[26] Clearly, this comparison does not implicate p_j^A, u_j^A, b_j^A, or u_j^R; neither does it depend on p_j^E.[27]

The characterization of directional effects in Proposition 1 is not quite complete. This is because the effects of b_j^A, u_j^A, and u_j^E on the disposition to use in state j can be positive or negative, depending on the parameter values.

2. *Changes in Groups of Parameters.*—To examine the effects of policy and environmental changes, and to make comparisons between optimal decision rules for different substances, we must typically consider the effects of varying many parameters simultaneously. For example, a general reduction in the cost of rehabilitation (due perhaps to the development of a new therapeutic drug) raises u_s^R for all s. Likewise, when one substance is more addictive than another, it is natural to assume that p_s^a is higher at every state s.

For any state-indexed variable z_s, we say that a change from $(z_s')_{s=0}^S$ to $(z_s'')_{s=0}^S$ represents a *general increase (decrease)* if $z_s' \geq z_s''$ $(z_s' \leq z_s'')$ for all s, with strict inequality for some s. Proposition 1 suggests that compound parameter changes of this type often have ambiguous effects on use. For example, a general increase in u_s^R or a general decrease in p_s^a can reduce the disposition to use in state j by making lower states (weakly) more attractive but can also increase this disposition by making higher states (weakly) more attractive.

It is nevertheless possible to reach a number of conclusions without imposing additional structure. Public policy discussions often emphasize initial use, choices among casual users who are at risk of becoming addicted, and patterns of behavior among hard-core addicts. To shed light on initial use, we study behavior in state 0. To shed light on the choices of casual users, we examine behavior in state 1, the length of the *first intentional use interval* (defined as $\{1, \dots, s^1 - 1\}$ where s^1 is the largest integer such that $(E, 1)$ is chosen for all $s \in \{1, \dots, s^1 - 1\}$ but not for s^1), and the length of the *initial resistance interval* (defined as $\{1, \dots, s^2 - 1\}$ where s^2 is the largest integer such that $(R, 0)$ is chosen for all $s \in \{1, \dots, s^2 - 1\}$ but not for s^2).[28] To shed light on the behavior of those with substantial cumulative exposure, we focus on choices in state S, as well as the length of the *final resignation interval* (defined as $\{s^3 + 1, \dots, S\}$ where s^3 is the smallest integer such that $(E, 1)$ is chosen for all $s \in \{s^3 + 1, \dots, S\}$ but not for s^3). While these aspects of behavior respond ambiguously to general changes in some parameters, other effects are unambiguous.[29]

PROPOSITION 2: *(i) A general increase in p_s^E or p_s^A, or a general reduction in u_s^A or b_s^A, weakly decreases the disposition to use in state 0 (and state 1 for p_s^E), weakly shortens the first intentional use interval, weakly lengthens the initial resistance interval, and weakly lengthens the final resignation interval.*

(ii) A general increase in u_s^R weakly increases the disposition to use in all states (including state 0) up to (but not including) the first state in which rehabilitation is an optimal choice after the increase. It also weakly lengthens the first intentional use interval, weakly reduces the disposition to use in state S, and weakly shortens the final resignation interval.

(iii) A general increase in u_s^E or b_s^E weakly shortens the initial resistance interval. In addition, a general increase in b_s^E weakly increases the disposition to use in states 0 and 1.

How do patterns of use compare for two substances that are the same in all respects, except

---

[26] If $(E, 0)$ yields a higher expected discounted payoff than $(E, 1)$, then $(E, 1)$ is obviously not the DM's best choice. Conversely, if $(E, 1)$ yields a weakly higher expected discounted payoff than $(E, 0)$, then $(E, 1)$ is necessarily preferred to both $(A, 0)$ and $(R, 0)$. To understand why, note that (a) $u_s^E + b_s^E + \delta V_{\min\{S, s+1\}}(\boldsymbol{\theta}) \geq u_s^E + \delta V_{\max\{1, s-1\}}(\boldsymbol{\theta}) > u_s^a + \delta V_{\max\{1, s-1\}}(\boldsymbol{\theta})$ for $a = A, R$ [where the first inequality follows because the DM weakly prefers $(E, 1)$ to $(E, 0)$, and the second inequality follows from Assumption 2], and (b) $u_s^E + b_s^E + \delta V_{\min\{S, s+1\}}(\boldsymbol{\theta}) > u_s^A + b_s^A + \delta V_{\min\{S, s+1\}}(\boldsymbol{\theta})$ (by Assumption 2). For $(R, 0)$, the desired conclusion follows from (a); for $(A, 0)$, it follows from (a) and (b).

[27] The DM prefers $(E, 1)$ to $(E, 0)$ if and only if he prefers it to E with the certainty of abstention. The probability p_j^E does not enter this comparison.

[28] At least one of these intervals is always empty. The length of the initial resistance interval is relevant only if parameters change after the DM starts using the substance (otherwise he would never advance beyond state 1). In that case, it sheds light on the DM's ability to achieve permanent recovery.

[29] To allow for multiple optima, we say that a parameter change weakly shortens (lengthens) an interval if it weakly reduces (increases) both the minimum and maximum length of the interval.

that one is more addictive than the other (higher values of p_s^E and p_s^A for all s)? Part (i) of the proposition provides a partial answer. Not surprisingly, an increase in addictiveness discourages use among new users (reducing the disposition to use in state 0, shortening the first intentional use interval, and lengthening the initial resistance interval). Strikingly, it always has the opposite effect on hard-core addicts, producing longer resignation intervals. One might think that an increase in addictiveness might discourage a relatively advanced user from taking actions likely to place him in an even more highly addicted state, but this effect never materializes in the final resignation interval. Instead, the DM is influenced by the increased futility of resisting use at lower states. He resigns himself to severe addiction because he recognizes his powerlessness to control his subsequent behavior adequately at lower states, despite intentions to abstain. According to part (i), general changes in the parameters governing payoffs from the avoidance activity (u_s^A and b_s^A) have similar effects.

How does an improvement in rehabilitation technology (higher values of u_s^R for all s) affect patterns of use? According to part (ii) of the proposition, use among those with low cumulative exposure increases in a strong sense (the disposition to use rising in all states up to the point where the DM enters rehabilitation). Since rehabilitation cushions the negative effects of addiction, this is not surprising. As in part (i), this development has the opposite effect on hard-core addicts, shortening the resignation interval. Notably, increasing u_s^R only for states in the resignation interval would have no effect on behavior. Thus, for a general increase in u_s^R, the DM turns away from intentional use in the resignation interval because rehabilitation becomes a more attractive option in lower states.

Part (ii) of Proposition 2 also implies that an improvement in rehabilitation technology can have the perverse effect of shifting the entire population distribution to more addicted states. Provided that all members of the population start out at $s = 0$, this occurs when a general increase in u_s^R raises the lowest state at which the DM selects rehabilitation.[30]

Part (iii) of the proposition concerns u_s^E and b_s^E. These parameters do not relate to the novel features of our model, but we have included their effects for completeness.

Proposition 2 underscores the fact that changes in the environment have complex effects on use, often driving consumption among new users and hard-core addicts in opposite directions. It is natural to wonder whether there are any general parameter changes that always have the same directional effect on the disposition to use in every addictive state. Our next result provides an example: if baseline well-being deteriorates more rapidly as the addictive state rises, then the disposition to use is lower in every state. This property holds in the standard rational addiction framework ($p_s^a = 0$) and is preserved in the presence of cue-triggered mistakes.

PROPOSITION 3: *Consider* $\bar{\theta}$ *and* $\underline{\theta}$ *derived, respectively, from* $\bar{w}_s(e, x, a)$ *and* $\underline{w}_s(e, x, a) = \bar{w}_s(e, x, a) - d_s$ *(with the same values of* y_s, r_s, q, *and* p_s^a). *If* d_s *is weakly increasing in* s, *the disposition to use is weakly higher with* $\bar{\theta}$ *than with* $\underline{\theta}$ *for all* s.

Propositions 2 and 3 shed light on the relationship between income and the consumption of addictive substances. While an increase in income raises the inclination to experiment recreationally, it can reduce the inclination to use at higher addictive states; accordingly, the model can generate higher rates of addiction among the poor. To see why, suppose the utility function has the following separable form: $w_s(e, x, a) = u(e) + v_s(x, a)$. What happens when we add a fixed increment, Δ, to income in all states? The parameters u_s^E and u_s^A all rise by $u(y_s + \Delta) - u(y_s)$, which is weakly increasing in s (assuming y_s is weakly decreasing in s). The parameters u_s^R increase by $u(y_s - r_s + \Delta) - u(y_s - \Delta) > u(y_s + \Delta) - u(y_s)$, and the parameters b_s^a increase by $[u(y_s - q + \Delta) - u(y_s + \Delta)] - [u(y_s - q) - u(y_s)] > 0$. Thus, we can decompose the effect into three pieces: (a) a fixed increase in u_s^a equal to $u(y_s + \Delta) - u(y_s)$ for each s and $a = E, A, R$; (b) a general

_____

[30] From simulations, we know that a general increase in u_s^R increases the lowest state at which the DM selects

rehabilitation for some parameter values and decreases it for others.

increase in u_s^R, and (c) a general increase in b_s^a for $a = E, A$. Proposition 3 tells us that the effect of the first piece is to weakly increase the disposition to use in every state. Part (ii) of Proposition 2 tells us that the second piece increases the disposition to use in state 0, and parts (i) and (iii) tell us the same thing for the third piece. Thus, the addictive substance is normal in state 0. Note, however, that the second and third pieces have ambiguous effects on the disposition to use in more advanced states of addiction, which is why the income effect can change signs.

B. *Patterns of Consumption*

According to our theory, the particular pattern of consumption that emerges in any instance depends systematically on the characteristics of the individual (including aptitude for cognitive control), the substance, and the environment. For reasonable parameter values, the model generates a wide variety of observed consumption patterns.[31]

Consider a highly addictive substance (p_s^a large). If baseline well-being declines rapidly with consumption, the DM may choose never to use [$(E, 0)$ at $s = 0$]. For most people, crack cocaine appears to be a good example of this (see Goldstein, 2001). In contrast, if the decline in baseline well-being is initially gradual but accelerates from one state to the next, the model can produce a pattern of progressive resistance. That is, the DM may begin using the substance intentionally, engage in half-hearted abstention (and therefore use intermittently) after reaching an intermediate addictive state, and shift to concerted abstention after a string of bad luck. If bad luck continues, precommitment to abstention through rehabilitation may follow with subsequent probabilistic recidivism. If baseline well-being flattens out for sufficiently advanced addictive states (the DM "hits bottom"), the model can also produce resignation. That is, a DM may give up, opting for $(E, 1)$ once he reaches a highly addicted state after an unsuccessful battle to abstain.

---

[31] We have generated each of the patterns described in this section through numerical simulations, which we omit to conserve space.

Now consider an enjoyable substance for which baseline well-being declines slowly with consumption. Irrespective of whether the probability of entering the hot mode is high or low, constant use often emerges. Caffeine potentially fits this description.

Finally, a sufficiently sharp drop in the pleasure generated by the substance from one addictive state to the next can produce *intentional* recidivism. That is, the DM may choose $(E, 1)$ in one state and $(R, 0)$ in the next, in which case he oscillates between the two. He enters rehabilitation in each instance without any desire to stay clean; he knows that he will resume using the substance upon release from rehabilitation, and fully expects to enter rehabilitation once again. This pattern is in fact observed among serious heroin users when repeated use dilutes the "high" (see Michael Massing, 2000). It is evidence of fairly sophisticated, forward thinking among junkies whose objective is to renew the high by temporarily getting clean, and who know that rehabilitation accomplishes this more reliably than abstention.

C. *Explaining the Distinctive Features of Addiction*

In Section I we argued that addiction is associated with five distinctive behavioral patterns. Our theory generates each of these patterns.

1. *Unsuccessful Attempts to Quit.*—Suppose life circumstances change over time, gradually shifting the parameters of the DM's problem from θ' to θ''. Suppose the DM's best choice for state 0 is $(E, 1)$ if θ' prevails forever, but that the optimal decision rule prescribes either $(E, 0)$, $(A, 0)$, or $(R, 0)$ for all s if θ'' prevails forever. If the shift from θ' to θ'' is either unanticipated or anticipated and sufficiently slow, the DM starts using the substance but subsequently decides to quit unconditionally. With $p_1^a > 0$, the attempt is unsuccessful when either $(E, 0)$ or $(A, 0)$ is chosen in state 1.

2. *Cue-Triggered Recidivism.*—For the setting described in the previous paragraph, unsuccessful attempts to quit are associated with high realizations of $c(a, \omega)$ (that is, exposure to relatively intense cues).

3. *Self-Described Mistakes.*—In our model, choices and preferences diverge whenever the DM selects $(E, 0)$ or $(A, 0)$ and then enters the hot mode. This constitutes a recognizable mistake.[32]

4. *Self-Control through Precommitment.*—The choice $(R, 0)$ is a costly precommitment; under our assumptions, its only purpose is to remove the option of consuming the substance.

5. *Self-Control through Behavioral and Cognitive Therapy.*—The choice $(A, 0)$ involves costly cue-avoidance. Its only purpose is to reduce the probability of encountering cues that trigger mistaken usage. Though not modeled explicitly, cognitive therapy would influence behavior in our setting by increasing M^{T} (that is, raising the threshold impulse required to defeat cognitive control).

Our three central premises play critical roles in accounting for each pattern. We can remove cue-triggered mistakes by setting $M^{T} = +\infty$, so that the DM always exercises cognitive control. With this change, $p_s^a = 0$ for all a and s, and the DM always choose either $(E, 1)$ or $(E, 0)$. All attempts to quit are successful, and there is no recidivism. Preferences and choices never diverge, so there are no mistakes. The DM never exercises self-control through precommitment by choosing $(R, 0)$, or through cue-management by choosing $(A, 0)$. Some sophistication is also essential; otherwise the DM would ignore his susceptibility to cue-triggered errors and make choices based on the mistaken assumption that $p_s^a = 0$ for all a and s.

We also observed in Section I that aggregate consumption of addictive substances responds to prices and information in the usual way. This too is consistent with our theory, as users sometimes make decisions in the cold mode.

V. Demand-Side Policy Analysis

In this section we study the welfare effects of various public policies concerning addictive substances. In keeping with the focus of the preceding sections, we restrict attention to "demand side" welfare effects, ignoring "supply side" consequences associated with the development of black markets, the spread of corruption, and enforcement costs.[33]

A. *The Welfare Criterion*

In formulating our model, we retain the standard assumption that each individual has a single coherent set of preferences. Our departure, which is grounded in the evidence from neuroscience presented in Section II, is to assume that there are imperfections in the *process* by which the brain makes choices, and that these imperfections give rise to mistakes in identifiable circumstances. Since the individual has only one set of preferences, discounted experiential utility, $\Sigma_{t=0}^{\infty} \delta^t w_{s_t}(e_t, x_t, a_t)$, accurately measures his well-being, and is unambiguously the appropriate welfare standard.[34]

It may be tempting to reinterpret our model as one with multiple selves, "hot" and "cold," where the preferences of the hot self can be inferred from choices in the hot mode. Under that interpretation, our use of cold preferences as a welfare standard is arbitrary. In our view, this interpretation commits a fallacy. By assuming that choices are *always* consistent with underlying preferences, it assumes away the possibility that individuals make systematic mistakes. This possibility is a central premise of our analysis and is justified based on the state of knowledge concerning the neuroscience of addiction. One can certainly dispute the validity of this premise. However, given the premise and our adherence to the standard formulation of preferences, the correct welfare criterion is unambiguous.

---

[32] When the HFM-generated forecast is sufficiently positive, cognitive override may not occur even when higher cognition forecasts undesirable consequences. Thus, an individual may use a substance while simultaneously recognizing (in terms of higher cognitive judgment) that this is a mistake.

[33] Supply-side effects are discussed elsewhere (see, e.g., MacCoun and Reuter, 2001; J. Miron and J. Zwiebel, 1995).

[34] This is in contrast with a number of the behavioral theories discussed in Section VI, for which one must either use a weak welfare standard such as the Pareto criterion (applied to multiple selves or multiple perspectives), or select a particular method of resolving conflicting preferences, for example, by respecting the tastes of only one self or perspective.

B. *Policy Objectives*

What might society hope to accomplish through public policies regarding addictive substances? Possible objectives include protecting third parties from externalities (e.g., second-hand smoke), combatting misinformation and ignorance, moderating the consequences of un-insurable risks, and helping consumers avoid mistakes. Both externalities and informational problems provide well-understood rationales for government intervention, and neither is intrinsically linked to the novel aspects of our framework. We therefore focus primarily on the last two sets of objectives.

1. *Amelioration of Uninsurable Risks.*—Risk and uncertainty relating to the effects of environmental cues on decision processes are central to our model. The DM's lack of knowledge concerning future states of nature, ω_t, prevents him from perfectly forecasting future decision modes and choices in states for which he plans to select either $(E, 0)$ or $(A, 0)$ and thereby creates uncertainty about subsequent addictive states. This translates into *monetary* risk because his resources depend on his addictive state, and because variation in expenditures on the addictive substance and rehabilitation imply variation in consumption of the nonaddictive good. Since the financial consequences of addiction—its effects on job retention, productivity, out-of-pocket medical costs (including rehabilitation), and, for some substances (e.g., cocaine, heroin), direct expenditures—are often substantial, this risk is quantitatively significant.[35]

From the perspective of risk (ignoring other considerations), policies that create actuarially fair redistributions over realizations of future states of nature are beneficial (harmful) when they distribute resources toward (away from) outcomes for which the marginal utility of nonaddictive consumption is relatively high. In subsequent sections, we initially impose the assumption that $w_s(e, x, a) = u(e) + v_s(x, a)$, which implies that the marginal utility derived from nonaddictive consumption depends only

(and inversely) on the level of nonaddictive consumption. We focus on this case because we regard it as a natural benchmark, but we also discuss the implications of relaxing separability.

We assume throughout that private insurance markets fail completely. As is well known, the welfare effects of public policies that redistribute resources across states of nature can depend on the specific factors that cause markets to fail (see, e.g., Mark V. Pauly, 1974). It is therefore important to specify the source of the market failure and to explain how it interacts with the policies considered.

We assume that private insurance companies are unable either to observe or to verify the state of nature ω_t, cues, the DM's decision mode, the addictive state, lifestyle activities, or consumption of the addictive substance.[36] The government is similarly handicapped. However, unlike private companies, it can observe transactions involving legal addictive substances (typically without identifying purchasers), and it can manipulate the prices of these commodities through taxation and subsidization.

Private companies can observe aspects of treatment (rehabilitation and medical costs), but we assume that treatment insurance is unavailable because (i) practical considerations preclude ex ante contracting at age zero when risks are homogeneous (e.g., before teen or even pre-teen exposure), and (ii) adverse selection arising from ex post heterogeneity precludes ex post contracting.[37] The government is similarly handicapped by the second problem but can avoid the first by imposing a universal policy on all consumers ex ante.

2. *Mistake Avoidance.*—Public policy can potentially improve welfare by creating conditions that reduce the frequency with which individuals experience decision process malfunctions or by forcing them to make alternative

---

[35] Among chronic users, average annual expenditures on cocaine and heroin exceeded $10,000 in 1999 (Office of National Drug Control Policy, 2001a).

[36] With respect to the addictive state, clinical diagnosis of addiction is both costly and imprecise. Consumption is potentially observable when the substance is dispensed as a prescription medicine, but in that case the same problems arise as for treatment (discussed in the next paragraph).

[37] In practice, private health insurance policies do provide some coverage for the treatment of addiction. Yet many people are not insured, and coverage is typically incomplete.

choices when malfunctions occur. In our model, these possibilities correspond, respectively, to reducing the probability of entering the hot mode and to ensuring abstention when appropriate.

In solving the DM's optimization problem, we treated the probabilities of entering the hot mode, p_s^a, as fixed parameters. Many of the policies considered below potentially change these parameters; for example, one can model a ban on advertising as a reduction in the volume of cues encountered. For some portions of our analysis, we also allow for the possibility that p_s^a depends on the price of the addictive good, q, and/or the DM's income, y. We reason that the frontal cortex is more likely to generate stronger cognitive incentives and override cues when the immediate consequences of use are more severe.[38] Accordingly, we assume that M^T weakly decreases with y, weakly increases with q, and weakly increases with an *equal* increase in y and q.[39] It is useful to state these assumptions more compactly in terms of p_s^a (with the added technical requirement of differentiability).

ASSUMPTION 3: *The probability of entering the hot mode, p_s^a, is differentiable in q and y with*

$$-\frac{\partial}{\partial q} p_s^a \geq \frac{\partial}{\partial y} p_s^a \geq 0.$$

C. *When Is Government Intervention Justified?*

When do the objectives discussed in Section V B potentially justify government inter-

vention? The following result provides an initial answer. Here and elsewhere, we say that use is *continual* if the DM selects $(E, 1)$ in every state.

PROPOSITION 4: *(i) Continual use solves the DM's choice problem if and only if it is first best (in the sense that it solves the maximization problem when $p_s^a = 0$ for all a and s). (ii) Suppose there is some state s' with $p_{s'}^E > 0$ such that $(E, 1)$ is not a best choice in s'. Then the DM's choices are not first best (in the sense that setting $p_s^a = 0$ for all a and s and reoptimizing strictly increases the value function for some states).*

Part (i) tells us that noncontinual use is *necessary* for the existence of a beneficial policy intervention. Laissez-faire is therefore the best policy for substance users who make no serious attempt to abstain (e.g., contented smokers or coffee drinkers). Notably, this conclusion follows even when the substance in question is highly addictive (p_s^a rises sharply with s) and well-being declines significantly with long-term use. The intuition is the same as for the final portion of Proposition 1.

Part (ii) tells us that noncontinual use is *sufficient* for the existence of a theoretical policy intervention that benefits the DM, provided the departure from continual use occurs in a state for which the DM is susceptible to cue-triggered mistakes. Of course, this intervention may be impractical given the government's information constraints.

D. *A Framework for Tax Policy Analysis*

The formal results below concern the desirability of various types of tax policies. Following standard practice, we evaluate these policies by embedding our decision-maker in a simple economy and studying effects on equilibrium allocations. Here we outline the structure of the economic environment. Notation and some additional formal details appear in Appendix A.

The economy consists of an infinite sequence of generations. In the absence of government intervention, every member of every generation is identical and confronts the decision problem

---

[38] Since cognitive control often must be asserted quickly if at all, and since extrapolation of future consequences is time consuming, we implicitly assume that the deployment of cognitive control responds to variation in immediate circumstance-specific consequences, but not to variation in future circumstance-specific consequences.

[39] When q is higher or y is lower, the immediate negative consequences of use are plainly greater, and potentially more likely to occupy the DM's awareness. When q and y rise by equal amounts, the immediate hedonic payoff from abstention rises while the immediate hedonic payoff from use is unchanged, so the immediate negative consequence of use is again more severe. It would also be natural to assume that cue exposure, $c(a, \omega)$, weakly declines in q (since use among social contacts declines), and this would reinforce Assumption 3. In principle, $c(a, \omega)$ could rise or fall with income.

described in Section III. We interpret the DM's discount factor δ as the product of a pure rate of time preference and a constant single-period survival probability. We assume that the size of each new generation is just sufficient to keep the total population constant. The total population is very large, and realization of hot and cold states are independent across DMs, so there is essentially no aggregate uncertainty. Income arrives in the form of the nonaddictive good, and the addictive good is produced under competitive conditions with constant-returns-to-scale technology, as are rehabilitation services. Thus, q and r_s are fixed and equal to unit production costs (where the costs of rehabilitation services potentially vary with the addictive state).

In each period, the government can tax (or subsidize) either the addictive substance or rehabilitation (we consider these instruments one at a time). There is no revenue requirement; taxes are purely corrective. By assumption, the government cannot condition the associated tax rates on either the DM's age or his addictive state. This could reflect either the practical difficulties associated with tailoring these taxes and subsidies to an individual's conditions (including the need for clinical diagnosis) or privacy concerns. Imagine in particular that the tax/subsidy for the addictive substance is either applied to anonymous transactions (like a sales tax) or imposed on producers (like a value-added tax), while the rehabilitation tax/subsidy nominally falls on service providers (again, like a value-added tax). The government can also use age-specific (equivalently, generation-specific) lump-sum instruments. An intertemporal policy specifies values for all available tax/subsidy instruments in every period.

Since there is no borrowing or lending in our model, and since we do not wish to advantage the government artificially, we assume that policies cannot redistribute resources across periods. We say that an intertemporal policy is *feasible* if there is, for each generation, an optimal decision rule such that the government's budget is balanced in every period. Feasible policies permit within-period transfers across generations, which can mimic borrowing and lending, thereby leaving the government in an artificially advantageous position. One could therefore argue for a stronger restriction requiring government budget balance for each gener-

ation within each period. While we impose the weaker requirement, our results also hold for the stronger requirement.[40]

A steady-state policy prescribes a constant tax rate and constant age-specific lump-sum taxes. Notably, each individual's problem is potentially nonstationary because steady-state lump-sum tax/subsidies may change with age. The set of feasible steady-state policies includes the zero-tax alternative, henceforth denoted ϕ, for which all tax/subsidy instruments are set to zero.

In the next two sections, we focus on the steady-state welfare effects of steady-state policies (often dropping the modifier "steady-state" for brevity). For any steady-state policy, we use the lifetime expected discounted hedonic payoff for the representative individual as our welfare measure. An optimal steady-state policy maximizes this payoff among all feasible steady-state policies. This objective function respects each individual's time preference over his own lifetime but is infinitely patient with regard to intergenerational comparisons, in effect placing equal weight on all generations. Since the DM's choice set is discrete, best choices are often insensitive to small parameter changes, so the optimal policy is typically not unique.

E. *Taxation and Subsidization of Addictive Substances*

Addictive substances are often heavily taxed (e.g. nicotine and alcohol) and occasionally subsidized (see, e.g., the description of a Swiss heroin prescription program in MacCoun and Reuter, 2001). Some policy analysts argue for taxation of addictive substances on the grounds that this discourages excessive use (e.g., Gruber and Koszegi, 2001). Others suggest that, in the absence of externalities, use is voluntary so laissez-faire is best (e.g., Becker and Murphy, 1988). Our theory of addiction suggests a more nuanced view.

Proposition 5 below relates the sign of the optimal tax rate on the addictive substance to

---

[40] In fact, the proof of Proposition 5 requires only minor adjustments when we impose the strong requirement; Proposition 6 holds as stated under either restriction.

observable patterns of consumption. Notably, the consumption patterns that determine optimal tax rates are endogenous, and the proposition requires us to assess them at the optimal tax rates.[41] This feature is common to many well-known optimal-tax results. For example, the Ramsey rule relates optimal commodity tax rates to compensated demand elasticities *evaluated at the optimal tax rates* (though the rule is frequently stated in a way that disguises this dependency).

The proposition refers to the following two possible patterns involving the likelihood of use.[42]

Condition A: For every age t, the likelihood of use is weakly increasing in s over states reached with positive probability at that age, and the DM does not enter rehabilitation in the lowest such state.[43] Moreover, at some age t, at least two addictive states are reached with positive probability.

Condition B: For every age t, the likelihood of use is weakly decreasing in s over states reached with positive probability at that age.[44] Moreover, at some age t, at least two addictive states are reached with positive probability, with neither expected use nor y_s constant over such states.

---

[41] Alternatively, one can make statements about welfare-improving changes, assessing usage patterns at arbitrary starting points. For example, from the proof of Proposition 5, we also have the following results: eliminating a positive tax is beneficial if initially Condition A holds; eliminating a positive subsidy is beneficial if initially Condition B holds and q is small; if Condition A holds with the no-tax policy ϕ, a small subsidy is welfare improving; if Condition B holds with ϕ and q is small, a small tax is welfare improving.

[42] In Appendix A, we define $\sigma_s^t(\chi)$ as the probability of use in state s at age t given a decision rule χ, accounting for the possibility of entering the hot mode. Here, the "likelihood of use" refers to $\sigma_s^t(\chi)$. Note that the likelihood of use is necessarily weakly increasing in s when the disposition to use is weakly increasing in s.

[43] Since the DM's decision problem is potentially nonstationary, it is possible for him to find himself in a state beyond the lowest one in which he would select rehabilitation *during the same period.*

[44] This occurs, for example, if best choices are unique, the disposition to use is weakly decreasing in s, the first intentional use interval is nonempty, and p_s^a is constant outside of this interval (i.e., the DM is fully addicted by the time he attempts to refrain from consuming the substance).

For each condition, the requirement "for all t" is less demanding than it might initially appear. Remember that, in the absence of taxes and subsidies, each DM's problem is stationary, and the best choices at each state are independent of age. In a steady state for the economy, age matters only because it affects the lump-sum tax (or subsidy). If the lump sums are relatively small, the general pattern of use will tend to be similar at different ages provided it is not too sensitive to small changes in income.

PROPOSITION 5: *Suppose that $w_s(e, x, a) = u(e) + v_s(x, a)$, that y_s is weakly decreasing in s, and that p_s^a does not depend on prices or income.*

(i) *Consider an optimal steady-state policy for which all budget-balancing optimal decision rules satisfy Condition A. The tax rate on the addictive substance is strictly negative.*

(ii) *Consider an optimal steady-state policy for which all budget-balancing optimal decision rules satisfy Condition B. If q is sufficiently small, the tax rate on the addictive substance is strictly positive.*

To develop intuition for this result, note that taxation (or subsidization) potentially affects welfare through three channels. First, it can change decisions in the cold mode. Second, it can redistribute resources across uncertain outcomes. Third, it can alter the effects of environmental cues on operational decision modes (through the trigger mapping T). With p_s^a independent of prices and income, the third channel vanishes (we discuss the implications of reinstating it below). Effects involving the second channel dominate welfare calculations for small taxes and subsidies because they are generally first order, while effects involving the first channel are not.[45] Accordingly, starting from a situation with no taxes, one can determine whether a small tax or subsidy improves welfare by focusing on the correlation between the taxed

---

[45] With continuous decision variables and interior solutions, the first channel would be second order for small taxes and subsidies. With discrete decision variables (as in our current model), it is literally zero for sufficiently small taxes and subsidies.

activity and the marginal utility of nonaddictive consumption.[46] This helps to build intuition concerning the signs of optimal tax rates. One should bear in mind, however, that the optimal-tax problem is more complex because, at the optimum, welfare effects involving the first channel (incentive effects) are also first order.

A policy provides de facto insurance if it redistributes age-t resources toward outcomes in which the marginal utility of nonaddictive consumption is relatively high, which in the benchmark case $[w_s(e, x, a) = u(e) + v_s(x, a)]$ means that the level of nonaddictive consumption is relatively low. This occurs when other expenditures are high and when the state is itself high (since y_s weakly declines with s).[47] A subsidy necessarily redistributes resources toward outcomes with relatively high expenditures on addictive substances. Moreover, if the likelihood of use increases with the addictive state (Condition A), it also redistributes resources toward outcomes for which income is relatively low. Since both effects are beneficial, a subsidy is desirable [part (i)]. Conversely, if the likelihood of use declines with the addictive state (Condition B), a positive tax (with a budget-balancing lump-sum payout) redistributes resources toward outcomes for which income is relatively low and rehabilitation expenditures are relatively high. It also redistributes resources away from outcomes with relatively high expenditures on addictive substances, but this effect is secondary when q is small, rendering the tax beneficial [part (ii)].

Proposition 5 underscores the fact that different policies are appropriate for different addictive substances, and that the characteristics of good policies are related to usage patterns. As we have seen in Section IV, usage patterns are in turn systematically related to aspects of the substance, the user, and the environment.

Part (i) suggests that a subsidy may be welfare improving in the case of a substance for which initial use tends to be "spur of the moment," but where an intention to use becomes increasingly predominant as the individual becomes more addicted. The argument for subsidization is stronger when the substance in question is more expensive. The apparent implication that the government might beneficially subsidize substances such as cocaine and heroin is provocative to say the least, and it should be tempered by several considerations, including the likely existence of externalities, the potential effects of price and income on the trigger mechanism (discussed below), and the fact that Condition A apparently does not hold universally, as many addicts seek treatment. Still, our analysis adds a potentially important cautionary note to existing discussions of the benefits of sin taxes (e.g., Gruber and Koszegi, 2001; Ted O'Donoghue and Matthew Rabin, 2004), which can violate social insurance principles by penalizing those who have experienced bad luck. It also provides a framework for understanding the potential benefits of somewhat more refined approaches, such as the Swiss policy of providing cheap heroin to users who cross some diagnostic threshold of addiction, and who are not interested in rehabilitation.

Part (ii) suggests that a tax may be welfare improving in the case of an inexpensive substance that people initially use regularly, for which attempts to abstain begin only after cue triggers are well established and stable (so that they change little with further use). Coffee, cigarettes, and alcohol arguably fall into this category.

How robust are these findings? Complementarity between addictive and nonaddictive consumption would raise the marginal utility of nonaddictive consumption whenever the DM uses the addictive good, strengthening the advantages of a subsidy, thereby reinforcing part (i) but potentially reversing part (ii). Substitutability would reduce the marginal utility of nonaddictive consumption whenever the DM uses the addictive good, strengthening the advantages of a tax, thereby reinforcing part (ii) but potentially reversing part (i).

---

[46] There are some subtleties here. A small tax or subsidy that changes cold-mode decisions can alter the correlation between a taxed activity and the marginal utility of nonaddictive consumption, thereby changing effects through the second channel. With discrete choice sets, the pertinent correlation can change dramatically even for tiny taxes and subsidies.

[47] In principle, the government could also redistribute resources through an income tax. Implicitly, we take the income-tax system as exogenously given. This is reasonable as long as addiction is not one of the primary factors influencing income distribution and the equity–efficiency trade-offs that an optimal income tax system is intended to address.

We can also relax the assumption that p_s^a is invariant with respect to taxation and subsidization. A tax of τ per unit increases price by τ and, from individuals of age t, raises less than τ in per capita revenues. Suppose the government distributes the revenue raised from each age group back to the same age group as a lump sum. Since the amount received by each individual is less than the price increase, Assumption 3 implies that p_s^a falls in every state. Any policy that reduces p_s^a weakly increases welfare through the third channel [strictly if the state is reached with positive probability and the DM selects $(a, 0)$]. This strengthens the advantages of a tax, reinforcing part (ii) of the proposition, and potentially reversing part (i).[48]

F. Harm-Reduction Policies

Subsidization of rehabilitation is relatively common. Popular justifications appeal to the notion that treatment should be affordable and universally available, though sometimes positive externalities are invoked.

In the context of our model, there are at least two reasons to subsidize rehabilitation. First, this provides de facto insurance for a large, uncertain expense. Second, under Assumption 3, the DM is less prone to make cue-triggered mistakes when he receives resources in kind (through a rehabilitation subsidy) rather than in cash.

There is, however, an additional consideration arising from the correlation between rehabilitation and income. If rehabilitation is more likely at advanced stages of addiction, then a subsidy beneficially redistributes resources toward low-income states. Since this reinforces the considerations discussed in the previous paragraph, subsidized rehabilitation is unambiguously desirable. If, however, rehabilitation is less likely at advanced stages of addiction, then a subsidy detrimentally redistributes resources toward high-income states, offsetting the considerations discussed in the previous paragraph. Formally, one can prove a result analogous to Proposition 5, relating the optimal tax/subsidy

treatment of rehabilitation to rehabilitation patterns.

Our next result deals instead with the welfare effects of small rehabilitation taxes and subsidies. It shows that a small rehabilitation subsidy is beneficial, and a small tax harmful, under extremely general conditions: at the no-tax alternative ϕ, rehabilitation must be chosen in some state, and there must be some randomness.[49] Here, we allow from the outset for the possibility that p_s^a depends on y_s and q.

PROPOSITION 6: *Suppose that $w_s(e, x, a) = u(e) + v_s(x, a)$, that y_s is weakly decreasing in s, and that $r_s > q$ for all s. Suppose also that, in the absence of taxes and subsidies (that is, with policy ϕ), the following conditions hold: first, there is at least one state in which rehabilitation is a best choice; second, rehabilitation is the unique best choice in the earliest of these; third, for some earlier state (other than 0), $(E, 1)$ is not a best choice. Then, within the class of policies that do not create net inter-cohort transfers, a small steady-state subsidy for rehabilitation is beneficial, and a small steady-state tax is harmful.*

Since Proposition 6 holds even when the cost of rehabilitation is very small, it is not primarily about the desirability of insuring a large, uncertain expense. For the correct intuition, note that with the no-tax alternative ϕ, each DM's problem is stationary, so best choices for each state are independent of age. This implies that the DM can never advance beyond the first state in which $(R, 0)$ is the best choice. Consequently, the likelihood of rehabilitation is positively correlated with the addictive state and negatively correlated with income, so the three effects discussed at the outset of this section work in the same direction, in favor of subsidization. The practical lesson is simple: if addiction is relatively unlikely to advance beyond the point where people start to seek rehabilitation, then subsidies are unambiguously desirable.

An appropriately modified version of our model could address the effects of other harm-

---

[48] The proof of Proposition 6 formally demonstrates a closely related point in the context of a subsidy for rehabilitation services.

[49] The assumption that $(E, 1)$ is not a best choice in every state up to the first in which rehabilitation is selected ensures some randomness.

reduction policies such as needle exchanges. We leave this for future work.

G. *Criminalization*

Historically, criminalization has been the cornerstone of U.S. drug policy, with more than 600,000 citizens incarcerated for drug-related offenses in 1999 (Office of National Drug Control Policy, 2001b). It affects users through two distinct channels: a *price effect* and a *rationing effect.* The price effect refers to changes in the marginal cost of using the substance resulting from penalties and other costs imposed on users and suppliers. The rationing effect refers to interference with the process of matching buyers and sellers: since criminalization forces buyers and sellers to carry out transactions secretively, buyers sometimes have difficulty locating supply.[50]

It is instructive to consider the price and rationing effects separately. The price effect is equivalent to a tax policy in which the revenue raised by the tax is destroyed. If criminalization only created a price effect, taxation would dominate it.

Now consider the rationing effect. Disrupting access to supply is potentially beneficial when the DM chooses $(E, 0)$ or $(A, 0)$, and potentially detrimental when he chooses $(E, 1)$. However, the impact of the rationing effect on consumption may be smaller when the DM chooses $(E, 1)$. An individual who *intends* to consume an illegal substance can set about locating supply deliberately and systematically and can maintain stocks in anticipation of transitory difficulties. In the extreme case where the rationing effect has no impact on consumption when the DM selects $(E, 1)$, it is unambiguously beneficial. This conclusion is obviously weakened, or even reversed, if unsuccessful search activity is costly (e.g., because it exposes the DM to physical harm).[51]

It follows that, in some circumstances, criminalization may be superior to taxation and to laissez-faire. This result deserves emphasis, since it is difficult to justify a policy of criminalization based on demand-side welfare considerations without adopting the nonstandard perspective that supply disruptions can avert mistakes.[52] Since it is better not to disrupt planned consumption, the case for criminalization is, ironically, strongest when enforcement is imperfect.

H. *Selective Legalization with Controlled Distribution*

Some policies permit transactions involving addictive substances in certain circumstances but not in others. Examples include a 1998 Swiss law legalizing the prescription of heroin for severe addicts and "blue laws" prohibiting alcohol sales on Sundays.

Policies of selective legalization with controlled distribution often make deliberate planning a prerequisite for availability, selectively disrupting impulsive use without disturbing planned use (assuming the hot mode only activates behaviors that target immediate consumption). This effect is potentially beneficial, if unintended. For example, with blue laws, alcoholics can make themselves less vulnerable to compulsive drinking on Sundays by choosing not to stock up in advance. These laws appear to reduce impulsive use in practice (Peter T. Kilborn, 2003; T. Norstrom and O. J. Skog, 2003).

A prescription requirement can play a similar role, provided prescriptions are filled with a lag. To represent this possibility formally, we modify our model as follows. Imagine that, in each period m, the DM must decide whether to "call in" a prescription for the substance. Taking this action makes the substance available in period $m + 1$; otherwise it is unavailable, and consumption is impossible.

With this option, the DM can always achieve the first-best outcome. Solving the

---

[50] Probabilistic consumption following the choice $(E, 1)$ changes the value function somewhat, but the results from Section IV extend to this case. See Goldstein and Kalant (1990) for evidence that drug usage declines as substances become less available.

[51] The costs of a *successful* search are part of the price effect.

[52] Though the mechanisms considered in this paper involve *stochastic* mistakes, the same conclusion would follow in a model with deterministic mistakes, for example, one in which the DM always errs by placing too much weight on the immediate hedonic reward.

dynamic programming problem with $p_s = 0$ for all s yields a deterministic consumption path. The DM can mimic this outcome by calling in his prescription in period m if and only if he consumes the substance in period $m + 1$ on the first-best path. In this way, he precommits to the first-best choice by optimally rationing himself.[53]

If the hot mode also activates behaviors that target future consumption, the preceding policy is ineffective. However, a small modification restores the first-best outcome: allow the DM to cancel irrevocably, at any point in period m, his prescription for $m + 1$. It is then optimal for him to cancel during period m while in the cold mode if and only if he does not wish to consume the substance in period $m + 1$.[54]

In more realistic settings, these policies might not permit consumers to achieve first-best outcomes. If, for example, the desirability of using a substance in period m depends upon conditions (e.g., mood) that are not resolved until the period is underway, the individual may sometimes regret failing to call in a prescription. However, the policy still weakly benefits consumers because it provides them with a tool for self-regulation without mandating its use.

Heterogeneity across individuals makes selective legalization with controlled distribution even more attractive relative to other policies. A prescription program accommodates heterogeneity by providing consumers with discretion: intentional users can continue to indulge without impediment, while unintentional users nevertheless benefit from improved self-control. In contrast, any feasible tax, subsidy, or criminal statute may be inappropriate—even harmful—for large subsets of consumers.

The policies considered in this section would be advantageous in any model where the DM makes similar types of mistakes and where he understands this proclivity. The particular stochastic mechanism discussed in this paper is not essential. Our conclusions do depend on the assumption that the government can limit resale of the substance (e.g., by requiring on-site administration) and can suppress illicit supply. Notably, selective legalization impairs black markets by siphoning off demand.

I. Policies Affecting Cue-Triggered Decision Processes

In our model, public policy can potentially help consumers by attenuating either exposure or sensitivity to cues (i.e., reducing $c(a, \omega)$ or $M(c, s, a, \omega)$ or raising M^T). Arguably, the producers of addictive substances raise the likelihood of triggering hot modes by exposing consumers to ubiquitous cues through billboards, television advertisements, product placement in stores, and so forth. Advertising and marketing restrictions of the type imposed on tobacco and alcohol may eliminate a cause of compulsive use. Restrictions on public consumption may have similar effects.

Other public policies may reduce cue-sensitivity by creating counter-cues. Brazil and Canada require every pack of cigarettes to display a prominent viscerally charged image depicting some deleterious consequence of smoking, such as erectile dysfunction, lung disease, or neonatal morbidity.[55] These counter-cues are designed to activate the cognitive control process described in Section II.

In our model, policies that reduce the likelihood of cue-triggered mistakes by removing problematic cues or establishing counter-cues unambiguously increase welfare. As with selective legalization, these policies are attractive because they are noncoercive, because they accommodate individual heterogeneity, and because they have the potential to reduce unintended use without distorting choice in the cold decision mode. Though individuals may have some ability to avoid problematic cues and create their own counter-cues, the government is arguably better positioned to do this.

---

[53] In a related analysis, Loewenstein et al. (2000) emphasize the role of "mandatory waiting periods" in a model where agents systematically overconsume durable goods.

[54] Alternatively, if the hot mode has a greater tendency to activate behaviors targeting future consumption when the planning horizon is short, one could restore (or at least enhance) the policy's efficacy simply by lengthening the lag between prescription requests and availability (e.g., calling in a prescription in period m makes the substance available in $m + k$, with $k > 1$).

[55] See ⟨http://www.hc-sc.gc.ca/hecs-sesc/tobacco/research/archive/⟩ for a description and some preliminary evidence on the effectiveness of the Canadian program.

VI. Related Literature

Existing economic theories of addiction include (i) variations on the standard model of intertemporal decision making (Becker and Murphy, 1988; Orphanides and Zervos, 1995), including generalizations that allow for random shocks and state-contingent utility (Hung, 2000; Laibson, 2001), (ii) models with "projection bias" wherein agents mistakenly assume that future tastes will resemble current tastes, but which otherwise conform to the standard model (Loewenstein, 1996, 1999; Loewenstein et al., 2001), (iii) models with present-biased preferences and either naive or sophisticated expectations (O'Donoghue and Rabin, 1999, 2000; Gruber and Koszegi, 2001), and (iv) models of "temptation" wherein well-being depends not only upon the chosen action, but also on actions not chosen (Gul and Pesendorfer, 2001a, b; Laibson, 2001). While all of these theories contribute to our understanding of addiction and share some important features with our model, none adheres to all of the central premises set forth and justified in Section II. In particular, none of these models depicts addiction as a progressive susceptibility to stochastic environmental cues that can trigger mistaken usage.

All models of rational addiction (beginning with Becker and Murphy, 1988) presuppose complete alignment of choices and time-consistent preferences, thereby denying the possibility of mistakes. Precommitments are never strictly beneficial, and a user would never state a sincere, unconditional intention to quit without following through. Stochastic environmental cues play a role in Laibson's (2001) extension, but the mechanism involves hedonic effects (cues trigger a change in taste for the substance) rather than mistakes. Laibson's framework can account for voluntary admission to rehabilitation clinics and related behaviors provided that these activities reduce the likelihood of experiencing cravings. However, it cannot account for the observation that many addicts seek in-patient treatment not because they expect to avoid cravings, but rather precisely because they anticipate cravings and wish to control their reactions. Furthermore, even in instances where entering a rehabilitation facility does reduce the likelihood of cravings (e.g., by removing environmental cues), the standard framework implies counterfactually that the addict would find the facility's program more attractive if it made the substance available upon demand (in case of cravings).

Adding projection bias to the standard model introduces the possibility that users may regard past actions as mistakes. For example, an addict may blame his initial drug use on a failure to anticipate the escalating difficulty of abstention. Coupled with state-contingent utility shocks (as in Laibson's model), projection bias could account both for the high frequency of attempted quitting (when not triggered, users underestimate the future difficulty of abstention), and the high frequency of failure (once triggered, users overestimate the future difficulty of continued abstention). However, even with projection bias, an otherwise standard decision-maker would never anticipate making mistakes in the future and sees no need for precommitments.

In models with present-biased decision-makers, choice is always aligned with the preferences prevailing at the moment when the choice is made. Even so, one can interpret present-bias as shorthand for considerations that lead to systematic mistakes in favor of immediate gratification, contrary to true (long-run) preferences (see, e.g., Gruber and Koszegi, 2001). As a model of addiction, this framework suffers from two main shortcomings. First, the decision-making bias is not domain-specific. A present-biased decision-maker mistakenly consumes *all* pleasurable commodities excessively; in this respect, there is nothing special about addictive substances. Second, the bias is always operative—it is not cue-conditioned.

In principle, one could formulate a model with a powerful, narrow-domain, cue-triggered present-bias. The resulting model (which does not appear in the literature) *would* conform to our premises; indeed, it would be nearly equivalent to our approach. Our model is somewhat simpler and more tractable than this alternative because we treat behavior in the hot mode as mechanical, whereas this present-bias approach would portray even triggered choices as optimal given well-behaved preferences. Naturally, for our model, one can say that the triggered decision-maker acts *as if* he optimizes subject to well-behaved preferences that attach enormous importance to consuming the addictive substance, but we think this "as if" representation is

unenlightening. Since the decision-maker is assumed always to consume the substance in the hot mode, and since we regard this as a mistake whenever he would behave differently in the cold mode, the representation illuminates neither choices nor welfare.

Finally, Gul and Pesendorfer (2001a, b) model addictive behaviors by defining preferences over both the chosen action and actions not chosen, thereby providing a potential role for "temptation" and a rationale for precommitment. Their axiomatic approach embraces the doctrine of revealed preference and therefore presupposes an alignment of choices and preferences, ruling out the possibility of mistakes. In addition, their model, as formulated, does not examine the role of stochastic cues in stimulating use.

VII. Final Remarks

This paper develops an economic model of addiction based on three premises: (i) use among addicts is frequently a mistake (a pathological divergence between choice and preference); (ii) experience with an addictive substance sensitizes an individual to environmental cues that trigger mistaken usage; and (iii) addicts understand their susceptibility to cue-triggered mistakes and act with some degree of sophistication. We argue that these premises find strong support in evidence from psychology, neuroscience, and clinical practice. Research indicates that addictive substances systematically interfere with the proper operation of an important process which the brain uses to forecast near-term hedonic rewards (pleasure), and this leads to strong, misguided, cue-triggered impulses that often defeat higher cognitive control. As a matter of formal mathematics, our model is tractable and involves a small departure from the standard framework. It generates a plausible mapping from the characteristics of the user, substance, and environment to dynamic behavior. It accounts for a number of important patterns associated with addiction, gives rise to a clear welfare standard, and has novel implications for public policy.

Our theory also has potentially important implications for empirical studies of addiction. It suggests that users of addictive substances may respond very differently to changes in prices, with dramatically different implications for welfare, depending on whether decisions reflect "hot" impulses or "cold" deliberation. In contrast, existing studies treat data on consumption as if it were generated by a single process.

The model could be extended in a variety of ways to improve realism and predictive power. Possibilities include: developing a more complete model of cognitive control in which future consequences may influence the likelihood of overriding HFM-generated impulses (through the threshold M^T); adding stochastic taste shocks realized at the outset of each period (to produce variation in the contingent plan chosen for each state); allowing payoffs (w_s) to depend directly on ω (to reflect the hedonic effects of cravings); allowing for imperfect information concerning an individual's susceptibility to cue-triggered mistakes; introducing partial, rather than full, self-understanding; modeling life-cycle changes (either anticipated or unanticipated) in preferences and susceptibilities resulting from aging and changes in circumstances; and modeling the long-term effects of early-life experiences.

It is natural to wonder whether the model applies not just to addictive substances, but also to other problematic behaviors such as overeating or compulsive shopping. These questions are currently the subject of study among neuroscientists and psychologists, and it is too early to say whether similar brain processes are at work.[56] Notably, people who suffer from pathological gambling, overeating, compulsive shopping, and kleptomania describe their experience as involving strong and often overwhelming cravings, they respond to cues such as stress and advertisements, and they exhibit cycles of binges and abstention.

APPENDIX A: THE DYNAMIC ECONOMY

This appendix contains additional technical details concerning the economy described in

---

[56] Some preliminary evidence suggests that there may be some connection. For example, compulsive gamblers and kleptomaniacs respond to drugs such as naltrexone which block the brain's ability to experience euphoric states; compulsive gamblers and bulimics experience sudden relapse even after many years of abstinence. See Holden (2001a) for a discussion of recent research concerning the commonalities between various behavioral pathologies and substance addiction.

Section V D and referenced in Propositions 5 and 6.

Let g denote generation, t denote age, and m denote time. Members of generation g are born in period $m = g$, and reach age t in period $m = g + t$. Let ρ denote the pure rate of time preference, and let π denote the constant single-period survival probability, so that $\delta = \rho\pi$. The size of each new generation at age $t = 0$ is $(1 - \pi)N$, where N is the constant size of the total population.

Let ψ^m denote the taxes/subsidies applied in period m, including either a tax on the addictive substance, τ^m, or a tax on rehabilitation, β^m, as well as age-specific lump-sum instruments, T^{tm}. The period-m policy determines tax-inclusive prices and incomes, from which we can compute (as described in Section III), for each generation $g \leq m$, a parameter vector $\boldsymbol{\theta}^{g,t} = (\boldsymbol{\theta}_1^{g,t}, \dots, \boldsymbol{\theta}_S^{g,t})$ with $t = m - g$ applicable in period m. An intertemporal policy Ψ assigns a policy ψ^m to each period m, and induces, for each generation, an infinite sequence of parameter vectors, $\boldsymbol{\Theta}^g = (\boldsymbol{\theta}^{g0}, \boldsymbol{\theta}^{g1}, \dots)$. Since $\boldsymbol{\theta}^{gt}$ can vary over t (in contrast to the case treated in sections III and IV), we must allow choice to vary with age as well as the addictive state. A decision rule χ maps age t and state s into a probability distribution over $\{(E, 1), (E, 0), (A, 0), (R, 0)\}$ (note that we allow for randomizations) and implies a probability $\sigma_s^t(\chi)$ of use in state s at age t. We use χ^g to denote the decision rule of generation g. The optimized value function $V_s^t(\boldsymbol{\Theta}^g)$ depends on the particular sequence of parameters confronted by generation g, and varies with age t. Since decisions are discrete, an optimal decision rule need not be unique and, indeed, is definitely not unique when it involves randomizations.

The optimized usage probabilities generate a state-transition probability matrix $\Lambda^t(\chi^g)$. For a large population of DMs starting in state 0 at age 0 and following decision rule χ^g, the population distribution across addictive states at age t is $z^t(\chi^g) = [\prod_{k=0}^{t-1} \Lambda^k(\chi^g)] \, \mathbf{z}^0$, where $\mathbf{z}^0$ is an S-dimensional vector with a 1 in the first position and zeros elsewhere.

We say that an intertemporal policy Ψ is *feasible* if there is, for each generation g, some decision rule χ^g solving the DM's choice problem given $\boldsymbol{\Theta}^g$ induced by Ψ, such that the government's budget is balanced in every pe-

riod. A steady-state policy ψ prescribes a constant tax rate, either τ or β, and constant age-specific lump-sum taxes, T^t. Each generation faces the same sequence of parameters, $\boldsymbol{\Theta} = (\boldsymbol{\theta}^0, \boldsymbol{\theta}^1, \dots)$, and $V_0^0(\boldsymbol{\Theta})$ is the lifetime discounted expected hedonic payoff for the representative individual.

APPENDIX B: PROOFS

Here we prove Propositions 3 and 4, and sketch the proofs of Propositions 1, 2, 5, and 6 to conserve space. Complete proofs are available on the *AER*'s web site.

SKETCH OF PROOF FOR PROPOSITION 1:

Sketch for parts (i-a) and (i-b).—The proof involves three steps.

Step 1: Consider $\boldsymbol{\theta}'$ and $\boldsymbol{\theta}''$ such that: (1) $\boldsymbol{\theta}_k' \neq \boldsymbol{\theta}_k''$, (2) $\boldsymbol{\theta}_i' = \boldsymbol{\theta}_i''$ for $i \neq k$, and (3) $V_s(\boldsymbol{\theta}') \geq V_s(\boldsymbol{\theta}'')$ for all s. Then (a) for all $j < k$, $V_j(\boldsymbol{\theta}') - V_j(\boldsymbol{\theta}'') \leq V_{j+1}(\boldsymbol{\theta}') - V_{j+1}(\boldsymbol{\theta}'')$, and (b) for all $j > k$, $V_j(\boldsymbol{\theta}') - V_j(\boldsymbol{\theta}'') \leq V_{j-1}(\boldsymbol{\theta}') - V_{j-1}(\boldsymbol{\theta}'')$. The argument, omitted, involves induction starting with $j = 1$ for part (a), and with $j = S$ for part (b).

Step 2: Consider $\boldsymbol{\theta}'$ and $\boldsymbol{\theta}''$ such that: (1) $\boldsymbol{\theta}_k' \neq \boldsymbol{\theta}_k''$, (2) $\boldsymbol{\theta}_i' = \boldsymbol{\theta}_i''$ for $i \neq k$, and (3) $V_s(\boldsymbol{\theta}') \geq V_s(\boldsymbol{\theta}'')$ for all s. Then (a) for $j < k$, the disposition to use in state j is weakly higher with $\boldsymbol{\theta}'$ than with $\boldsymbol{\theta}''$, and (b) for $j > k$, the disposition to use in state j is weakly lower with $\boldsymbol{\theta}'$ than with $\boldsymbol{\theta}''$. These conclusions follow from step 1, which implies that, for $j < k$ ($j > k$), the difference in continuation values following abstention and use, and hence the disincentive to use, is weakly greater (smaller) with $\boldsymbol{\theta}''$ than with $\boldsymbol{\theta}'$.

Step 3: It is easy to verify that $V_s(\boldsymbol{\theta})$ is weakly increasing in u_k^a and b_k^a, and weakly decreasing in p_k^a. Combining this with step 2 completes the proof of parts (i-a) and (i-b).

Sketch for part (i-c).—Consider two parameter vectors, $\underline{\boldsymbol{\theta}}$ and $\bar{\boldsymbol{\theta}}$, such that $\bar{b}_j^E > \underline{b}_j^E$ with all other components equal, or $\bar{p}_j^E < \underline{p}_j^E$ with all other components equal. We argue, in two steps, that the disposition to use in state j is weakly higher with $\bar{\boldsymbol{\theta}}$ than with $\underline{\boldsymbol{\theta}}$.

Step 1: (a) If $(E, 1)$ is optimal in state j with $\underline{\boldsymbol{\theta}}$, then it is optimal in state j with $\bar{\boldsymbol{\theta}}$, and (b) if $(E, 1)$ is the unique optimal choice in state j with $\underline{\boldsymbol{\theta}}$, then it is the unique optimal choice in state j

with $\bar{\theta}$. When $\bar{p}_j^E < p_j^E$ and all other components of θ and $\bar{\theta}$ are equal, (a) and (b) follow from part (ii) of the proposition. When $\bar{b}_j^E > \underline{b}_j^E$ and all other components of θ and $\bar{\theta}$ are equal, one can show that the difference in discounted continuation values following abstention and use in state j, and hence the disincentive to use, is strictly less than the increase in the immediate benefits of use, $\bar{b}_j^E - \underline{b}_j^E$. If the DM weakly prefers use to abstention with θ, he must therefore strictly prefer it with $\bar{\theta}$.

Step 2: (a) If neither $(E, 1)$ nor $(E, 0)$ are optimal choices in state j for $\bar{\theta}$, then the sets of optimal state j choices are identical with $\bar{\theta}$ and θ; (b) if either $(A, 0)$ or $(R, 0)$ is optimal in state j with $\bar{\theta}$, it is also optimal with θ. In each case, the result follows from the easily verified fact that the same value function, $V_s(\bar{\theta})$, continues to satisfy the valuation equation (1) when the parameter vector is changed from θ to $\bar{\theta}$.

From part (a) of step 1 and part (a) of step 2, the maximum disposition to use in state j is weakly greater with $\bar{\theta}$ than with θ. From part (b) of step 1 and part (b) of step 2, the minimum disposition to use in state j is weakly greater with $\bar{\theta}$ than with θ.

Now consider two parameter vectors, θ and $\bar{\theta}$, such that $\bar{u}_j^R < u_j^R$ with all other components equal. We claim that if something other than $(R, 0)$ is optimal in state j with θ, then it is also optimal in state j with $\bar{\theta}$ (from which it follows that the maximum disposition to use cannot be higher with θ); moreover, if $(R, 0)$ is not optimal in state j with θ, then the sets of optimal state j choices are identical with $\bar{\theta}$ and θ (from which it follows that the minimum disposition to use cannot be higher with θ). Analogously to step 2, these conclusions follow from the easily verified fact that the same value function, $V_s(\theta)$, continues to satisfy the valuation equation (1) when the parameter vector is changed from θ to $\bar{\theta}$.

Sketch for part (ii).—Suppose $\bar{\theta}$ coincides with θ except for p_j^E, p_j^A, u_j^A, u_j^R, and/or b_j^A (subject to the restrictions imposed by Assumptions 1 and 2). We claim that, if $(E, 1)$ is optimal in state j for θ, it is also optimal in state j for $\bar{\theta}$; moreover, if $(E, 1)$ is the unique optimum in the first instance, it is also the unique optimum in the second instance. Analogously to step 2 of part (i-c), these conclusions follow from the

easily verified fact that the same value function, $V_s(\theta)$, continues to satisfy the valuation equation (1) when the parameter vector is changed from θ to $\bar{\theta}$. Part (ii) follows directly.

SKETCH OF PROOF FOR PROPOSITION 2:

The proposition is proved by breaking each change into components, where the effect of each component is either neutral or described by Proposition 1.

To illustrate, we consider the effect of changing p_s^E on the length of the final resignation interval. Consider θ and $\bar{\theta}$ with $p_s^E \leq \bar{p}_s^E$ for all s (and all other parameters fixed). With θ, let $\underline{s}^3$ denote the first state (working backward from S) in which $(E, 1)$ is not an optimal choice; this defines the longest possible resignation interval. Let $\underline{s}_0^3 \geq \underline{s}^3$ denote the first state (working backward from S) in which something other than $(E, 1)$ is an optimal choice; this defines the shortest possible resignation interval. Remember that $\underline{s}^3$ and $\underline{s}_0^3$ may differ because the optimal choice in each state is not necessarily unique. Consider moving from θ to $\bar{\theta}$ in two steps. (1) Change from p_s^E to $\bar{p}_s^E$ for $s > \underline{s}^3$. Since $(E, 1)$ is initially optimal for all such states, this leaves all optimal choices unchanged [Proposition 1, part (ii), coupled with the observation that, when $(E, 1)$ is optimal, neither $(A, 0)$ nor $(R, 0)$ is ever optimal]. (2) Change from p_s^E to $\bar{p}_s^E$ for $s \leq \underline{s}^3$. This weakly increases the disposition to use in states $\underline{s}^3 + 1$ through S [Proposition 1, part (i-b)]. Thus, the disposition to use in all states $s > \underline{s}^3$ is weakly lower with θ than with $\bar{\theta}$. It follows that $(E, 1)$ continues to be an optimal choice in states $s > \underline{s}^3$ with $\bar{\theta}$, so the maximum final resignation interval is weakly longer with $\bar{\theta}$ than with θ. Since nothing other than $(E, 1)$ is optimal in states $s > \underline{s}_0^3$ with θ, nothing other than $(E, 1)$ can be optimal in states $s > \underline{s}_0^3$ with $\bar{\theta}$, so the minimum final resignation interval is weakly longer with $\bar{\theta}$ than with θ.

PROOF OF PROPOSITION 3:

Select any state s'. We can decompose the change from $\bar{\theta}$ to θ into two components: (i) a change from $\bar{\theta}$ to θ' derived from $w_s(e, x, a) = \bar{w}_s(e, x, a) - d_{s'}$, and (ii) a change from θ' to θ. The first change reduces u_s^a by $d_{s'}$ for all states s and actions a. This is simply a renormalization, and has no effect on choices. The second change weakly increases u_s^a by $d_{s'} - d_s$ for all

$s < s'$, which weakly reduces the disposition to use in state s' by Proposition 1 part (i-b), and weakly decreases u_s^a by $d_s - d_{s'}$ for all $s > s'$, which also weakly reduces the disposition to use in state s' by Proposition 1 part (i-a). Thus, the disposition to use in state s' weakly decreases.

PROOF OF PROPOSITION 4:

Part (i).—Consider some parameter vector $\bar{\boldsymbol{\theta}}$, and let $\underline{\boldsymbol{\theta}}$ denote the parameter vector obtained by setting $p_s^a = 0$ for all a and s, leaving all other elements of $\bar{\boldsymbol{\theta}}$ unchanged. By part (ii) of Proposition 1, continual use solves the DM's choice problem for $\bar{\boldsymbol{\theta}}$ if and only if it does so for $\underline{\boldsymbol{\theta}}$.

Part (ii).—Consider some parameter vector $\bar{\boldsymbol{\theta}}$, and suppose there is some state s' with $p_{s'}^E > 0$ such that $(E, 1)$ is not a best choice in s'. Applying (1) for $\boldsymbol{\theta} = \bar{\boldsymbol{\theta}}$ and using the fact that $(E, 1)$ is not a best choice in s', we have

$$V_{s'}(\bar{\boldsymbol{\theta}}) = \max\{(1 - \bar{p}_{s'}^E)(\bar{u}_{s'}^E + \delta V_{\max\{1,s'-1\}}(\bar{\boldsymbol{\theta}}))$$
$$+ \bar{p}_{s'}^E(\bar{u}_{s'}^E + \bar{b}_{s'}^E + \delta V_{\min\{S,s'+1\}}(\bar{\boldsymbol{\theta}})), (1 - \bar{p}_{s'}^A)$$
$$\times (\bar{u}_{s'}^A + \delta V_{\max\{1,s'-1\}}(\bar{\boldsymbol{\theta}}))$$
$$+ \bar{p}_{s'}^A(\bar{u}_{s'}^A + \bar{b}_{s'}^A + \delta V_{\min\{S,s'+1\}}(\bar{\boldsymbol{\theta}})), \bar{u}_j^R$$
$$+ \delta V_{\max\{1,s'-1\}}(\bar{\boldsymbol{\theta}})\}.$$

Since $(E, 1)$ is not a best choice in s', the first term in braces is strictly less than $\bar{u}_{s'}^E + \delta V_{\max\{1,s'-1\}}(\bar{\boldsymbol{\theta}})$; given Assumption 2, so are the other two terms. Thus, $V_{s'}(\bar{\boldsymbol{\theta}}) < \bar{u}_{s'}^E + \delta V_{\max\{1,s'-1\}}(\bar{\boldsymbol{\theta}})$. Let $\underline{\boldsymbol{\theta}}$ denote the parameter vector obtained by setting $p_s^a = 0$ for all a and s, leaving all other parameters unchanged. Since the DM could select $(E, 0)$ in s', we have $V_{s'}(\underline{\boldsymbol{\theta}}) \geq \underline{u}_{s'}^E + \delta V_{\max\{1,s'-1\}}(\underline{\boldsymbol{\theta}}) \geq \bar{u}_{s'}^E + \delta V_{\max\{1,s'-1\}}(\bar{\boldsymbol{\theta}})$, so $V_{s'}(\underline{\boldsymbol{\theta}}) > V_{s'}(\bar{\boldsymbol{\theta}})$.

SKETCH OF PROOF FOR PROPOSITION 5:

Without loss of generality, we can proceed as if, for the optimal policy, the net transfer to each cohort is zero in each period. If this is not the case, simply redefine income in state s at age t as $y_{st} = y_s - L_t$, where L_t is the net transfer received at age t; the original policy remains optimal.

Sketch for part (i).—We prove this in two steps.
Step 1: An optimal tax rate must be *weakly*

negative. To prove this, we assume that there is a strictly positive optimal tax rate and establish a contradiction by showing that this policy must be strictly inferior to ϕ (the no-tax policy).

Consider a decision rule χ (where we drop the generational superscript g because we are examining steady states) that is optimal and satisfies budget balance with the optimal policy, and any age t' at which neither use nor nonuse is a certainty from the perspective of period 0 (under the stated assumptions, there is always at least one such age). Now suppose that policy ϕ prevails but that the DM nevertheless continues to follow χ. Through a series of algebraic steps, one can show that $E_0[u'(e^{t'})|x^{t'} = 1] > E_0[u'(e^{t'})|x^{t'} = 0]$. That is, the expectation, as of age zero, of the marginal utility of nonaddictive consumption in t' is greater when conditioned on use than when conditioned on nonuse. Intuitively, use tends to occur when income is lower, and it also entails a cost.

Suppose we switch from the optimal policy to ϕ. Assume for the moment the DM continues to follow χ. From the perspective of age 0, the result is an actuarially fair redistribution across age-t realizations of (s, ω), from realizations in which the DM does not use the substance to realizations in which he does. Since $E_0[u'(e^t)|x^t = 1] > E_0[u'(e^t)|x^t = 0]$ for the last dollar redistributed, and since u is strictly concave, the transfer makes him strictly better off. Thus, his discounted expected hedonic payoff weakly increases for every age t and strictly increases for some. Reoptimizing the decision rule given ϕ reinforces this conclusion.

Step 2: The no-tax policy, ϕ, is not optimal. Intuitively, for the same reasons as in step 1, a small subsidy coupled with lump-sum transfers that achieve budget balance within each cohort and period should generate a first-order welfare improvement by creating an actuarially fair redistribution from realizations in which the DM does not use the substance to realizations in which he does. Formally, this reasoning encounters two technical issues. First, we must establish that policies with small tax rates and budget balance within each cohort and period are feasible. Allowing for randomized choices, this is accomplished through standard arguments and a routine application of the Kakutani Fixed Point Theorem. Second, any such redis-

tribution must be actuarially fair relative to probabilities associated with a decision rule that is optimal with the new policy, not with ϕ.

To deal with this second issue, we consider a sequence of tax rates, associated age-specific lump-sum taxes, and optimal decision rules with budget balance within each cohort and period, (τ_j, T_j, χ_j), with $\tau_j, T_j \to 0$, and χ_j converging to some limit χ_∞. By standard arguments, χ_∞ is optimal with ϕ. Let b_∞^t denote the likelihood of use at age t with χ_∞. Fixing the choice rule at χ_∞, a small tax τ generates per capita revenue $b_\infty^t \tau$ from a cohort of age τ. Distributing this back to the same cohort as a lump sum, and taking the derivative of the expected age-t payoff with respect to τ, we obtain $(1 - b_\infty^t)b_\infty^t(E_0[u'(e^t)|x^t = 0] - E_0[u'(e^t)|x^t = 1])$. This is zero when b_∞^t is 0 or 1, and by the same arguments as in step 1, is strictly positive for intermediate values. For large j, χ_j is arbitrarily close to χ_∞, so holding the choice rule fixed at χ_∞, a switch from ϕ to (τ_j, T_j) creates a redistribution that is *almost* actuarially fair for the probabilities implicit in χ_∞. We therefore know that redistribution is almost neutral for t such that $b_\infty^t \in \{0, 1\}$, and strictly beneficial for t such that $b_\infty^t \in (0, 1)$. Accordingly, there exists j sufficiently large such that the expected present value of the DM's payoff is higher with (τ_j, T_j) than with ϕ, assuming he chooses χ_∞. Reoptimizing for (τ_j, T_j) reinforces this conclusion.

Sketch for part (ii).—The argument parallels that given for part (i), except we use the fact that $E_0[u'(e^{t'})|x^{t'} = 0] > E_0[u'(e^{t'})|x^{t'} = 1]$ when q is sufficiently small.

SKETCH OF PROOF FOR PROPOSITION 6:

First consider small subsidies. The argument generally parallels step 2 of the sketch for Proposition 5, part (i). Take any sequence of rehabilitation tax rates, associated age-specific lump-sum taxes, and optimal decision rules with budget balance within each cohort and period, (β_j, T_j, χ_j), with $\beta_j < 0$, $\beta_j \to 0$, $T_j \to 0$, and χ_j converging to some limit χ_∞. Let B_∞^t denote the likelihood of rehabilitation at age t with χ_∞. Fixing the choice rule at χ_∞, a small tax β generates per capita revenue $B_\infty^t \beta$ from a cohort of age t. Distributing this back to the same cohort as a lump sum, and taking

the derivative of the expected age-t payoff with respect to β, we obtain $(1 - B_\infty^t)B_\infty^t(E_0[u'(e^t)|a^t \neq R] - E_0[u'(e^t)|a^t = R])$. This equals zero when $B_\infty^t \in \{0, 1\}$, and it is strictly negative for $B_\infty^t \in (0, 1)$ (under the conditions stated in the proposition, the DM chooses R only in the highest state reached with positive probability in t; rehabilitation therefore occurs when income is lower, and it entails a cost greater than q, so the expected marginal utility of nonadditive consumption must be greater when conditioned on rehabilitation than when conditioned on no rehabilitation).

We evaluate the change from (ϕ, χ_∞) to (β_j, T_j, χ_j) in three steps. First, change the hot-mode probabilities to those prevailing under (β_j, T_j), leaving everything else constant. Second, change the policy from ϕ to (β_j, T_j), still holding the choice rule fixed at χ_∞. Third, reoptimize, changing the choice rule to χ_j. The third change is obviously weakly beneficial, as is the first (with $\beta_j < 0$, the lump-sum transfers are negative, so, under Assumption 3, the hot-mode probabilities fall). Now consider the second step. For large j, (β_j, T_j, χ_j) is arbitrarily close to (ϕ, χ_∞), so the hot-mode probabilities are almost unchanged, and we compute expected utility using almost the same probabilities as with (ϕ, χ_∞); moreover, holding the choice rule fixed at χ_∞, a switch from ϕ to (τ_j, T_j) creates a redistribution that is *almost* actuarially fair for the probabilities implicit in (ϕ, χ_∞). Thus, $\beta_j(1 - B_\infty^t) B_\infty^t(E_0[u'(e^t)|a^t \neq 0] - E_0[u'(e^t)|a^t = R])$ approximates the period-t welfare effect. From the concluding sentence of the previous paragraph, we therefore know that the second step improves the DM's expected discounted payoff for sufficiently large j.

Now consider small taxes. For policy ϕ, let s^* denote the earliest state in which rehabilitation is an optimal choice (recall that it is the unique optimal choice in s^*), and let $s' \in \{1, \dots, s^* - 1\}$ denote a state in which $(E, 1)$ is not a best choice (both states are referenced in the proposition). For sufficiently small $\beta > 0$, one can show that, for all t, s^* is also the earliest state in which rehabilitation is an optimal choice (and that it is the unique optimal choice in s^*), and $(E, 1)$ is not a best choice in s'. This means that, for any budget-balancing optimal decision rule χ_β, there is at least one t in which both rehabilitation and no rehabilitation occur

with strictly positive probability; for any such t, the same considerations as above imply $E_0[u'(e^t)|a^t = R] < E_0[u'(e^t)|a^t \neq R]$, where we take expectations assuming the policy ϕ is in place, but the hot-mode probabilities associated with β prevail, and the DM continues to follow χ_β.

We evaluate the elimination of a small tax in three steps. First, eliminate the tax (and associated lump-sum transfers) without changing the hot-mode probabilities, and keeping the choice rule fixed at χ_β. This is strictly beneficial (it creates an actuarially fair redistribution from realizations with no rehabilitation to realizations with rehabilitation; from the preceding paragraph, we know this is strictly beneficial for the t at which both types of realizations occur with positive probabilities, and neutral otherwise). Second, reoptimize the decision rule; this is weakly beneficial. Third, change the hot-mode probabilities to those prevailing with the policy ϕ and reoptimize the decision rule; under Assumption 3, this is also weakly beneficial.

REFERENCES

Becker, Gary and Murphy, Kevin. "A Theory of Rational Addiction." *Journal of Political Economy*, 1988, *96*(4), pp. 675–700.

Behara, A. and Damasio, H. "Decision-making and Addiction. Part I. Impaired Activation of Somatic States in Substance Dependent Individuals When Pondering Decisions with Negative Future Consequences." *Neuropsychologia*, 2002, *40*(10), pp. 1675–89.

Behara, A., Dolan, S. and Hindes, A. "Decision-making and Addiction. Part II. Myopia to the Future or Hypersensitivity to Reward?" *Neuropsychologia*, 2002, *40*(10), pp. 1690–1705.

Berke, J. P.; et al. "Drug Addiction and the Hippocampus." *Science*, 2001, *294*(9), p. 1235–39.

Bernheim, B. Douglas and Rangel, Antonio. "Toward Welfare Analysis with Fallible Decision-Makers," in Peter Diamond and Hannu Vartiainen, eds., *Economic institutions and behavioral economics*. Princeton: Princeton University Press, 2005 (forthcoming).

Berridge, K. "Food Reward: Brain Substrates of Wanting and Liking." *Neuroscience and Biobehavioral Reviews*, 1996, *20*(1), pp. 1–25.

Berridge, K. "Pleasure, Pain, Desire, and Dread: Hidden Core Processes of Emotion," in D. Kahneman, E. Diener, and N. Schwarz, eds., *Well-being: The foundations of hedonic psychology*. New York: Russell Sage Foundation Publications, 1999, pp. 525–57.

Berridge, Kent and Robinson, Terry. "What Is the Role of Dopamine in Reward: Hedonic Impact, Reward Learning, or Incentive Salience?" *Brain Research Review*, 1998, *28*(3), pp. 309–69.

Berridge, Kent and Robinson, Terry. "Parsing Reward." *Trends in Neurosciences*, 2003, *26*(9), pp. 507–13.

Bolla, J. I.; Cadet, J. L. and London, E. D. "The Neuropsychiatry of Chronic Cocaine Abuse." *Journal of Neuropsychiatry and Clinical Neuroscience*, 1998, *10*(3), pp. 280–89.

Brauer, L. H.; et al. "Dopamine Ligands and the Stimulus Effects of Amphetamine: Animal Models versus Human Laboratory Data." *Psychopharmacology*, 1997, *130*(1), pp. 2–13.

Brauer, L. H.; et al. "Haloperidol Reduces Smoking of Both Nicotine-Containing and Denicotenized Cigarettes." *Psychopharmacology*, 2001, *159*(1), pp. 31–37.

Center for Disease Control. "Smoking-Attributable Mortality and Years of Potential Life Lost—United States, 1990." *Morbidity and Mortality Weekly Report*, 1993, *42*(33), pp. 645–48.

Chaloupka, F. and Warner, K. "The Economics of Smoking," in J. Newhouse and D. Cutler, eds., *Handbook of health economics*. Amsterdam: North-Holland, 2001, pp. 1539–67.

Cohen, J. D. and Blum, K. I. "Reward and Decision." *Neuron*, 2002, *36*(2), pp. 193–98.

Diagnostic and statistical manual of mental disorders. Washington, DC: American Psychiatric Association, 1994.

Di Chiara, G. "Drug Addiction as Dopamine-Dependent Associate Learning Disorder." *European Journal of Pharmacology*, 1999, *375*(1–3), pp. 13–30.

Gardner, Eliot and David, James. "The Neurobiology of Chemical Addiction," in Jon Elster and Ole-Jorgen Skog, eds. *Getting hooked: Rationality and addiction*. Cambridge: Cambridge University Press, 1999, pp. 93–136.

Gawin, F. H. "Cocaine Addiction: Psychology

and Neurophysiology." *Science*, 1991, *251*(5001), pp. 1580–86.

Goldstein, A. *Addiction: From biology to drug policy*, 2nd Ed. New York: Oxford University Press, 2001.

Goldstein, A. and Kalant, H. "Drug Policy: Striking the Right Balance." *Science*, 1990, *249*(4976), pp. 1513–21.

Grim, J. W.; et al. "Neuroadaptation: Incubation of Cocaine Craving after Withdrawal." *Nature*, 2001, *412*(6843), pp. 141–42.

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

Gul, Faruk and Pesendorfer, Wolfgang. "Temptation and Self-Control." *Econometrica*, 2001a, *69*(6), pp. 1403–35.

Gul, Faruk and Pesendorfer, Wolfgang. "A Theory of Addiction." Unpublished Paper, 2001b.

Harris, J. E. *Deadly choices: Coping with health risks in everyday life*. New York: Basic Books, 1993.

Holden, C. "'Behavioral' Addictions: Do They Exist?" *Science*, 2001a, *294*(5544), pp. 980–82.

Holden, C. "Zapping Memory Centers Triggers Drug Craving." *Science*, 2001b, *292*(5519), p. 1039.

Hser, Y. I.; Anglin, D. and Powers, K. "A 24-Year Follow-up Study of California Narcotics Addicts." *Archives of General Psychiatry*, 1993, *50*(7), pp. 577–84.

Hser, Y. I.; Hoffman, V.; Grella, C. and Anglin, M. D. "A 33 Year Follow-up of Narcotics Addicts." *Archives of General Psychiatry*, 2001, *58*(7), pp. 503–08.

Hung, Angela. "A Behavioral Theory of Addiction." Unpublished Paper, 2000.

Hyman, Steven and Malenka, Robert. "Addiction and the Brain: The Neurobiology of Compulsion and Its Persistence." *Nature Reviews: Neuroscience*, 2001, *2*(10), pp. 695–703.

Jentsch, J. D. and Taylor, J. R. "Impulsivity Resulting from Frontostriatal Dysfunction in Drug Abuse: Implications for the Control of Behavior by Reward-Related Stimuli." *Psychopharmacology*, 1999, *146*(4), pp. 373–90.

Kaczmarek, H. J. and Kiefer, S. W. "Microinjections of Dopaminergic Agents in the Nucleus Accumbens Affect Ethanol Consumption but not Palatability." *Pharmacology, Biochemistry, and Behavior*, 2000, *66*(2), pp. 307–12.

Kelley, A. E. "Neural Integrative Activities of Nucleus Accumbens Subregions in Relation to Learning and Motivation." *Psychobiology*, 1999, *27*, pp. 198–213.

Kilborn, Peter T. "Hard-Pressed States Try Sunday Liquor Sales." *New York Times*, 19 May 2003, Section A, p. 1.

Krawczyk, D. C. "Contributions of the Prefrontal Cortex to the Neural Basis of Human Decision Making." *Neuroscience Biobehavioral Reviews*, 2002, *26*(6), pp. 631–64.

Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 1997, *112*(2), pp. 443–77.

Laibson, David. "A Cue-Theory of Consumption." *Quarterly Journal of Economics*, 2001, *116*(1), pp. 81–119.

Loewenstein, George. "Out of Control: Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes*, 1996, *65*(3), pp. 272–92.

Loewenstein, George. "A Visceral Account of Addiction," in Jon Elster and Ole-Jorgen Skog, eds., *Getting hooked: Rationality and addiction*. Cambridge: Cambridge University Press, 1999, pp. 235–64.

Loewenstein, George; O'Donoghue, Ted and Rabin, Matthew. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics*, 2003, *118*(4), pp. 1209–48.

MacCoun, R. and Reuter, P. *Drug war heresies: Learning from other vices, times, and places*. Cambridge: Cambridge University Press, 2001.

Massing, Michael. *The fix*. Los Angeles: University of California Press, 2000.

McAuliffe, W. E. "A Test of Wikler's Theory of Relapse: The Frequency of Relapse Due to Conditioned Withdrawal Sickness." *International Journal of Addiction*, 1982, *17*(1), pp. 19–33.

Miron, J. and Zwiebel, J. "The Economic Case against Drug Prohibition." *Journal of Economic Perspectives*, 1995, *9*(4), pp. 175–92.

National Institute on Alcohol Abuse and Alcoholism. "Economic Perspectives in Alcoholism Research." *Alcohol Alert*, National Institutes of Health, No. 51.

National Institute on Drug Abuse. *The economic costs of alcohol and drug abuse in the United*

States, 1992, NIH Publication No. 98-4327. Bethesda: National Institutes of Health, 1998.

Nestler, E. J. "Molecular Basis of Long-Term Plasticity Underlying Addiction." *Nature Reviews Neuroscience*, 2001, *2*(2), pp. 119–28.

Nestler, E. and Malenka, Robert. "The Addicted Brain." *Scientific American*, 2004, *290*(3), pp. 78–85.

Norstrom, T. and Skog, O. J. "Saturday Opening of Alcohol Retail Shops in Sweden: An Impact Analysis." *Journal of Studies on Alcohol*, 2003, *64*(3), pp. 393–401.

O'Brien, C. "Experimental Analysis of Conditioning Factors in Human Narcotic Addiction." *Pharmacological Review*, 1975, *27*(4), pp. 533–43.

O'Brien, C. "A Range of Research-Based Pharmacotherapies for Addiction." *Science*, 1997, *278*(5335), pp. 66–70.

O'Donoghue, Ted and Rabin, Matthew. "Doing It Now or Later." *American Economic Review*, 1999, *89*(1), pp. 103–24.

O'Donoghue, Ted and Rabin, Matthew. "Addiction and Present-Biased Preferences." Unpublished Paper, 2000.

O'Donoghue, Ted and Rabin, Matthew. "Optimal Sin Taxes." Unpublished Paper, 2004.

Office of National Drug Control Policy. *What American users spend on illegal drugs*. Publication No. NCJ-192334. Washington, DC: Executive Office of the President, 2001a.

Office of National Drug Control Policy. *The economic costs of drug abuse in the United States, 1992–1998*. Publication No. NCJ-190636. Washington, DC: Executive Office of the President, 2001b.

Olds, J. and Milner, P. "Positive Reinforcement Produced by Electrical Stimulation of Septal Area and Other Regions of Rat Brain." *Journal of Comparative and Physiological Psychology*, 1954, *47*(6), pp. 419–27.

Orphanides, Athanasios and Zervos, David. "Rational Addiction with Learning and Regret." *Journal of Political Economy*, 1995, *103*(4), pp. 739–58.

Pauly, Mark V. "Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection." *Quarterly Journal of Economics*, 1974, *88*(1), pp. 44–54.

Pecina, S.; et al. "Pimozide Does Not Shift Palatability: Separation of Anhedonia from Sensorimotor Suppression by Taste Reactivity." *Pharmacological Biobehavioral Reviews*, 1997, *58*(3), pp. 801–11.

Pickens, R. and Harris, W. C. "Self-Administration of d-Amphetamine by Rats." *Psychopharmacologia*, 1968, *12*(2), pp. 158–63.

Robbins, T. W. and Everitt, B. J. "Interaction of Depaminergic System with Mechanisms of Associative Learning and Cognition: Implications for Drug Abuse." *Psychological Science*, 1999, *10*(3), pp. 199–202.

Robins, L. "Vietnam Veterans' Rapid Recovery from Heroin Addiction: A Fluke or Normal Expectation." *Addiction*, 1994, *88*(8), pp. 1041–54.

Robins, L.; Davis, D. and Goodwin, D. "Drug Use by U.S. Army Enlisted Men in Vietnam: A Follow-up on their Return Home." *American Journal of Epidemiology*, 1974, *99*(4), pp. 235–49.

Robinson, T. and Berridge, K. "The Neural Basis of Drug Craving: An Incentive-Sensitization Theory of Addiction." *Brain Research Reviews*, 1993, *18*(3), pp. 247–91.

Robinson, T. and Berridge, K. "The Psychology and Neurobiology of Addiction: An Incentive Sensitization View." *Addiction*, Supplement 2000, *2*, pp. 91–117.

Robinson, Terry and Berridge, Kent. "Addiction." *Annual Reviews of Psychology*, 2003, *54*, pp. 25–53.

Rolls, E. T. "The Orbitofrontal Cortex and Reward." *Cerebral Cortex*, 2000, *10*(3), pp. 284–94.

Schultz, W. "Predictive Reward Signal of Dopamine Neurons." *Journal of Neurophysiology*, 1998, *80*(1), pp. 1–27.

Schultz, W. "Multiple Reward Signals in the Brain." *Nature Reviews Neuroscience*, 2000, *1*(3), pp. 199–207.

Schultz, W.; Dayan, P. and Montague, P. R. "A Neural Substrate of Prediction and Reward." *Science*, 1997, *275*(5306), pp. 1593–99.

Shaham, Y.; Erb, S. and Steward, J. "Stress-Induced Relapse to Heroin and Cocaine Seeking in Rats: A Review." *Brain Research Reviews*, 2000, *33*, pp. 13–33.

Shalev, U.; Grimm, J. W. and Shaham, Y. "Neurobiology of Relapse to Heroine and Cocaine Seeking: A Review." *Pharmacological Reviews*, 2002, *54*(1), pp. 1–42.

Shefrin, H. M. and Thaler, Richard. "An Eco-

nomic Theory of Self-Control." *Journal of Political Economy*, 1981, *89*(2), pp. 392–406.

Stewart, J. and Wise, Roy. "Reinstatement of Heroin Self-Administration Habits: Morphine Prompts and Naltrexone Discourages Renewed Responding after Extinction." *Psychopharmacology*, 1992, *108*(1–2), pp. 79–84.

Trosclair, A.; Huston, C.; Pederson, L. and Dillon, I. "Cigarette Smoking Among Adults—United States, 2000." *Morbidity and Mortality Weekly Report*, 2002, *51*(29), pp. 642–45.

Ungless, M.; Whistler, J.; Malenka, R. and Bonci, A. "Single Cocaine Exposure In Vivo Induces Long-Term Potentiation in Dopamine Neurons." *Nature*, 2001, *411*(6387), pp. 583–87.

U.S. Census Bureau. *Statistical abstract of the United States*. Washington, DC: U.S. Government Printing Office, 2001.

Vorel, S.; et al. "Relapse to Cocaine-Seeking after Hippocampal Theta Burst Stimulation." *Science*, 2001, *292*(5519), pp. 1175–78.

Wachtel, S. R. "The Effects of Acute Haloperidol or Risperidone on Subjective Responses to Methamphetamine by Healthy Volunteers." *Drug Alcohol Dependency*, 2002, *68*(1), pp. 23–33.

Wagner, F. A. and Anthony, J. C. "From First Drug Use to Drug Dependence: Developmental Periods of Risk for Dependence upon Marijuana, Cocaine, and Alcohol." *Neuropsychopharmacology*, 2002, *26*(4), pp. 479–88.

Watanabe, Masataka. "Coding and Monitoring of Motivational Context in the Primate Prefrontal Cortex." *Journal of Neuroscience*, 2000, *22*(6), pp. 2391–2400.

Wickelgren, Ingrid. "Getting the Brain's Attention." *Science*, 1997, *276*(5335), pp. 35–37.

Wise, Roy. "The Brain and Reward." in J. M. Liebman and S. J. Cooper, eds., *The neuropharmacological basis of reward*. New York: Oxford University Press, 1989, pp. 377–424.

Wyvell, C. L. and Berridge, K. C. "Intra-accumbens Amphetamine Increases the Conditioned Incentive Salience of Sucrose Reward: Enhancement of Reward Wanting without Enhancing Liking or Positive Reinforcement." *Journal of Neuroscience*, 2000, *20*(21), pp. 8122–30.