

NBER WORKING PAPER SERIES

A THEORY OF GOVERNMENT REGULATION OF ADDICTIVE BADS:
OPTIMAL TAX LEVELS AND TAX INCIDENCE FOR CIGARETTE EXCISE TAXATION

Jonathan Gruber
Botond Koszegi

Working Paper 8777
<http://www.nber.org/papers/w8777>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2002

We thank Alan Auerbach, Peter Diamond, David Laibson, Sendhil Mullainathan, as well as seminar participants at the NBER's Public Economics program and the University of Kentucky for discussion. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2002 by Jonathan Gruber and Botond Koszegi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Theory of Government Regulation of Addictive Bads:
Optimal Tax Levels and Tax Incidence for Cigarette Excise Taxation
Jonathan Gruber and Botond Koszegi
NBER Working Paper No. 8777
February 2002
JEL No. H2, I1

ABSTRACT

The traditional normative analysis of government policy towards addictive bads is carried out in the context of a “rational addiction” model, whereby the only role for government is in correcting the external costs of consumption of such goods. But available evidence is at least as consistent, if not more so, with an alternative where individuals are “time inconsistent” about decisions such as smoking, having a higher discount rate between this period and the next than between future periods. We develop this time inconsistent model, and show that this alternative formulation delivers radically different implications for government policy towards smoking. Unlike the traditional model, our alternative implies that there is a role for government taxation of addictive bads even if there are no external costs; we estimate that the optimal tax on cigarettes is \$1 or more higher than that implied by the traditional model. And we estimate that cigarette excise taxes are much less regressive than previously believed, and indeed for most parameter values are progressive, since lower income groups are much more price elastic and therefore benefit more from the commitment device provided by higher excise taxes.

Jonathan Gruber
Department of Economics
MIT
E52-355
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
gruberj@mit.edu

Botond Koszegi
UC Berkeley

One of the major sources of debate among policymakers at all levels of government is the appropriate role of governments in regulating addictive bads. This is an important topic because a large share of consumer spending is on activities which potentially fit this classification. This category certainly includes smoking, illicit drug use, and (at least high levels of) alcohol consumption, but may also include activities such as risky driving and poor eating habits.

The traditional economic approach to such activities has been to treat consumers as “rational addicts”, to use the term of Becker and Murphy (1988). Their seminal article codified what had become the standard approach among economists to thinking about regulation of addictive bads. In their model, consumption of addictive bads is governed by the same decisionmaking process as is consumption of all other goods. Consumers trade off the utility gains from consuming the good against the costs of doing so, and as rational forward-looking agents they recognize that those costs include the damage that they are doing to themselves through consumption, as well as the additional future damage to which they are driving themselves by consuming more of an addictive good.

In this standard revealed preference framework, there is no rationale for government regulation of addictive bads other than interpersonal externalities. Just as the government has no cause, absent market failures, for interfering with revealed preference in the realm of non-addictive goods, there is no reason to take addictiveness per se as a call to government action, if individuals are pursuing these activities “rationally”. It is this framework which implicitly underlies the well known efforts of Manning et al. (1991) and others to measure the external costs to society of drinking and smoking. These estimates, which are frequently cited and influential in debates over excise taxation, suggest that the optimal tax rate for cigarettes in particular is fairly low, since the net external costs of smoking are small.

In this paper we consider an alternative formulation of consumption of addictive bads. Our model follows that of Becker and Murphy, with one exception: we allow agents to be time inconsistent in their smoking decisions. Laboratory evidence on preferences uniformly indicates that individuals use lower discount rates in evaluating future intertemporal tradeoffs, relative to the discount rate that they use in evaluating similar tradeoffs today. For example, in smoking decisions, the agent might want to enjoy her cigarette today, but would prefer to exercise self-control tomorrow. Since she will have similar preferences for immediate rewards in the future, there is a conflict between the intertemporal selves. This kind of time inconsistency has been modeled as quasi-hyperbolic discounting by Laibson (1997) and O’Donoghue and Rabin (1999a), and it has

been applied in the context of savings decisions (Laibson 1997, Laibson, Repetto, and Tobacman 1998, O'Donoghue and Rabin 1999b), retirement decisions (Diamond and Kőszegi 1998), and even growth (Barro 1999).

The goal of this paper is to explore the implications of applying quasi-hyperbolic discounting to the analysis of government regulation of addictive bads. Although our theoretical model is general, we focus in particular on the case of smoking. We do so because the available evidence, reviewed below, suggests that smoking decisions are better modeled in the time inconsistent framework than in the time consistent one.

We begin, in Part I, with some background on smoking, addiction modeling, and time inconsistency. Part II then introduces the model, and discusses the implications for optimal taxation. The key insight of this analysis is that, with time inconsistent agents, the presence of “internalities” due to smoking justify non-zero optimal taxes even in the absence of externalities. These parts summarize our earlier work on this topic (Gruber and Kőszegi 2001), and extend that earlier analysis to consider a variety of questions around optimal policy design towards addictive bads.

We then turn to a new question: how does time inconsistency affect the distributional implications of excise taxation? Standard incidence analysis, in which a marginal increase in taxes is equivalent to a negative income effect due to the higher prices the agent has to pay, is invalid in our model. Taxation also affects how the agent's self-control problem plays out. In particular, a price-induced decrease in consumption may be good for the agent, because it softens the overconsumption due to the desire for immediate gratification. Our model adjusts for this theoretically, and under simplifying conditions yields a simple multiplicative adjustment to the standard incidence measure.

The theoretical framework allows us to reexamine a classic argument against excise taxation of bads such as cigarettes. Since lower income groups spend a larger share of their incomes on addictive bads, such taxes are viewed as regressive. When taxation plays a role as a self-control device for time inconsistent agents, the possibility that lower income individuals have a lower short-term discount factor (β), and the fact that they are more price sensitive, can mitigate or reverse this conclusion.

On the other hand, as we highlight, this self-control adjustment does not unambiguously improve the progressivity of excise taxes. Several of the reasons typically offered for why the poor smoke more than the rich would actually lead this adjustment to increase regressivity. If the poor have lower values of life, or if smoking does less damage to them (shortens their life by fewer years), then

the future damage from smoking, and thus the self-control adjustment we develop, is less important to them. Surprisingly, the same is the case if the poor have a lower long-term as opposed to short-term discount factor (δ); thus, if the poor are more impatient, the impact on regressivity depends critically on whether this impatience arises through a lower long term discount factor (which would raise regressivity) or more short term time inconsistency (which would lower regressivity).

We therefore, in Part III, turn to calibrations of our model. We first calibrate the optimal tax model, and show that the optimal corrective tax is at least \$1, and likely much higher. We also show that, so long as the poor do not value their lives very little relative to the rich and are not much more (long run) impatient, this adjustment reduces and most likely reverses the conclusion that cigarette taxes are regressive. For even modest levels of time inconsistency, we find that excise taxes can be progressive, delivering a larger benefit to lower than to higher income groups.

1 Background

1.1 Addiction and The Case for Government Intervention

There has been a long-standing interest in the economics community in modeling the consumption of addictive goods. Until the mid-1980s, most of this literature modeled addiction as habit formation, whereby past consumption of the addictive good increases taste for current consumption. In a pathbreaking article, Becker and Murphy (1988) explored the dynamic behavior of the consumption of addictive goods, and pointed out that many phenomena previously thought to have been irrational are consistent with optimization according to stable preferences. In the Becker and Murphy model, individuals recognize the addictive nature of choices that they make, but may still make them because the gains from the activity exceed any costs through future addiction. In this “rational addiction” framework, individuals recognize the full price of addictive consumption goods: both the current monetary price, and the cost in terms of future harm and addiction.

Rational addiction has subsequently become the standard approach to modeling consumption of goods such as cigarettes. This standard has been reinforced by a sizeable empirical literature, beginning with Chaloupka (1991) and Becker, Grossman and Murphy (1994), which has tested and generally supported the key empirical contention of the Becker and Murphy (1988) model: that consumption of addictive goods today will depend not only on past consumption but on future consumption as well. More specifically, this literature has generally assessed whether higher prices next year lead to lower consumption today, as would be expected with forward-looking addicts. The

consistent findings across a variety of papers that this is the case has been cited in the widespread use of this framework for modeling addiction.

The key normative implication of the Becker and Murphy (1988) model is that the optimal regulatory role for government is solely a function of the interpersonal externalities induced by smoking. Since smoking, like all other consumption decisions, is governed by rational choice, the fact that smokers impose enormous costs on themselves is irrelevant; it is only the costs they impose on others that gives rise to a mandate for government action.

There is a large literature that is devoted to measuring the magnitude of these externalities. Key contributors include Manning et al. (1989, 1991) and Viscusi (1995), and the literature is nicely reviewed in Chaloupka and Warner (1998) and Evans, Ringel and Stech (1999). The estimates of Manning et al. suggest that the net externalities associated with smoking are only about 16 cents per pack in 1986. This low-sounding estimate reflects the fact that the increased health costs imposed by smokers on others in group insurance and public programs are offset by their premature death, which reduces the costs of health insurance for the elderly under Medicare and of defined benefit pensions, most notably Social Security. The Congressional Research Service updated this analysis to 1995 and estimated externalities of 37 cents per pack (in current dollars), which is well below the average level of state and federal excise taxation (76 cents per pack, according to Gruber (2001)). Even this low figure is the subject of considerable controversy, as Viscusi (1995) claims that smoking actually generates net positive externalities for society.

However, these estimates may be too low for a number of reasons. The first is the health costs of secondhand smoke, arising through increased lung cancer and (even more importantly) cardiac disease risk through exposure to the smoking of others. The size of the health costs of second hand smoke are quite ambiguous and controversial. While Viscusi (1995) claims that there is little credible evidence of significant costs associated with second hand smoke, Chaloupka and Warner (1998) suggests that such costs may amount to as much as 70 cents per pack. A second issue involves the case of pregnant women. Smoking leads to an increased incidence of low birthweight babies, which imposes both short-run costs of medical care and long run costs of special education. Evans, Ringel and Stech (1999) present a detailed calculation showing that these costs may amount to 42 to 72 cents per pack. Of course, both of these cases raise the difficult issue of what share of costs to the family from second-hand smoke (the predominant source of the damage from second hand smoke) and to one's children from low birthweight should be considered internal and what share external.

Another consideration for policy-makers is the incidence of tobacco taxation; given an overall government revenue requirement, excise taxation must be compared to other taxes in composing a distributionally attractive revenue raising package. In fact, however, excise taxes on cigarettes are very regressive. This is illustrated in Table 1, which is tabulated from the 1997-98 Consumer Expenditure Survey (CEX), a nationally representative survey which provides the highest quality and most comprehensive micro-data on consumption in the U.S. We use data from these two years of the CEX, and measure real (in 1999 dollars) spending on cigarettes, real spending on all goods, and real income. Since the data on income in the CEX are of poor quality, particularly at the bottom of the income distribution, we compare spending on cigarettes from the CEX to incomes measured from the Current Population Survey (CPS). So, for example, the cigarette expenditures as a fraction of income in the bottom quartile is the mean expenditure on cigarettes in that quartile from the CEX divided by the mean income in that quartile from the CPS.

The first column shows the distribution of cigarette expenditures as a share of income, by income quartile. By this measure, cigarette taxes are indeed quite regressive. Cigarette expenditures amount to 3.2% of incomes in the bottom quartile, but only 0.4% of incomes in the top income quartile.

This seeming inequity is much smaller than it appears when considered from the lifetime perspective, however. The lifetime burden of excise taxes is typically much smoother than the annual incidence, since income levels vary more over a lifetime than smoking levels (Poterba 1989). We consider lifetime incidence in two ways in the remaining columns of Table 1. The first is to follow Poterba (1989) and to use cigarette expenditures as a share of consumption expenditures (the appropriate proxy for lifetime income in the life cycle model). The problem with this approach is that if the lifecycle model does not hold in practice, as most empirical analyses suggest is the case for the bulk of the U.S. population, then consumption is not an adequate proxy for lifetime income. We therefore consider another proxy as well, educational attainment. In this case we divide the sample into four attainment groups: High School dropouts; High School graduates with no college education; those with some college but who did not obtain a bachelor's degree; and those who are college graduates.

In either case, the tobacco share difference across income groups is much smaller. The expenditures of the bottom group are now only about four times as large as those of the top group, rather than eight times as large as in the first column. Nevertheless, even using these permanent income measures, excise taxes are more regressive than many other forms of revenue raising.

A third consideration in setting excise taxes on cigarettes is the optimal distribution of commodity taxes. According to the classic Ramsey formulation, commodities should be taxed in inverse proportion to their elasticities. We will ignore this consideration in the current paper, as it is derived from a qualitatively different government optimization problem. While the Ramsey problem consists of raising an externally given amount of revenue with the least distortion, we are concerned with using taxes to *eliminate* distortions.

Thus, the case for extending the excise tax on tobacco further appears quite weak, based at least on traditional economic grounds. The average tax level of 76 cents per pack already exceeds most estimates of the externalities associated with smoking, particularly if most of the damage through second hand smoke and low birthweight infants is internalized through family decision-making. And cigarette taxes are much more regressive than other means of revenue raising. But we will argue below that under the alternative (more justified) time inconsistent version of addiction modeling, there is a quite compelling case for higher excise taxes.

1.2 The Case for Time Inconsistency in Smoking

There are four types of evidence for time inconsistency in smoking behavior. The first is laboratory experiments. Laboratory experiments document overwhelmingly that consumers are time inconsistent (Ainslie 1992, Ainslie and Haslam 1992, Thaler 1981, for example). In experimental settings, consumers consistently reveal a lower discount rate when making decisions over time intervals further away than for ones closer to the present, raising the specter of inter-personal conflict over decisions that have implications for the future.

The second is calibrating real world behavior against models with and without time inconsistency, to assess which type of model does the best job of explaining observed patterns. For example, Angeletos et al. (2001) show that a hyperbolic discounting model fits observed consumption and savings patterns much better than an exponential one; in particular, the pattern of low liquid but high illiquid wealth holdings is consistent with the self-control problems inherent in time inconsistent models. And Della Vigna and Malmendier (2001) show that the behavior of health club members, such as paying a flat fee (rather than a per use charge which almost always adds up to less ex post), is best explained by time inconsistent models.

The third is an econometric test in Gruber and Mullainathan (2001). Drawing on the model developed below, they argue that one means of empirically distinguishing time inconsistent agents from time consistent agents is the impact of cigarette taxation on their measured well-being. Time

consistent smokers will be made worse off by cigarette taxation, by the standard arguments that underlie the Becker-Murphy model. But time inconsistent agents, as we model here, can be made better off by higher taxes, as they provide the self-control device the agents demand. Gruber and Mullainathan use data on self-reported well-being from the General Social Survey, matched to information on cigarette excise taxes, to show that higher levels of excise taxes raise reported well-being among smokers, but not among others, which provides some empirical support for the time inconsistent model.

This is only a limited set of evidence, and much more is needed before the time inconsistent model will be accepted as the appropriate formulation of preferences. But it is important to note that there is *no* evidence, psychological or other, that supports time consistent preferences over these time inconsistent ones. This suggests that alternative formulations such as the one we develop in this paper be taken seriously, particularly given the radically different implications for government policy we show below.

Since smoking is a short-term pleasure, and the psychological evidence indicates that time inconsistency is most prevalent with short horizons, this formulation should be especially fruitful in the context of addictive bads such as smoking. Our fourth source of evidence in favor of time inconsistency comes from this particular application.

Two key features distinguish time consistent and time inconsistent agents. The first is the use of *commitment devices* or *self-control techniques*. We distinguish a self-control device from an alternative technology for smoking cessation, quitting aids: whereas quitting aids decrease the disutility from not smoking, self-control devices lower the utility from smoking. Time consistent decisionmakers might use a quitting aid, but in general they will not use a self-control device—with time consistency, lowering the utility of an undesired alternative is irrelevant for decisionmaking. But for some types of time inconsistent agents (what we label below *sophisticated* agents, who recognize their own time inconsistency), self-control devices are valued as a means of combating one's own time inconsistent tendencies.

In the relatively small medical literature on self-initiated attempts at quitting smoking, the voluntary use of self-control devices figures prominently. People regularly set up socially managed incentives to refrain from smoking by betting with others, telling them about the decision, and otherwise making it embarrassing to smoke (Prochaska, Crimi, Lapsanski, Martel, and Reid 1982). Various punishment and self-control strategies for quitting are also widely studied in controlled experiments on smoking cessation (Miller 1978, Murray and Hobbs 1981, and see Bernstein 1970

for a variety of “aversive stimulus” techniques), and they are recommended by both academic publications (Grabowski and Hall 1985) and self-help books (CDC various years). In one study, for example, subjects tore up a dollar bill for every cigarette they smoked above their given daily limit, and reduced that limit gradually. Presumably, these experiments are incorporating self-control devices because they are seen as the best option for helping individuals quit smoking, as could be the case if individuals were time inconsistent.

A second feature that distinguishes time consistent agents from time inconsistent agents is an inability to actualize predicted or desired future levels of smoking. The former phenomenon is specific to a class of hyperbolic discounters whom we label *naive* below, in that they do not understand that they cannot make consistent plans through time.

In fact, unrealized intentions to quit at some future date are a common feature of stated smoker preferences. According to Burns (1992), eight of ten smokers in America express a desire to quit their habit. Unfortunately, these desires can be interpreted in a number of ways, and we are not aware of any evidence for adults on their specific predictions or intentions about future smoking behavior. For youths, however, there is clear evidence that they underestimate the future likelihood of smoking. For example, among high school seniors who smoke, 56 percent say that they will not be smoking 5 years later, but only 31 percent of them have in fact quit five years hence. Moreover, among those who smoke more than 1 pack/day, the smoking rate five years later among those who stated that they would be smoking (72 percent) is actually *lower* than the smoking rate among those who stated that they would not be smoking (74 percent) (U.S. Department of Health and Human Services 1994).

One might question the relevance of this argument; after all, as noted earlier, there is an empirical literature which presumes to test the rational addiction model by documenting forward looking behavior by consumers. Gruber and Köszegi (2001) accept that smokers are forward-looking; indeed, they develop even more convincing evidence than previous papers that this is the case.¹ But, as they highlight, forward looking behavior by smokers also arises in time inconsistent models, so that this evidence does not necessarily support the Becker and Murphy model and

¹In particular, Becker, Grossman and Murphy (1994), among others, rely on smokers knowing about unannounced price changes as much as one year in advance, assess the impact of these changes on sales and not consumption, and develop results which are very sensitive to the assumptions of their estimation strategy. In contrast Gruber and Koszegi (forthcoming) study the reaction of monthly cigarette consumption to tax increases which are legislatively enacted but not yet effective, and find a much more robust reaction of current consumption to known future price increases.

its normative implications. That is, this empirical literature tests one premise of the Becker and Murphy model, showing that smokers are not fully myopic, but not the second key premise, time consistency. And the evidence discussed above suggests that time inconsistency may be a better assumption to adopt for modeling the smoking decision.

Much of this evidence is consistent with at least two other recent models of individual decision-making as well, Bernheim and Rangel (2001) and Gul and Pesendorfer (forthcoming, 2000). Both of these papers rule out an effective tax policy essentially by assumption, therefore arriving at different policy implications than those developed below.² There has been little attempt to distinguish these models, although the evidence in Gruber and Mullainathan (2001) is more consistent with our formulation than with these alternatives. It seems likely that behavior is some combination of the three models, and the best model depends on the addictive good in question.³ In this paper, we focus our attention on the consequences of hyperbolic discounting, acknowledging that future work should try to appropriately distinguish and combine these alternative approaches.

2 Government Policy Towards Addictive Goods

2.1 A Consumption Model for Addictive Goods

To address the main questions of government intervention in the market for an addictive good, we first need to introduce a model of individual choice in these products. We use the basic model that we introduced in our earlier paper (Gruber and Köszegi 2001). This model marries what are in our opinion the two most important aspects of the consumption of addictive goods. First, as documented in physiological and psychological studies, the instantaneous utility from consuming a good such as cigarettes depends in specific ways on past consumption of the same good. The literature distinguishes two of these effects: reinforcement and tolerance. Reinforcement means that having consumed a cigarette in the past increases the “craving” for one today—it increases the

²Bernheim and Rangel (2001) assume that in the “visceral” state when the agent overconsumes the drug, she is not price sensitive. Gul and Pesendorfer (forthcoming, 2000), whose model is driven by disutility from temptation, assume that the agent is tempted equally strongly by the drug for all possible prices, as long as she has enough wealth to pay for it. These models do not deny a role for government intervention; for example, Bernheim and Rangel suggest that the optimal policy would be to force individuals to place advance orders for addictive substances, in order to mitigate the influence of the “visceral” state.

³For example, Bernheim and Rangel’s (2001) model emphasizes the visceral factors in behavior, and assumes that sometimes the drive to take drugs outpaces all rational optimization. This may be true for drugs like cocaine and heroin, but probably not for cigarettes.

marginal utility of consumption. And tolerance refers to the notion that our body tends to “get used to” drugs, and therefore it gets harder and harder over time to achieve the same “high”—one needs to consume more and more of the good to reach the same level of utility. Both of these can be captured in an instantaneous utility function of the form

$$U_t = U(a_t, c_t, S_t) = v(a_t, S_t) + u(c_t). \quad (1)$$

a_t and c_t are the levels of consumption of the addictive and ordinary goods, respectively. Reinforcement and tolerance are incorporated into the utility function through the dependence of v on S_t , the so-called stock of past consumption, a measure of the amount of past consumption of a_t . S_t evolves according to

$$S_{t+1} = (1 - d)(S_t + a_t), \quad (2)$$

where $0 < d < 1$ is the depreciation rate of the stock. Reinforcement means simply that $v_{aS}(a_t, S_t) > 0$, while the formal equivalent of tolerance is $v_S(a_t, S_t) < 0$.

The second important aspect of addictive goods consumption concerns the nature in which instantaneous utilities are integrated into a global utility function. Suppose we are in a T -period model. A *time consistent*, exponential discounter agent makes decisions at time t according to the discounted utility function

$$\sum_{i=0}^{T-t} \delta^i U_{t+i}. \quad (3)$$

We will contrast this type of discounting with the alternative recently popularized by Laibson (1997), quasi-hyperbolic discounting. For quasi-hyperbolic discounters, discounted utility becomes

$$U_t + \beta \sum_{i=1}^{T-t} \delta^i U_{t+i}. \quad (4)$$

β and δ are usually assumed to be between zero and one. This formulation is intended to capture the idea that decisionmakers might have self-control problems regarding the consumption of addictive goods. Under this specific form of *time inconsistency*, the discount factor between consecutive future periods (δ) is larger than between the current period and the next one ($\beta\delta$). Thus, the agent is “impatient” when faced with a choice between today and tomorrow, but she would like to “become patient” in the future. The arguments for this formulation were presented above.

In each period of her life, a quasi-hyperbolic discounter wants to be impatient for that period and become patient later. This creates a conflict between the current self and the future ones—which is the essence of the self-control problem. There are two extreme assumptions one can

make about how this conflict is resolved. One can assume naivité, in which the decisionmaker is completely unaware that she will be impatient again in the future. Such an agent maximizes the utility function 4 in each period, and changes her plans over and over again. At the other extreme, one can assume sophistication, where the agent realizes that she will change her mind and she behaves strategically according to this. Formally, the successive intertemporal selves play the subgame-perfect equilibrium in an extensive-form game, in which the choice variable of each self is consumption in that period.

To complete the setup of the model, let p_t be the price of the addictive good in period t , and denote income in period t by I_t . To make it easier to solve our model for sophisticated quasi-hyperbolic discounters, we assume that there are no savings: the income I_t is consumed in each period.

Based on simple examples and intuition, there is reason to believe that relaxing our liquidity constraint assumption would actually *increase* the optimal tax. If the agent was allowed to borrow from future income, she would take advantage of the opportunity, and would do so partly to finance consumption of the addictive good. Thus, the tax would have to offset the tendency to borrow in addition to the tendency to overconsume the drug. But the agent's use of savings and addictive goods consumption as a self-control device considerably complicates the analysis,⁴ and we have not solved for the dynamics of our model with savings.

For a more thorough introduction of this model, see Gruber and Kőszegi (2001); for a contrast of sophistication and naivité, see O'Donoghue and Rabin (1999a). In an important paper on this topic, O'Donoghue and Rabin (2000) specifically consider the effects of sophistication and naivité for addiction. Their main insight is that in an environment where agents face occasional strong urges to consume the addictive good, naifs are more likely to get severely addicted. Intuitively, naifs start consuming the good, believing that they will quit soon, but then procrastinate in doing so. O'Donoghue and Rabin's model uses a discrete-choice setup, however, which allows only a crude analysis of price effects and taxation, so that they are unable to analyze the implications for optimal government policies.

In Gruber and Kőszegi (2000), we solve for the Euler equations for the three types of consumers, time consistent agents, naive hyperbolic discounters, and sophisticated hyperbolic discounters. For

⁴In particular, the agent may want to deprive herself of savings, so that the future self cannot afford the addictive good. Conversely, if, for example, being addicted decreases the agent's tendency to spend too much, she may want to 'overaddict' herself to improve her savings behavior.

future reference, we quote the sophisticates' first-order condition here, which is valid as long as utility functions and equilibrium strategies are differentiable:

$$\begin{aligned} & v_a(a_t, S_t) - p_t u'(c_t) = \\ = & (1-d)\delta \left[\left(1 + (1-\beta) \frac{\partial a_{t+1}}{\partial S_{t+1}} \right) (v_a(a_{t+1}, S_{t+1}) - p_{t+1} u'(c_{t+1})) - \beta v_s(a_{t+1}, S_{t+1}) \right]. \quad (5) \end{aligned}$$

2.2 Efficient Policies Towards Addictive Goods

The goal of this section is to examine the implications of the quasi-hyperbolic framework for optimal government policy. Of course, the question arises why we consider only government interventions to combat self-control problems. If a sophisticated agent had access to an effective private self-control device, she would take advantage of it, reducing the value of a government intervention. However, we find it unlikely that fully effective self-control devices can be found in this context. Market-provided self-control mechanisms are probably undercut by the market mechanism itself: although firms have a financial incentive to provide self-control to agents, other firms have a financial incentive to break it down. For example, if a firm developed a self-control shot that causes pain when the consumer smokes, another firm would have an incentive to develop a drug that relieves these effects for agents who temporarily want to get rid of their commitment. Other problems arise in contracting setups. If there are ex post gains to be made, the future self might want to renegotiate today's contract. But even if there are none, there is an ex post incentive to cheat on the contract: smoking is hard to verify in court. This leaves us with privately provided self-control mechanisms like betting with others or becoming involved in situations where it is very difficult to smoke, but these mechanisms are likely to run into similar enforcement problems to those discussed above.

Government intervention in the market has its own problems. The large taxes implied by our calibrations might create considerable black market and smuggling activity that undercut tax policy, just as incentives to cheat undercut private self-control devices. Our calibrations, therefore, indicate the optimal level of corrective increase in the price of cigarettes, assuming that this corrective increase can be enforced—either through public or through private channels. We will use the shorthand of optimal taxation, but really this is the optimal combination of enforceable taxation and private correction.

As in our earlier paper (Gruber and Köszegi 2001) and similarly to Becker and Murphy (1988), we restrict attention to quadratic utility functions; this is done only to simplify the analysis. Thus,

the functions v and u take the form

$$\begin{aligned} v(a_t, S_t) &= \alpha_a a_t + \alpha_s S_t + \frac{\alpha_{aa}}{2} a_t^2 + \alpha_{as} a_t S_t + \frac{\alpha_{ss}}{2} S_t^2 \\ u(c_t) &= \alpha_c c_t \end{aligned} \tag{6}$$

where α_a, α_{as} , and α_c are positive and α_s, α_{aa} , and α_{ss} are negative. The key parameter is $v_{aS}(a_t, S_t) = \alpha_{as}$, which measures the effect of past consumption on the marginal utility of current consumption. To ensure that first-order conditions are sufficient to find the equilibrium in our model, we assume that $U(a_t, c_t, S_t)$ is strictly concave, that is, we suppose its Hessian is negative definite.

In this case, it is very easy to prove by backward induction that a_t is linear in S_t : $a_t = \lambda_t S_t + \mu_t$, where λ_t and μ_t are constants. The following theorem, quoted from Gruber and Köszegi (2001), establishes that for a general class of parameter values, marginal propensities to addiction are stationary far from the end of the horizon, giving a consumption function $a_t = \lambda^{*s} S_t + \mu^{*s}$. Similarly, $a_t = \lambda^{*n} S_t + \mu^{*n}$ for naifs and $a_t = \lambda^{*TC} S_t + \mu^{*TC}$ for time-consistent agents.

Theorem 1 *Suppose $\beta \geq \frac{1}{2}$, $U(a_t, c_t, S_t)$ is strictly concave, and $p_t = p$, a constant. Then, $\lim_{j \rightarrow \infty} \lambda_{T-j} = \lambda_s^*$, where λ^{*s} is given as the unique solution on the interval $(-1, \frac{\alpha_{as}}{-\alpha_{aa}})$ of*

$$\lambda^{*s} = -1 + \frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \delta(1-d)^2[(1+(1-\beta)\lambda^{*s})(\alpha_{aa}\lambda^{*s} + \alpha_{as}) - \beta\alpha_{as}\lambda^{*s} - \beta\alpha_{ss}]}. \tag{7}$$

Similarly, for β close enough to 1, the constant in the agent's consumption function is (approximately) stationary far from the end of the horizon (Gruber and Köszegi 2000).

Notice that in order to carry out our analysis, we assume that there are no income effects. This makes the analysis considerably simpler without affecting the basic conclusions we reach. Without income effects, the price of the addictive good has no influence on λ^{*s} . However, the social planner can effectively choose μ_t by setting the appropriate taxes. If the social planner had a full set of instruments—she could set the tax in each period at will—then she could achieve any combination μ_t 's in the different periods. We assume that the government is restricted to a set of tax instruments such that $(\mu_0, \mu_1, \dots, \mu_t, \dots) \in M \subset \mathfrak{R}^\infty$, where M is closed according to the product topology.

As in any model where different socially relevant actors have different tastes, a discussion of optimal government policy must start with the setup of the social welfare function. In the context of hyperbolic discounting, these actors are not separate individuals, but different intertemporal incarnations of the same individual. For this application, we take the agent's long-run preferences

as those relevant for social welfare maximization. Since the discount factor β applies to all future periods, if the representative agent were to vote in a tax change today that is instituted starting tomorrow, these are the preferences she would use in choosing the new tax rate.⁵ A hyperbolic discounter would like to exert self-control on her future selves, but not on her current self, so she would actually prefer a cigarette tax that becomes effective only in the next period. Therefore, a social welfare function that maximizes the exponentially discounted sum of instantaneous utilities may be the most relevant for policy analysis.

An alternative is to maximize the current self's utility, as when the tax chosen is effective today. This would decrease the optimal tax somewhat, because the agent would only want to tax her future selves, not her current self. As long as β is not very small and δ is sufficiently large (so that the agent cares about the future to a significant extent), this makes little difference in our analysis.

Thus, the social planner solves

$$\begin{aligned} \max_{(\mu_t) \in M} \sum_{t=0}^{\infty} \delta^t [v(\lambda^{*s} S_t + \mu_t, S_t) + \alpha_c (I_t - p(\lambda^{*s} S_t + \mu_t))] \\ \text{s.t. } S_0, S_{t+1} = (1-d)(S_t + \lambda^{*s} S_t + \mu_t) \end{aligned} \quad (8)$$

In our earlier paper (Gruber and Köszegi 2001), we consider the case when the government is restricted to a tax τ that is constant over time, and which is assumed to be passed through to consumers one-for-one.⁶

This leads to μ_t being constant over time, and consequently $M = \{(\mu, \mu, \mu, \dots) | \mu \in \mathfrak{R}\}$. Then, the first-order condition for the optimal tax policy becomes

$$(1-\beta) \sum_{t=0}^{\infty} \delta^t (v_a(a_t, S_t) - (p+\tau)\alpha_c) = \beta \frac{1}{1-\delta} \cdot \frac{1-\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*s})} \tau \alpha_c. \quad (9)$$

It is easy to show that the optimal tax is positive: the derivative of 8 with respect to μ at $\mu = \mu^{*s}$ can be written in the form

$$\sum_{t=0}^{\infty} \delta^t (v_a(a_t, S_t) - p\alpha_c + \delta(1-d)V_S^s(S_{t+1})), \quad (10)$$

where $V^s(S_t)$ stands for the exponentially discounted utility from leaving stock S_t and consuming according to the sophisticated consumption function from then on. A hyperbolic discounter agent

⁵Interestingly, if the tax was instituted later than in the next period, the agent would actually vote in a *higher* tax than what is implied by her long-term discount factor. This is due to the intertemporal substitutability of taxes, discussed below.

⁶Past experience suggests that cigarette excise taxes are passed through more than one-for-one; see Gruber (2001) for a discussion of this evidence. This suggests that the optimal taxes presented below are overstated.

solves $v_a(a_t, S_t) - p\alpha_c + \beta\delta(1-d)V_S^s(S_{t+1}) = 0$, and since by assumption $V_S^s(S_{t+1})$ is negative, the above derivative is negative for $\beta < 1$. Therefore the optimal μ is lower than μ^{*s} , and consequently the optimal tax is greater than zero. Since we assumed that there are no consumption externalities, this conclusion contrasts with the implications of a rational addiction framework: that optimal taxes on cigarettes should depend only on their externalities. In our formulation, optimal taxes depend on their internalities as well, and so are positive so long as such internalities exist.

There are several important additional implications of this framework for the optimal taxation of addictive goods. *First*, since the optimal tax is positive for $\beta < 1$, the left-hand side of the equation 9 is positive. But $v_a(a_t, S_t) - (p + \tau)\alpha_c > 0$ means that the addiction is harmful—higher consumption lowers utility from future periods. Therefore, at least in an average sense, the optimal tax is not so large so as to make the addiction harmless on the margin. The reason is that the tax is there to correct a marginal self-control problem. If there was no self-control problem (on average), there would be nothing to correct—the agent’s different intertemporal selves would not disagree, so the losses to consuming more would be second-order. But then, the selves would be consuming too little, since their private costs are higher than the social costs due to the tax.

Second, we can determine the dependence of the optimal tax on S_0 . For any t , the total derivative of $v_a(a_t, S_t)$ with respect to S_t is $\alpha_{aa}\lambda^{*s} + \alpha_{as}$. This is greater than zero because $\lambda^{*s} < \lambda_T = -\frac{\alpha_{as}}{\alpha_{aa}}$. Therefore the derivative of the left-hand side of equation 9 with respect to S_0 is positive, and so the optimal tax is increasing in the level of initial addiction. The reason for an optimal tax that is increasing in S_0 follows from the nature of the tax. It is solely a “self-control tax,” a tax that is intended to aid in overcoming the agent’s self-control problem. As such, it has to increase as the self-control problem becomes more serious. And in our model the marginal harm done by smoking more, and so the self-control problem, increases with S_0 .

This observation has two relevant implications for policy. There is a potential problem in implementing the optimal policy over time. Suppose the representative agent has reached her steady state level of consumption before the tax is implemented. In the first period of the new regime—with the tax in place—the agent cuts back on her consumption, leading to a decrease in the stock of past consumption. Thus, in the next period, the government wants to decrease its tax rate. In other words, the government’s tax policy is *time inconsistent*—without commitment on the government’s part, the agent will not believe that the government will carry through its plan, undermining the effectiveness of the policy. The key intuition for this result is that the government’s tax is not only intended to curb current, but also *past* consumption. Anticipating that the tax will

be high (and consequently that she will smoke less) tomorrow leads the agent to cut back today, in order to reduce her craving tomorrow. But once tomorrow rolls around, the agent has already cut back, and there is less reason to impose the high tax. It is sufficient to “scare” the agent into quitting—it is not necessary to actually carry out the threat.⁷

Relatedly, suppose that the government can set a different tax rate in period 0 and the rest of the periods (and commit to this entire policy). If the agent has already reached her steady-state consumption level when the tax is instituted, then the period-0 tax will be higher than the long-term tax. This occurs exactly because the tax breaks the agent’s habit, so $S_1 < S_0$. For distributional reasons, the traditional prescription is exactly the opposite. According to this traditional view, a high up-front tax for cigarettes is undesirable because it hurts addicted consumers too much, and these consumers tend to be disproportionately poor. But we show below that these distributional considerations may be misplaced.

Third, we can assess the implications for the time pattern of excise taxes if the government is constrained from implementing the optimal pattern. Suppose that the government has a free hand in choosing τ_t , but has some restrictions in the choice of taxes in future periods. If the government sets τ_t in period t , and the choice does not affect future taxes, then the first-order condition for the optimal choice of τ_t is a variant of condition 9:

$$\begin{aligned} \tau_t \alpha_c &= -(1 - \beta) \left(\sum_{k=1}^{\infty} \delta^k (1 - d)^k (1 + \lambda^{*s})^{k-1} [(v_a(a_{t+k}, S_{t+k}) - p\alpha_c)\lambda^{*s} + v_s(a_{t+k}, S_{t+k})] \right) \\ &\quad - \beta \delta (1 - d) \lambda^{*s} \sum_{i=0}^{\infty} \delta^i (1 - d)^i (1 + \lambda^{*s})^i \tau_{t+1+i} \alpha_c. \end{aligned} \quad (11)$$

From this formulation, we can easily prove that if the good is addictive ($\lambda^{*s} > 0$), taxes in different periods of time are *substitutes*. Suppose we increase $\tau_{t'}$ for some $t' > t$. First, this has the direct effect of decreasing the right-hand side of the above expression. In addition, it decreases consumption at t' as well as before and after, decreasing both a and S in each of these periods. But notice that the derivative of the right-hand side with respect to $a_{t''}$ (where $t'' > t$) is $-\alpha_{aa}\lambda^{*s}$, while the derivative with respect to $S_{t''}$ is $-(\alpha_{as}\lambda^{*s} + \alpha_{ss})$. The former expression is clearly positive, while the latter is positive by the concavity of the utility function, so an increase in $\tau_{t'}$ unambiguously decreases τ_t . And if we increase $\tau_{t'}$ for some $t' < t$, the decrease in consumption pushes down S_t ,

⁷Once consumption of cigarettes has reached the (optimal) steady state, the time inconsistency problem disappears. In steady state, new smokers enter the population as others leave, but the overall distribution remains the same. Therefore, there is no reason to change the tax. But according to our calibrations below, we are far from the optimal steady state, so time inconsistency is likely to be a problem for a while.

decreasing the optimal tax in period t .

Once again, this is a consequence of the intertemporal complementarity in consumption levels. In order to induce the agent to consume less in period $t + 1$, the government can use other tools than the tax in that period. It can increase the tax in period t as well, since that will decrease the stock of past smoking inherited in period $t + 1$. In addition, and perhaps more surprisingly, it can also instead increase the tax in period $t + 2$, changing the agent's expectations about the future: since she expects to have to cut back in period $t + 2$, she starts by cutting back in period $t + 1$. Thus, just as the standard Ramsey model would suggest that we "overtax" complements of untaxed goods such as leisure, the complementarity of smoking in different periods suggests that we "overtax" smoking in some periods if taxation is more difficult in others. For example, if it is politically infeasible to optimally tax adult smoking, but politically attractive to reduce youth smoking, the substitutability of tax policy over time suggests over-taxing youth smoking.

This point raises a novel justification for the type of "clean air regulations" that exist in almost all states in the U.S., banning smoking in public places such as private workplaces, restaurants, bars, or grocery stores. The traditional justification for such regulations is reducing exposure to second-hand smoke, but this justification is somewhat weakened by the lack of convincing evidence of health damage from such exposure (Viscusi, 1995). But if our model is written over space—which is isomorphic to writing it over time as long as agents are in a given place at a given time—and if it is impossible to regulate smoking in private homes, then over-regulating (or even restricting) smoking in observable public places can be justified.⁸ That is, since smoking is a complement across spaces, then taxes across spaces are substitutes. So if we are undertaxing in one space, we may want to overtax in other spaces. In that sense, clean air restrictions are not only protecting others from the smoker; they are also part of a package of tools to protect the smoker from himself.

Fourth, similar conclusions hold for naive quasi-hyperbolic discounters, but there is an important qualitative distinction. In that case, the tax not only corrects a self-control problem, but also a misperception problem—the agent is wrong in predicting her future behavior. This has a very important implication for optimal taxation. Whereas in the sophisticated case taxation that eliminates all harmful consumption can never be justified, even if the good is very addictive and people

⁸If the decisionmaker's chosen location depends on the regulatory policy, our result has to be qualified somewhat. For example, if the regulation leads the agent to go outside to smoke, there is no one-to-one map between space and time, even for the purposes of our model. Of course, to the extent that such substitution is possible, "clean air regulations" are less effective.

have severe self-control problems (low β 's), it might be the best policy for the naive case. Thus, if adults are sophisticated, then only taxation, and never a complete ban, of smoking is justified for them. If youths are naive, a ban might be appropriate, even if they are aware of the negative health consequences of smoking. Interestingly, however, there is a force that tends to make the optimal tax lower for naifs. Since $\lambda^{*n} \geq \lambda^{*s}$ (lemma 2 in the appendix), reducing consumption today will have a greater effect on future consumption for naifs than for sophisticates. Thus, taxation has a more detrimental effect on future revenues, decreasing the optimal tax. We return to this effect in the calibration section.

Finally, it is interesting to note how this framework differs formally from one with interpersonal externalities. It might seem that each self's failure to take full account of the future harm of consumption is equivalent to an externality, and the theory (of Pigouvian taxation) developed for that case should apply. This is not the case. Even when the government has access to a full set of instruments (that is, it can set taxes differently in different periods as it wishes), there is no easy interpretation of the optimal period-specific tax as being equal to the marginal externality caused by extra consumption. More specifically, and contrary to what one might think, the tax is not equal to $1 - \beta$ times the marginal effect of a higher stock. This is true for both naifs and sophisticates, but for different reasons.

For sophisticates, the reason is the “incentive effect” discussed in Gruber and Kőszegi (2001). Even when future selves are consuming in an optimal way from the long-run self's point of view, a sophisticated self feels that they are not: since she sees higher prices than the social ones, the level of consumption that is just right given the true cost of the good seems too high for her. Therefore, she feels a need to exert control on the future selves by consuming less. This effect helps the government, and so the optimal tax is less than $1 - \beta$ times the marginal externality.

Naifs, on the other hand, are not driven by the incentive effect. However, even with optimal corrective taxes, they tend to underestimate how much they will consume in the future. By adjacent complementarity, this “optimism effect” induces them to consume less today, which again helps the social planner's cause and reduces the optimal tax.

Both of these effects depend on the intertemporal complementarity of consumption. Therefore, strikingly, the optimal tax on an addictive harmful good is *lower* than the tax on an equally harmful good that is not addictive.

2.3 Incidence of Taxes on Addictive Goods

In this section we argue that the traditional economic methods for incidence analysis are incomplete for addictive goods in the presence of time inconsistency. We also propose an alternative incidence measure, based on what we believe is the goal of incidence analysis.

Broadly speaking, the goal of incidence analysis is to determine who is “hurt” by different tax policies. For an economist, the appropriate measure for this analysis is utility—how the tax policy affects each person’s utility in society. For example, when we say that an increase in the tax on gasoline would be borne heavily by low-income consumers, we (should) mean that if the tax was instituted, their utility would be seriously affected.

So why is incidence traditionally done in terms of prices and quantities consumed? The reason is the envelope theorem. Consider a consumer facing a maximization problem of the form

$$\begin{aligned} \max_{q_1, \dots, q_n} \quad & U(q_1, \dots, q_n) \\ \text{s.t.} \quad & p_1 q_1 + \dots + p_n q_n \leq W \end{aligned} \tag{12}$$

We know that the derivative of the maximum in this problem with respect to p_1 is $-\alpha_c q_1$, where α_c is the Lagrange multiplier on the constraint. That is, for a maximizing consumer, the utility effect of a small price increase is equal to the product of the price increase, the quantity consumed, and the marginal utility of wealth. Thus, for standard utility maximizers, incidence analysis in terms of prices and quantities is justified.

The same is *almost* true for a time consistent consumer of addictive goods. For simplicity, assume that the price is constant: $p_t = p$. Then, the derivative of self t ’s discounted utility with respect to p equals

$$-\alpha_c a_t - \alpha_c \delta a_{t+1} - \alpha_c \delta^2 a_{t+2} - \dots - \alpha_c \delta^{T-t} a_T. \tag{13}$$

Once again, the utility impact of a price increase depends on the marginal utility of wealth and the amounts consumed. As for any good that is consumed in multiple periods, the utility cost of a cigarette tax depends not only on current, but on future consumption as well. The only wrinkle is that for non-addictive goods it is reasonable to say that current consumption is an unbiased measure of future consumption, whereas for addictive goods, consumption often tends to increase until the individual reaches her steady-state level of consumption.

For quasi-hyperbolic discounters, the envelope theorem does not hold in the above form, so standard incidence analysis fails to capture the utility cost of taxation. Once again we focus here

on the more complicated case of sophisticates, and mention only briefly how our results differ for naifs. In addition, as in section 2.2 we evaluate incidence in terms of the agent's discounted utility according to her long-term discount factor δ .

Some arithmetic reveals that the derivative of the agent's discounted utility with respect to p is

$$-\alpha_c \underbrace{\left(\sum_{j=t}^T \delta^{j-t} a_j \right)}_{\text{standard term}} - \underbrace{\frac{1-\beta}{\beta} \sum_{j=t}^T \delta^{j-t} \frac{\partial a_j}{\partial p} (v_a(a_j, S_j) - p\alpha_c)}_{\text{self-control adjustment}}. \quad (14)$$

The first sum in this expression is similar to what we had in the time consistent case: since the agent now has to buy her consumption at a higher price, her utility is affected by the extra cost of this. The additional term, which we call the “self-control adjustment” to incidence, is new. It captures the value time-inconsistent agents attach to the self-control tool provided by a higher price. Each self consumes “too much” from the long-run self's point of view, so the price-induced decrease in consumption increases discounted utility. This term is absent in the time consistent case because each self consumes “just the right amount” from the long-run self's point of view.

By repeated application of the sophisticates' first-order condition 5, we can rewrite a key part of the self-control adjustment:

$$\begin{aligned} v_a(a_j, S_j) - p\alpha_c &= -\beta(1-d)\delta v_S(a_{j+1}, S_{j+1}) \\ &- \beta(1-d)^2\delta^2 (1 + (1-\beta)\lambda_{j+1}) v_S(a_{j+2}, S_{j+2}) \\ &- \beta(1-d)^3\delta^3 (1 + (1-\beta)\lambda_{j+1}) (1 + (1-\beta)\lambda_{j+2}) v_S(a_{j+3}, S_{j+3}) \\ &- \dots \end{aligned} \quad (15)$$

We will use expressions 14 and 15 to study the self-control adjusted incidence of cigarette taxation. But even though our analysis is motivated by data on cigarette consumption, the methods are applicable to other goods as well. We once again restrict ourselves to the quadratic utility function setup introduced in section 2.2.

We will need the following preliminary lemma for our analysis:

Lemma 1 1. λ_t is decreasing in β for each t .

2. λ_t is decreasing in δ for each t .

3. λ_t is increasing in t .

Proof. Appendix. \square

We are now ready to discuss our main results. First, for any good that is harmful and addictive, this self-control adjustment will be positive, so that the incidence is lower than in the time consistent case. Moreover, for a wide range of parameter values this adjustment will actually be larger than the standard incidence effect, so that incidence is not only reduced but reversed; that is, taxes are beneficial on net.

Next, notice that the right-hand side of equation 15 depends on the constants λ_i^s in the agent's consumption function. Thus, the self-control adjustment is larger if $\lambda_i^s > 0$ —when the good is addictive as opposed to merely harmful. Surprisingly, for a person suffering from self-control problems, the taxation of addictive harmful goods imposes less of a burden than the taxation of goods that are *exactly as harmful* but not addictive. The reason is that the consumption of a harmful addictive good imposes two kinds of future costs; it causes direct harm and triggers an increase in future consumption, which is too high anyway. In a sense, consuming a harmful addictive good makes the future self-control problem regarding this good worse, because it increases the short-run desire to consume and pushes the costs into the future. Both of these effects are bad from a long-run perspective, and a quasi-hyperbolic discounter does not take either of them sufficiently into account.

This result is all the more striking given our earlier claim that addictiveness decreases the optimal tax. The difference is driven by the revenue side of the government's problem. Inducing the agent to reduce smoking decreases her future consumption as well, eroding the government's tax base. Incidence analysis does not take this into account, whereas optimal taxation does, creating the opposite implications of addictiveness.

The conclusion that the true burden of taxation is lower for addictive harmful goods is true for naifs as well, albeit in a somewhat different form. For naifs, the fact that future selves respond to an increase in stock by consuming more has no negative implications for perceived discounted utility—since they believe that future selves will behave optimally, such an increase in consumption is just an optimal response to current consumption. Thus, there is no sense in which naifs do not take the addictiveness of the good sufficiently into account. Rather, their problem stems from perception: since they do not know that the future selves will consume too much, they fail to realize that current consumption just exacerbates this overconsumption problem.

One difficulty with giving crisp conclusions about incidence in this model is that the marginal damage $v_S(a_t, S_t)$, through its dependence on consumption and the stock, depends on the other

parameters of the model. Our (theoretical *and* empirical) knowledge of these changes is limited, so we abstract from them in the rest of our discussion and in the calibration.

Ultimately, our interest is in exploring the implications this incidence adjustment for the regressivity of cigarette taxation. In the context of our model, there are four reasons why the magnitude of this adjustment might differ for the poor and the rich. First, the self-control adjustment in expression 14 is proportional to the price responsiveness of the agent. Thus, as long as $\beta < 1$, more price elastic consumers bear less of the “true” burden of taxation. The intuition is simple: since the agent consumes too much in each period, the price hike increases utility by restraining the overconsumption. This self-control is more effective if the agent is more responsive to price incentives. As we document below, the poor are much more price sensitive than the rich in their smoking decisions. Thus, this factor will tend to reduce the regressivity of excise taxation, all else equal.

The remaining three factors all relate to a critical question that has been unanswered by the literature on smoking: why do the poor smoke more? One reason may be that lower-income individuals smoke more mostly because they have a lower β . In this case self-control incidence adjustment is unambiguously larger for them, and therefore the regressivity of cigarette taxes is overestimated in standard analysis. The first, direct effect of a decrease in β is an increase in the multiplier $\frac{1-\beta}{\beta}$ in the second term in expression 14. The intuition is obvious: if β is lower, each self is ignoring more of the future harm she causes by smoking more, and therefore a decrease in consumption is more valuable. In addition, equation 15 indicates that there is an indirect effect as well, coming from the fact that the good is addictive. As we have noted above, the harm from current smoking is not restricted to the direct health costs of smoking, but includes the cost of induced future smoking as well. This cost is higher if β is lower, since in that case the future selves are making a worse decision. Exacerbating the problem is that λ_t^s is greater for agents with a lower β ; therefore, the effect of current consumption on future consumption, which the current self is not sufficiently taking into account, is higher.

On the other hand, another reason why the poor might smoke more is that they have the same level of time inconsistency as the rich (the same value of β), but a lower long run discount factor (a lower δ). In fact, if the poor have a lower δ , then, all else equal, the burden of taxes falls *more heavily* on them. Lowering δ has the direct effect of decreasing the right-hand side of equation 15, thus decreasing the self-control adjustment. With the future “less important,” each intertemporal

self ignores part of a less important thing, and so the tax is less desirable on self-control grounds.⁹ Lowering δ also lowers the dollar value of damage people attach to consuming cigarettes, since this damage tends to come at the end of life, further raising incidence. Though the standard incidence measure is also lower when δ is lower, the sum-total of these effects on the multiplicative adjustment is greater in relative terms. The conclusion might be surprising: if the poor's higher smoking rate comes from a smaller *short-term* discount factor (β), then a cigarette tax is not as regressive as currently believed, while if it comes from a lower *long-term* discount factor (δ), then such a tax is even more regressive than recent estimates, all else equal.

If lower-income individuals smoke more because they attach a lower value to life, once again we have to conclude that their self-control adjustment is smaller. We model a decrease in the value of life as an increase in α_s , which reduces the marginal harm from smoking. An increase in α_s does not affect λ_t^s for any t , so from equation 15, a decrease in the marginal harm of a cigarette decreases the value of the self-control gain a price increase achieves. The intuition is simply that a person who places a lower value on life cares less about self-control aimed at protecting that life, decreasing the self-control adjustment.¹⁰

Thus, in summary, the impact of the self-control adjustment to standard incidence measures is unclear. The adjustment itself will tend to lower incidence, and it will do so more for groups where smoking is more price sensitive (which is true for the poor). But on the other hand, this adjustment is reduced as δ is lower and the value of life is lower, both of which are also potentially true for the poor. In the next section, we turn to a calibration exercise which can help assess the relative importance of these offsetting influences.

3 What Difference Does It Make?

The previous sections discussed some of the qualitative theoretical implications of a hyperbolically discounted utility function for efficiency and incidence analysis. In the current section we attempt to quantify these implications, using available information and otherwise assessing the sensitivity

⁹An effect acting in the opposite direction arises from the addictiveness of the good: lowering δ increases λ_j^s . Therefore, current consumption exerts a larger influence on future consumption, and the fact that the intertemporal selves are not taking this sufficiently into account makes the tax easier to bear. However, the former (direct) effect of a decrease in δ always outweighs the latter (indirect) effect. The proof is omitted, but is available upon request.

¹⁰As we note below, a related reason that the poor may smoke more is because smoking is less damaging when life is shorter, since fewer years of life are lost. This operates in a parallel fashion to the lower value of life point.

to parameter variation.

In order to carry out the analysis, we make three simplifying assumptions in addition to working with a quadratic utility function. First, we assume that decisionmakers start off from a steady-state consumption level: a_t does not depend on t . Second, we ignore end-of-life effects operating through λ_t and the price elasticity of consumption. Third, we assume that the disutility associated with smoking, $v_S(a_t, S_t)$, is constant.^{11,12} Let h_S denote the money equivalent of the per-period future marginal utility of an extra cigarette (so it should be negative.) This accounts for the pure disutility effect of the stock, but not the impact of current consumption on future smoking decisions.

3.1 Calibrating Optimal Taxes

In this section, we attempt to calibrate the model from Section 2 in order to calculate an actual optimal tax for sophisticated agents. Starting from equation 9 for the optimal tax, substituting equation 15, and then setting $h_s = \frac{v_s(a_t, S_t)}{\alpha_c}$ for each t gives

$$(1 - \beta) \frac{\delta(1 - d)}{1 - \delta(1 - d)(1 + (1 - \beta)\lambda^{*s})} (-h_S) = \frac{1 - \delta(1 - d)}{1 - \delta(1 - d)(1 + \lambda^{*s})} \tau. \quad (16)$$

In our calibration, we will use the *combined* discounted damage of a cigarette in all future periods, $H_S = \frac{(1-d)\delta}{1-(1-d)\delta} h_S$. Rewriting the above expression

$$\tau = \frac{1 - \delta(1 - d)(1 + \lambda^{*s})}{1 - \delta(1 - d)(1 + (1 - \beta)\lambda^{*s})} (1 - \beta) (-H_S). \quad (17)$$

One difficulty with estimating the optimal tax is parameterizing H_S . Clearly, there is a lot of disutility associated with smoking that is hard to quantify, such as that from constant coughing and increased vulnerability to various illnesses. We will ignore all these, and assume that the only disutility from smoking is in the increased chance of early death. Viscusi (1993) reviews the

¹¹This assumption is not inconsistent with the assumption of addictiveness of cigarettes. There seem to be two different stocks relevant for smoking: ‘addictiveness’ stock and ‘health’ stock, with different depreciation rates. Addictiveness depreciation, which (loosely speaking) measures how long a cigarette consumed today affects craving in the future, appears quite fast. On the other hand, the health effects of that cigarette last much longer. Thus, while the addictive process might be governed by the addictiveness stock for which marginal utility is not constant, most of the important long-term effects could come from the health stock for which marginal utility is constant. For simplicity, we do not model behavior with two stocks of past consumption, but instead assume that v_S is constant.

¹²A possible non-linearity is that quitting completely has an extra health benefit relative to just reducing smoking to a low level. In as much as this is true, our incidence adjustment is understated: since the observed price elasticity in consumption is partly on the participation margin, assuming a linear v_S understates the self-control benefits of taxation.

literature on life valuation and suggests a consensus range of 3-7 million 1990 dollars for the value of a worker's life; choosing the midpoint value and expressing it in current dollars gives a figure of \$6.8 million. Presumably, this is a present discounted value for all remaining years. We assume that the average worker is 40 years old and would live to age 79 if a nonsmoker, and use a 3% discount rate to calculate the implied value of a year of life at each age from these facts. We use the fact that smokers die on average roughly 6 years earlier from Cutler et al. (2001), and compute for each age 15-73 the PDV of the cost of losing 6 years at the end of life. We then take a weighted average of these costs at each age, where the weights are the share of cigarettes smoked at each age from the May 1999 Current Population Survey Tobacco Use Supplement, a nationally representative survey of smokers. Finally, we divide this weighted average by the average number of cigarettes smoked over one's lifetime; that is, we assume that average and marginal damage is equal.¹³

At these figures, the cost in terms of life years lost per pack of cigarettes is \$35.64. Of course, this figure will vary with the agent's δ that we are considering; since costs are at the end of life, they will fall as δ falls. But it is an enormous figure for any reasonable δ , and is on the order of 100 times as large as estimates of the interpersonal externalities from smoking.

There are several offsetting biases to using this figure as an estimate of the damage per pack. This estimate is too high to the extent that smokers value their lives less than nonsmokers.¹⁴ On the other hand, this estimate is too low to the extent that we have ignored all non-mortality related damage due to smoking. Smoking not only shortens lives but lowers quality of years spent alive as well through reduced health. Moreover, our quasi-hyperbolic framework implies that the hedonic valuation estimates of a life that we are using are too low. Under hedonic analysis, life valuations are backed out of revealed preference in the market. With quasi-hyperbolic discounting, this approach is theoretically unfounded: agents are not maximizing their discounted utility, so

¹³Age 79 and the use of 6 years of reduced life reflects an averaging of effects for men and women. The average cigarettes smoked over the lifetime (19,418) is derived by subtracting the average starting age of current smokers (age 17) from age 73, multiplying by 365, and then multiplying by average cigarettes per day smoked among daily smokers (19). Note also that it is unclear whether the reference point for mortality reduction should be the first pack smoked (so that mortality reductions come from the perspective of age 79), the last pack smoked (so that they come from the perspective of age 73), or in between. We use the sum of damage over ages 73-79 as an average.

¹⁴The evidence in Viscusi and Hersch (2001) suggests that this is true, in that the compensating differentials that smokers require for risky jobs are lower than the differentials required by nonsmokers. But, as they discuss at length in their article, their estimate of the reduced value of injury for smokers, approximately 50%, reflects both the demand side role of preferences and the supply side response of the labor market to those preferences. So it is unclear how large the pure downward risk preference adjustment should be.

market behavior will in general not reflect true valuations. Unless the *only* job-related risk faced by the worker is concentrated in the current period *and* the worker is completely liquidity constrained, quasi-hyperbolic discounters will accept a compensating differential that is “too small” relative to the true long run value of their life, because they are excessively tempted by the short-run rewards from accepting the risks.

In Table 2, we present a set of calibrations for the optimal tax for a variety of values of β , δ , d and λ^{*s} . As in our previous paper (Gruber and Kőszegi 2001), we take the relevant period length to be one month. We consider a range of values of β of 0.6 and 0.9. $\beta = 0.6$ parameterizes a very serious self-control problem, albeit one we believe could well apply to smoking. Kirby and Herrnstein (1995) find weekly discount rates of 10 to 30 percent, while Thaler’s (1981) estimates indicate monthly discount rates of 20 to 30 percent. Given this evidence, $\beta = 0.9$ represents a mild self-control problem. We also assess the sensitivity of our results to varying δ between 0.9 and 0.97 (on an annualized basis); a 10% long run annual discount rate is quite substantial discounting, while 3% is more similar to that typically used in public policy analysis. Physiological and empirical evidence suggests that λ^{*s} is fairly high for smoking. Evidence is less clear on the depreciation rate. As the results are fairly linear with respect to variations in δ and β , but not so with respect to d and λ^{*s} , we show the full range of taxes for d and λ^{*s} , and for two values of β and δ each.

As Table 2 shows, the results are quite sensitive to parameter values. To fix ideas further, Gruber and Kőszegi (2001) use their data on the timing of response to tax rate changes to rule out combinations of high λ^{*s} and low d . As discussed in more detail in that paper, such a combination would imply a very slow adjustment to steady state smoking levels—even long after the price change, agents should be reducing their consumption from period to period. But they observe fairly quick adjustments to steady state smoking levels after taxes change, definitely ruling out combinations such as $d = 0.5$ and $\lambda^{*s} = 0.9$, and compatible with a combination of $d = 0.6$ and $\lambda^{*s} = 0.7$.

At $d = 0.6$ and $\lambda^{*s} = 0.7$, the optimal tax is very large for all of the combinations of β and δ shown in Table 2. The optimal tax is \$2 for $\beta = 0.9$ and $\delta = 0.97$; it rises \$9.37 for $\beta = 0.6$ and $\delta = 0.97$. Thus, lower β leads to a higher optimal tax, as the externalities rise with the extent of time inconsistency. We also find that the optimal tax is \$1.02 for $\beta = 0.9$ and $\delta = 0.9$; and it is \$4.76 for $\beta = 0.6$ and $\delta = 0.9$. Lowering δ has three effects on the optimal tax. First, a lower δ reduces the long run revenue cost to the government of driving down smoking, which raises the optimal tax. Second, a lower δ reduces the discounted distortion from addicting future selves to smoking, lowering the optimal tax; the mistakes made by future selves are less important today as

δ decreases. Third, a lower δ lowers the present discounted value of damage done at the end of life, lowering the optimal tax. Calibrations reveal that the first effect outweighs the second (e.g. when value of life effects are ignored, the optimal tax rises as δ falls), but that the third effect swamps either of the first two, leading to lower optimal taxes as δ falls.

Most importantly, in all cases, the optimal tax is very large relative to externalities. This is not surprising given that the internal damage from smoking is as much as one hundred times as large as the estimates of externalities. Thus, for even small deviations from the time consistent case, there are enormous self-control gains from higher prices of cigarettes.

The optimal tax would be lower for an alternative social welfare function, which maximizes utility according to the preferences of the self when the tax change is instituted, thus using a quasi-hyperbolically discounted welfare function. With the exponentially discounted social welfare function that we use, the social planner wants to correct every self's self-control problem, including self 0's, whereas with hyperbolic discounting self 0 wants to respect her own preferences and just correct future selves' behavior. But the difference is not large so long as δ is close to one and β is far from zero.

3.2 Calibrating Tax Incidence - Theory

We next turn to a calibration of tax incidence in our model. As with efficiency, incidence is potentially quite different in the time inconsistent formulation.

Setting $\lambda_t = \lambda^{*s}$ and $v_S(a_t, s_t) = v_S$, expression 15 becomes

$$v_a(a_j, S_j) - p\alpha_c = \beta(1-d)\delta v_S \frac{1 - [(1-d)\delta(1 + (1-\beta)\lambda^{*s})]^{T-j}}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})}. \quad (18)$$

Using the above and setting $a_t = a$, expression 14 for the utility-based measure of incidence becomes (for self 1)

$$-\alpha_c \frac{1-\delta^T}{1-\delta} a - \frac{1-\beta}{\beta} \sum_{j=1}^T \delta^{j-1} \frac{\partial a}{\partial p} \left(\beta(1-d)\delta v_S \frac{1 - [(1-d)\delta(1 + (1-\beta)\lambda^{*s})]^{T-j}}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})} \right), \quad (19)$$

which can be rewritten as

$$-\alpha_c \frac{1-\delta^T}{1-\delta} a - (1-\beta)(1-d)v_S \frac{\partial a}{\partial p} \frac{1}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})} \left(\frac{1-\delta^T}{1-\delta} \delta + \delta^T \frac{1 - [(1-d)(1 + (1-\beta)\lambda^{*s})]^{T-j}}{1 - (1-d)(1 + (1-\beta)\lambda^{*s})} \right). \quad (20)$$

Finally, this expression has a more convenient form as $-\alpha_c \frac{1-\delta^T}{1-\delta} a$ times:

$$1 - (1-\beta)(1-d) \underbrace{\frac{-v_S}{\alpha_c p} \left(-\frac{\partial a}{\partial p} \frac{p}{a} \right)}_{\eta} \underbrace{\frac{\delta}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})}}_{\text{horizon-independent effect}} \quad (21)$$

$$- (1 - \beta)(1 - d) \frac{-v_S}{\alpha_c p} \underbrace{\left(-\frac{\partial a}{\partial p} \frac{p}{a} \right)}_{\eta} \underbrace{\frac{\delta^T (1 - \delta)}{1 - \delta^T} \cdot \frac{1 - [(1 - d)(1 + (1 - \beta)\lambda^{*s})]^T}{1 - (1 - d)(1 + (1 - \beta)\lambda^{*s})}}_{\text{length-of-horizon effect}} \cdot \frac{1}{1 - (1 - d)\delta(1 + (1 - \beta)\lambda^{*s})}.$$

The quantity $-\alpha_c \frac{1 - \delta^T}{1 - \delta} a$ is the traditional measure of tax incidence in a dynamic framework—it is simply the product of consumption, the marginal utility of income, and a horizon term depending on how long the tax affects the agent. Therefore, expression 21 gives a multiplicative factor that can be used to adjust our traditional measures of the regressivity of cigarette taxes when we want to use a utility-based incidence measure for quasi-hyperbolic discounters. In other words, the difference in the magnitudes of these adjustment factors across income groups tells us the degree to which the usual incidence measures are “off” when the utility gains from extra self-control are taken into account.

The latter term in expression 21 can be decomposed into two terms, one that does not depend on T and one that does. This allows us to immediately assess the relative importance of length-of-horizon effects of incidence relative to other effects in the model. Note that this term is less than $\frac{\delta^{T-1}(1-\delta)}{(1-\delta^T)(1-(1-d)(1+(1-\beta)\lambda^{*s}))}$ times the horizon-independent term. Suppose that the average smoker’s horizon is 25 years—this smoker is 40 years old, and will live to age 65. Even for relatively large δ ’s, the above term will be small. For example, for $\delta = 0.95$, and assuming $\beta = 0.6$, $d = 0.5$, and $\lambda^{*s} = 0.9$ —the combination in our range of parameter values that maximizes the relative importance of the length-of-horizon effect—it is equal to only about 3.7% of the horizon-independent effect. In addition, the *differences* in this term between income groups will be even smaller, since the difference in life expectancy between income groups is low. Therefore, we work with a formula that ignores the length-of-horizon term.

As in our optimal tax analysis, we calculate the dollar value H_S of the extra mortality risk generated by consuming an extra pack of cigarettes. This gives the final expression for the adjustment factor:

$$1 - (1 - \beta) \frac{-H_S}{p} \underbrace{\left(-\frac{\partial a}{\partial p} \frac{p}{a} \right)}_{\eta} \frac{1 - (1 - d)\delta}{1 - (1 - d)\delta(1 + (1 - \beta)\lambda^{*s})} \quad (22)$$

With the simplifying assumptions we made for sophisticates, we can put the self-control adjustment in a convenient form for naifs as well, and we do so in the Appendix. As we show there, we get a much larger adjustment to standard incidence for naifs. The intuition for this result derives from naifs’ misperception about their future behavior combined with the addictiveness of the good. Since the good is addictive and future selves do not consume optimally, an increase in

current consumption not only decreases future utility, but sparks further increases in consumption that are also harmful. Since naifs believe that their future selves will behave optimally, they do not take this into account, whereas sophisticates do (albeit only partially). This effect is aggravated by the fact that naifs are more responsive to increases in the stock of past smoking anyway. Thus, a price increase that forces the agent to consume less is more beneficial for a naif. Another way to put this difference is that the tax acts partially like an incentive effect: by decreasing the current self's consumption, it decreases future selves' consumption as well. Since sophisticates' behavior is already shaped by the incentive effect, a tax is less beneficial for them.

Our calculations in the appendix do ignore the “optimism” effect for naifs described earlier, by making v_S constant. Although we have no way of estimating this optimism effect, we can prove that it is not as large as the incentive effect operating for sophisticates, and thereby that incidence is unambiguously lower for naifs than for sophisticates (e.g. the incidence adjustment would be even larger for naifs than what is presented below).¹⁵ Although incidence is unambiguously lower for naive agents, we should emphasize that this does not imply a higher optimal tax for them. Since $\lambda^{*n} \geq \lambda^{*s}$, breaking a naive agent's addiction has a more detrimental effect on future revenues than does breaking a sophisticate's addiction. Incidence analysis does not take this into account, but optimal tax analysis does.¹⁶

3.3 Calibrating Tax Incidence - Parameters

The calibration of incidence involves not only the preference parameters discussed above, but some new parameters as well. Moreover, when discussing distributional issues, it becomes critical to then consider the distribution of these parameters as well. In this section we bring evidence to bear on several of these underlying parameters and how they differ along the income distribution, and discuss our assumptions on the remaining parameters.

A central variable for computing incidence is the price elasticity of demand for cigarettes, and

¹⁵We can compare the optimism effect to the incentive effect by thinking of both as price increases. The incentive effect acts as a current price increase, because the agent realizes that current consumption triggers future harmful consumption. The optimism effect, on the other hand, acts as a future price increase, because the agent believes that future selves will perceive a higher health cost from smoking than they actually will. Using the results in our earlier paper, we can study how the agent responds to these price changes. It turns out that since naifs underestimate the addictiveness of cigarettes (lemma 2), the optimism effect is lower. Therefore, incidence of taxes is lower on naive agents, although we do not know exactly how much lower.

¹⁶Without the optimism effect, the optimal tax for naifs is simply $-(1 - \beta)H_S$, the Pigouvian corrective tax

in particular how it varies across income groups. Standard micro-data estimates for the elasticity of consumption center around -0.45, and the previous literature has indicated that the elasticity is much higher for lower than for higher income smokers. For example, Evans, Ringel and Stech (1999) estimate the elasticity for those with missing income (which they presume to be on average a very low income group) is -0.73, for those below median income is -0.53, while it is -0.13 for those above.

We provide updated estimates of elasticities and their distribution by matching to the CEX data from 1980 through 1998 information on cigarette prices and taxes in each state in each survey month.¹⁷ We then estimate models of cigarette expenditures as a function of price, instrumented by excise tax; as discussed by Gruber and Kőszegi (2001), such an instrumental variables strategy is required because there may be state-specific pricing that is endogenous to cigarette demand. We control in our model for a set of demographic characteristics (age, education, sex, and race of the household head; dummies for number of persons in the household), and a full set of state dummies, year dummies, and calendar month dummies.

The results of this exercise are shown in Table 3. The coefficient estimates show the impact of a one dollar price increase on expenditures; the standard errors are in parentheses. Below each estimate is the elasticity of the quantity of cigarettes consumed with respect to price implied by this consumption response, at the mean price and quantity.¹⁸ Across the full sample, each dollar price increase leads to an increase in cigarette expenditures of only 16.8 cents, for an implied elasticity of -0.66. This is larger than the traditional estimated elasticity of -0.45 for cigarette expenditures, but it is very close to the estimates using more recent data in Gruber and Kőszegi (2001).

Across groups, we see a clear pattern of higher price sensitivity for lower income, consumption, or education groups. In every case, for the bottom group expenditures *decline* as price rises, implying an elasticity of less than minus 1. For income categories, the elasticities decline monotonically as income rises, with a top elasticity that is roughly one third that of the bottom group. For consumption and education categories, the decline is monotonic to the third category, but elasticities then increase again for the top group.

While we discussed above our computation of the health damage from a pack of cigarettes, we

¹⁷Data on taxes and prices by state are from Tobacco Institute (1999). Taxes are measured monthly using information on state excise tax histories; prices are measured annually each November, so we take a weighted average of the past and future November prices in assigning a price to each month.

¹⁸The elasticity implied by the estimated response of total consumption is $(\beta/\bar{q}) - 1$, where β is the estimated coefficient and \bar{q} is the sample mean of expenditure.

now need to assess how this varies by group. This in turn has two components: differences by group in the value of a life; and differences by group in the marginal damage of smoking. The only evidence on the former comes from Viscusi and Evans (1990), who use survey evidence on the value of injuries to estimate an income elasticity of one. Applying this same income elasticity to deaths would imply that the value of life for the lowest income quartile is only 10% that of the highest income quartile.

While we will consider estimates that use this income elasticity, there are at least two strong arguments for using values of life that are much closer across groups. First, public policy goals are fundamentally predicated (at least in principle) on equal valuations of life across groups. No analysis of policy in the U.S. of which we are aware uses varying valuation of life across income or demographic groups for computing costs and benefits. Our tax is a self-control tax, so in principle it should use the agent's own valuation of their life. But, even if those valuations differ, public policy may want to correct for those differences in using more equal life values for computing optimal taxes. Of course, if this is the case, taxes are less regressive than our estimates below, even less than those that assume equal valuations of life. A tax on the addictive good not only corrects low-income agents' self-control problem, but also their relative undervaluation of life, and so is more beneficial for them.

Second, such massive differences in the value of life are inconsistent with observed smoking behavior. Evans, Ringel and Stech (1999) show that smoking rates in the top income quartile are about one-half those in the bottom income quartile. While this difference is large, it is quite small relative to what would be implied by a unitary income elasticity of value of life, assuming that all other parameters are the same across income or education groups. Suppose, for example, that the value of life for individuals in the bottom quartile is 75% of that of those in the top quartile, a number that would be large according to Viscusi and Evans' (1990) estimates. Given that we estimate the average damage from cigarettes to be \$34.86, a reasonable pair of valuations would be \$40 for the rich and \$30 for the poor. Given our elasticity estimates—which are admittedly derived from small price changes—a \$10 difference in prices should lead to a much greater difference in consumption levels than we actually observe. According to the elasticity of the top quartile (-0.387), a \$10 price increase (off the mean price in this period of \$2 per pack in our CPS data) should lead agents to cut their consumption by 194%, yet smoking is 50% as high among the highest income as the lowest income groups. Even if we accept that responses to small price changes cannot be extrapolated to make conclusions about large price changes, it is unlikely that smokers would

only cut their consumption by 50% were the price of cigarettes to increase to \$13 per pack today.¹⁹

There is, to date, little evidence on the latter question of how the marginal damage of smoking differs by income group. The differential impact of smoking by group is unclear *ex ante*. On the one hand, since lower income groups live for fewer years, there are fewer years of life lost from a shift up in the hazard curve of death. On the other hand, to the extent that the damage from smoking interacts with other disease, there could be more damage to lower income groups who are in worse health for other reasons as well. Given the lack of evidence, we subsume this point in our variation in the value of life.

There is also little evidence on the difference in preference parameters across income groups. We therefore assume that preference parameters are the same. Given the evidence presented above, we will present all calibrations for $d = 0.6$ and $\lambda^{*s} = 0.7$. In fact, we find that the basic pattern of our results is not sensitive to the values chosen for these parameters. We will consider the range of β and δ used in Table 2.

3.4 Calibrating Tax Incidence - Results

Putting together the results of the last two subsections, Table 4 shows our results for incidence. There is one panel for each income definition (current income, expenditure, and education). In each case, we use the estimated elasticities from Table 3 by group, and we use three ranges of value of life: the range implied by an income elasticity of one, where we use the ratio of income or expenditure in each category to the top category to get relative life value; a second case where we use relative life value of 0.5, 0.667, 0.883, and 1; and a final case where we assume equal life values (and, implicitly, relative health damages from smoking).

The figures in the table represent the incidence of a \$1 tax per pack of cigarettes, as a share of income in panels 1 and 3, and as a share of consumption in panel 2. *Ex ante* incidence is shown in the first column of each table, and the proportions replicate Table 1. The remainder of the results show incidence after applying our time inconsistency adjustment.

Two notes about interpretation. First, in many cases below, we find that our adjustment actually reverses incidence, which we will show as a negative burden; that is, the corrective benefits

¹⁹It is not clear whether the values of life exhibited by the agents should be multiplied by β when considering the total price of a cigarette. On the one hand, if the estimates represent the true long-run valuations of life, people will discount them by β when making decisions. On the other hand, as we have mentioned, the short-run impatience would also affect the behavior on which estimates of the value of life are based. Therefore, they already incorporate β in some way.

are so large for that parameter combination that the tax is a benefit, so that a larger negative number implies a larger benefit. Second, in some cases below we find that the ratio of the burden on the poor relative to the rich grows, but the absolute gap between the two narrows. It is not clear whether this is appropriately interpreted as a rise or a reduction in the regressivity of the tax.

For the income concept shown in panel 1 of Table 4, taxes are the most regressive ex ante, with the burden on the lowest income group almost ten times that on the highest. In the first set of columns, we show the adjusted incidence measure for $\beta = 0.9$ and $\delta = 0.97$. For this modest deviation from time inconsistency, we find that the implications of the adjustment for incidence depend on the distribution of values of life. In the case where the poor have a very low value of life relative to the rich (column 2), this adjustment worsens a ratio measure of regressivity, but narrows the absolute gap. For the medium set of life values, regressivity is reduced both relatively and absolutely. And for the case of equal life values, the tax becomes *progressive*, having a beneficial impact for the lowest group, with small positive incidence on other groups. This represents the effects of the more elastic smoking decision among the lowest income smokers.

Varying β has an enormous effect. For $\beta = 0.6$, taxes are now beneficial for every group, and much more so for the lowest income group so long as life values are not too unequal. For equal life values, the adjustment implies an enormous benefit to the lower income group, with large benefits for the next two groups and modest benefits for the top group.

The next two panels vary δ . A lower value for δ significantly mitigates the self-control adjustment, for the reasons discussed earlier. For low values of life, and $\beta = 0.9$, there is little change in regressivity. But regressivity is reduced as life values become more equal, although the tax remains regressive in the third panel even with equal values of life. But for a lower $\beta = 0.6$, taxes are progressive so long as the value of life of the poor is not too low, even with this lower $\delta = 0.9$.

When distribution is analyzed in expenditure terms, in panel 2, the impacts of the adjustment are less dramatic. Indeed, for the lowest value of life case with $\beta = 0.9$ and $\delta = 0.97$, regressivity is constant in absolute terms and worsens in ratio terms. But once again we find that so long as either we use the medium or high ratios of life values, or set β lower, we reduce or reverse regressivity.

Perhaps the strongest findings are obtained for the education distribution in panel 3. This is because the income differences are much more modest by education group, so that our low life value case (using the implied income elasticity of -1) yields much more uniform life values by group than the other two tables. In this case regressivity is everywhere reduced in both ratio and absolute terms, or reversed through more beneficial taxes on lower income groups.

Thus, while these calibrations are sensitive to parameter values and assumptions, the overall message is clear: so long as the poor do not have life values and/or marginal damage from smoking very far below the rich, and so long as their long run discount rate δ is not much lower (or, if it is, their β is lower as well), the regressivity of cigarette taxes is reduced for sophisticated time inconsistent smokers. For more equal values of life and higher degrees of time inconsistency regressivity is generally reversed, with larger benefits from cigarette taxation for lower income groups. That is, the same forces which lead us to a large optimal tax on cigarettes lead as well to the conclusion that cigarette taxes do not necessarily impose burdens on smokers, and that the benefits of taxation are largest for those low income groups that smoke the most and have the most price sensitivity. Moreover, if lower income groups are more naive, as some analysts suggest, then the conclusions here are strengthened, as we can prove that incidence is even lower on naive than on sophisticated time inconsistent consumers.

4 Conclusions

Appropriate government policy towards addictive bads such as smoking has been, and will continue to be, a major source of debate among both policy makers and academics. Cigarette excise taxes have risen dramatically over the past decade relative to the recent past, yet taxes have only recently returned to their real level of the mid-1950s, and the share of taxes as a percentage of price remains well below its historical levels (Gruber, 2001). And in just the past five years we have seen a massive payment from the tobacco industry to the states, and an attempt at a comprehensive tobacco regulation bill at the federal level.

Economists should be important participants in the debate over government policy in this arena. The guidance provided by economists to date has been guided by the notion that smoking decisions are made in a rational, time consistent fashion along the lines of the Becker-Murphy model. But available evidence, albeit quite weak by empirical economics standards, does not support this formulation. The purpose of this paper was to write down a model which was more consistent with the available evidence. Our model deviates in only one way from the Becker Murphy formulation, by introducing specific time inconsistent preferences. We then assess the differential implications of this model for analyzing government policy towards smoking.

We find that this change in the model has radical implications for government policy, since government regulation provides a commitment device that is valued by time inconsistent consumers.

First, we estimate that the optimal tax on cigarettes, above and beyond externalities, is at least \$1 per pack, and quite likely much higher. We also find that taxes should be high up front, in order to break addiction, but that this can lead to government time inconsistency problems. With a constrained time pattern of taxes, taxes can be used as substitutes over time; likewise, they can be substitutes over space, suggesting a novel justification for clean air regulations that limit smoking in public places.

Second, we show that the traditional conclusion that cigarette taxes are regressive is reduced, and most likely reversed, once the self-control benefits of taxes are accounted for. We develop an adjustment to the standard measure of tax incidence which accounts for these benefits. We note that, while lowering incidence overall, this adjustment does not unambiguously reduce regressivity; that depends on the relative price sensitivities, values of life, degree of time inconsistency, and degree of impatience across income groups. But our calibrations show that, given the much higher price elasticities of lower income smokers, in almost all cases taxes on cigarettes are much less regressive and in some case are indeed progressive.

One concern about this model and this set of conclusions is that it presents a “slippery slope” towards justifying excessive regulation of a host of economic behaviors, ranging from smoking and drinking to driving and fast food consumption. But there are at least three reasons why smoking is a more appropriate platform for our model than other behaviors. First, there is significant casual evidence (and the one econometric study of Gruber and Mullainathan (2001)) to suggest that smoking decisions are taken in a time inconsistent fashion. Second, smoking is clearly harmful at all levels, and the harm rises monotonically with the amount consumed; drinking, for example, is sometimes argued to be beneficial at low levels of consumption, and only harmful at very high levels. Finally, the internal costs of smoking dwarf its external costs; the vast majority of harm done by a smoker is to himself or herself. At standard values of the value of a life/year, we estimate above that a pack of cigarettes costs \$35.64 in terms of lost life expectancy, roughly 100 times the level of externalities from smoking. This suggests that simply relying on externalities to determine optimal policy can lead to very large mistakes; in other words, this is a place where getting the model right matters a lot. We think that it is valuable to consider the implications of models such as this in other arenas; but the argument for doing so in the context of smoking is most strong.

Of course, the empirical evidence for time inconsistency in smoking remains weak, and much more work is needed here. At the same time, the fact that there is no empirical support, or even laboratory support, for exponential discounting in this or related contexts suggests that alternative

models of the type that we have derived be taken seriously. The important general point is that, when standard public finance analyses suggest that the tax on addictive bads is simply equal to their external costs, and that such taxes are highly regressive, those analyses are implicitly embracing a rational addiction model. Given the enormous magnitude of the internal costs to smoking, however, alternative models such as ours must be considered in designing regulatory policy towards addictive goods.

A Proofs

Lemma 1 1. λ_t is decreasing in β for each t .

2. λ_t is decreasing in δ for each t .

3. λ_t is increasing in t .

Proof. Since the utility function is quadratic, we have

$$S_{t+1} = (1-d)(S_t + a_t) = (1-d)(S_t + \lambda_t S_t + \mu_t) \quad (23)$$

$$a_{t+1} = \lambda_{t+1} S_{t+1} + \mu_{t+1} = \lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1} \quad (24)$$

Plugging this into the sophisticates' first-order condition, equation 5, and assuming $p_t = p$ in each period:

$$\begin{aligned} & \alpha_a + \alpha_{aa}(\lambda_t S_t + \mu_t) + \alpha_{as} S_t - p\alpha_c = \\ &= (1-d)\delta[(1 + (1-\beta)\lambda_{t+1})[\alpha_a + \alpha_{aa}(\lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1}) \\ &+ \alpha_{as}(1-d)(S_t + \lambda_t S_t + \mu_t) - p\alpha_c] \\ &- \beta(\alpha_s + \alpha_{as}(\lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1}) + \alpha_{ss}(1-d)(S_t + \lambda_t S_t + \mu_t))] \end{aligned} \quad (25)$$

The above has to be true for all S_t , so the coefficient of S_t in the expression has to be zero. After 'some' manipulation, this implies

$$\lambda_t = -1 + \frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \delta(1-d)^2[(1 + (1-\beta)\lambda_{t+1})(\alpha_{aa}\lambda_{t+1} + \alpha_{as}) - \beta\alpha_{as}\lambda_{t+1} - \beta\alpha_{ss}]} \quad (26)$$

Define the function $f_s(\lambda)$ according to equation 26. We will prove that

$$-f_s(\lambda_T) < \lambda_T = \frac{\alpha_{as}}{-\alpha_{aa}}.$$

$$-f(-1) > -1.$$

$-f_s$ is continuous and increasing on $(-1, \lambda_T)$.

$-f_s$ is decreasing in β and δ .

These are sufficient to establish all parts of the lemma.

First, notice that the second term of $f_s(\lambda)$ is the reciprocal of a quadratic with a negative coefficient on λ^2 . Then if this term is positive for two points on the real line, it is also positive in-between these two points. Moreover, it is easy to show that on the interval where this term is positive, f_s is strictly convex.²⁰ Therefore, it is sufficient to show that $f_s(-1) > -1$, $\lambda_T > f_s(\lambda_T) > -1$, and that $f'_s(-1) \geq 0$. The first two ensure that we are on the continuous and strictly convex section of f_s , and the last one (together with convexity) ensures that f_s is increasing on $(-1, \lambda_T)$.

The rest is just carrying out the above. We have

$$f_s(-1) = -1 + \frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \beta\delta(1-d)^2[-\alpha_{aa} + 2\alpha_{as} - \alpha_{ss}]} > -1 \quad (27)$$

as both the numerator and the denominator are positive in the second term. Proceeding,

$$\begin{aligned} f_s(\lambda_T) &= f_s\left(\frac{\alpha_{as}}{-\alpha_{aa}}\right) = -1 + \frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]} \\ &= \frac{\alpha_{as} - \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]}{-\alpha_{aa} + \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]} \end{aligned} \quad (28)$$

This being $< \lambda_T = \frac{\alpha_{as}}{-\alpha_{aa}}$ is equivalent to $-\alpha_{as}\lambda_T - \alpha_{ss} > 0$. But the latter can be rewritten as $\alpha_{as}^2 < \alpha_{ss}\alpha_{aa}$, and since owing to the concavity of $U(a_t, c_t, S_t)$ we have $\alpha_{as}^2 < \alpha_{ss}\alpha_{aa}$, this inequality holds. $-\alpha_{as}\lambda_T - \alpha_{ss} > 0$ also implies that the second term is positive, so that $f(\lambda_T) > -1$.

Moving on,

$$f'_s(\lambda) = -\frac{(\alpha_{as} - \alpha_{aa})(2(1-\beta)\lambda\alpha_{aa} + \alpha_{aa} + (1-2\beta)\alpha_{as})}{[-\alpha_{aa} + \delta(1-d)^2[(1+(1-\beta)\lambda)(\alpha_{aa}\lambda + \alpha_{as}) - \beta\alpha_{as}\lambda - \beta\alpha_{ss}]]^2}, \quad (29)$$

which gives

$$f'_s(-1) = -(1-2\beta) \left[\frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \delta(1-d)^2[(1+(1-\beta)\lambda)(\alpha_{aa}\lambda + \alpha_{as}) - \beta\alpha_{as}\lambda - \beta\alpha_{ss}]} \right]^2 \geq 0. \quad (30)$$

Finally, we prove that $f_s(\lambda_t)$ is decreasing in β and δ . For δ , this is quite simple: the term multiplying δ in the denominator is positive since $\lambda_t \leq \lambda_T$ and the utility function is strictly

²⁰The second derivative of the reciprocal of a quadratic q is $-\frac{q^2q'' - q(q')^2}{q^4}$, which is positive as long as q is positive and concave.

concave. It takes slightly more work to prove that $f_s(\lambda_t)$ is decreasing in β . The derivative of the denominator with respect to β is a positive constant times

$$-[\alpha_{aa}\lambda_{t+1}^2 + 2\alpha_{as}\lambda_{t+1} + \alpha_{ss}]. \quad (31)$$

Notice that this is equal to

$$-(\lambda_{t+1} \ 1) \begin{pmatrix} \alpha_{aa} & \alpha_{as} \\ \alpha_{as} & \alpha_{aa} \end{pmatrix} \begin{pmatrix} \lambda_{t+1} \\ 1 \end{pmatrix}, \quad (32)$$

which is greater than zero since the utility function's Hessian is negative definite.

This completes the proof. \square

Theorem 2 $\lambda^{*n} \geq \lambda^{*s} \geq \lambda^{*TC}$.

Proof. To prove $\lambda^{*s} \geq \lambda^{*TC}$, define f_{TC} as f_s except with $\beta = 1$. The difference between the denominators of the second terms of f_s and f_{TC} , ignoring the positive multiplicative constant $\delta(1-d)^2$, is

$$(1-\beta)[\lambda^2(\alpha_{aa} + p^2\alpha_{cc}) + 2\alpha_{as}\lambda + \alpha_{ss}] = (1-\beta)(\lambda \ 1 \ p\lambda) \begin{pmatrix} \alpha_{aa} & \alpha_{as} & 0 \\ \alpha_{as} & \alpha_{ss} & 0 \\ 0 & 0 & \alpha_{cc} \end{pmatrix} \begin{pmatrix} \lambda \\ 1 \\ p\lambda \end{pmatrix} \leq 0 \quad (33)$$

by the concavity of $U(a_t, c_t, S_t)$. This implies that for any $\lambda \in (-1, \lambda_T)$, $f_s(\lambda) > f_{TC}(\lambda)$, and thus $\lambda^{*s} \geq \lambda^{*TC}$.

To prove the inequality $\lambda^{*n} \geq \lambda^{*s}$, let us introduce relevant value functions for sophisticates and time-consistent agents:

$$\begin{aligned} V^s(S_{t+1}) &= \sum_{i=1}^{T-t} \delta^i U(a_{t+i}^s, c_{t+i}^s, S_{t+i}^s) \\ V^{TC}(S_{t+1}) &= \sum_{i=1}^{T-t} \delta^i U(a_{t+i}^{TC}, c_{t+i}^{TC}, S_{t+i}^{TC}), \end{aligned} \quad (34)$$

where the superscripts refer to the two different agents. Since the utility function is quadratic and strategies are linear, both of these are quadratic in S_{t+1} . Furthermore, it is clear that $V^{TC}(S_{t+1}) \geq V^s(S_{t+1})$ for all S_{t+1} , so $V^{TC}(S_{t+1})$ has at least as large a prime coefficient as $V^s(S_{t+1})$. In addition, these coefficients are negative ²¹.

Agent i ($i = n, s$) solves

²¹The easiest way to see this is to check that $V_S^{TC}(S_{t+1}) = \delta v_s(a_{t+1}^{TC}, S_{t+1}^{TC})$ is negative for a sufficiently high S_{t+1} .

$$\max_{a_t} v(a_t, S_t) + u(I_t - pa_t) + \beta V^i((1-d)(S_t + a_t)), \quad (35)$$

leading to the first-order condition

$$v_a(a_t, S_t) - pu'(c_t) + \beta(1-d)V_S^i(S_{t+1}) = 0. \quad (36)$$

Differentiating this totally with respect to S_t gives

$$\begin{aligned} \frac{\partial a_t}{\partial S_t} &= \frac{v_{as}(a_t, S_t) + \beta\delta(1-d)^2 V_{SS}^i(S_{t+1})}{-v_{aa}(a_t, S_t) - p^2 u''(c_t) - \beta\delta(1-d)^2 V_{SS}^i(S_{t+1})} \\ &= -1 + \frac{v_{as}(a_t, S_t) - v_{aa}(a_t, S_t) - p^2 u''(c_t)}{-v_{aa}(a_t, S_t) - p^2 u''(c_t) - \beta\delta(1-d)^2 V_{SS}^i(S_{t+1})}. \end{aligned} \quad (37)$$

Since v_{as} , $-u''$, $-v_{aa}$, and $-V_{SS}^i$ are positive constants, and $-V_{SS}^s \geq -V_{SS}^{TC}$, the above implies $\lambda^{*n} \geq \lambda^{*s}$. \square

B Incidence Calibration for Naifs

When considering the effect of a price increase on the utility of naive hyperbolic discounters, even with the discount structure given, we can use at least two measures for their utility: when discounted utility is evaluated according to their true instantaneous utilities and when it is evaluated according to naifs' perceived instantaneous utilities. We believe that the appropriate measure for welfare analysis is agents' true utility, so we calibrate the incidence of addictive goods taxation on naifs using this measure. For a quadratic utility function, and ignoring length-of-horizon effects (taking the horizon to be infinite), the derivative of a naif's exponentially discounted utility function with respect to price p is then

$$-\alpha_c \sum_{t=1}^{\infty} \delta^t a_t + \sum_{t=1}^{\infty} \delta^t \frac{\partial a_t}{\partial p} \left[(v_a(a_t, S_t) - p\alpha_c) + \sum_{k=1}^{\infty} \delta^k (1-d)^k (1 + \lambda^{*n})^{k-1} [(v_a(a_{t+k}, S_{t+k}) - p\alpha_c) \lambda^{*n} + v_s(a_{t+k}, S_{t+k})] \right] \quad (38)$$

This expression is derived by considering the total utility effect of changing the agent's consumption in each period (through the change in price), including the utility effect coming through changing later selves' consumption levels. As before, we assume that the disutility from stock $v_S(a_t, S_t)$ is constant. Let this constant be v_S . Combining time-consistent agents' and naifs' first-order conditions, we then have

$$v_a(a_t, S_t) - p\alpha_c = -\beta \frac{\delta(1-d)}{1 - \delta(1-d)} v_S. \quad (39)$$

Combining this with the rest of the above expression yields

$$-\alpha_c \sum_{t=1}^{\infty} \delta^t a_t + \sum_{t=1}^{\infty} \delta^t \frac{\partial a_t}{\partial p} \left[\frac{\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})} v_S - \beta \frac{\delta(1-d)}{1-\delta(1-d)} v_S - \beta \frac{\delta(1-d)}{1-\delta(1-d)} \frac{\delta(1-d)\lambda^{*n}}{1-\delta(1-d)(1+\lambda^{*n})} v_S \right]. \quad (40)$$

Rearranging gives

$$-\alpha_c \sum_{t=1}^{\infty} \delta^t a_t + \sum_{t=1}^{\infty} \delta^t \frac{\partial a_t}{\partial p} (1-\beta) \frac{\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})} v_S. \quad (41)$$

If the agent started off in a steady state with $a_t = a$ for all t , the multiplicative adjustment to standard incidence is

$$1 - (1-\beta) \frac{-\partial a}{\partial p} \frac{p}{a} \cdot \frac{-h_S}{p} \cdot \frac{\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})}. \quad (42)$$

The total discounted future utility cost of an extra cigarette, expressed in monetary terms, is $H_S = \frac{\delta(1-d)}{1-\delta(1-d)} h_S$. Substituting this in the above yields our final expression

$$1 - (1-\beta) \frac{-\partial a}{\partial p} \frac{p}{a} \cdot \frac{-H_S}{p} \cdot \frac{1-\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})}. \quad (43)$$

This expression is almost identical to the one for sophisticated agents, with two crucial differences. Naturally, the naifs' self-control adjustment features the naifs' responsiveness to stock, λ^{*n} . Perhaps more surprisingly, the consumption response to an increase in stock is multiplied by $1-\beta$ for sophisticates, and not for naifs. We prove in the appendix that $\lambda^{*n} \geq \lambda^{*s}$ (theorem 2), so both of these tend to decrease the multiplicative adjustment for naive quasi-hyperbolic discounters relative to their sophisticated counterparts. Therefore, the incidence of a tax tends to be smaller on naifs than on sophisticates.

References

- AINSLIE, G. (1992): *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person*. Cambridge, UK: Cambridge University Press.
- AINSLIE, G., AND N. HASLAM (1992): "Hyperbolic Discounting," in *Choice Over Time*, ed. by G. Loewenstein, and J. Elster, chap. 3, pp. 57–92. New York, NY: Russell Sage Foundation.
- ANGELETOS, G.-M., D. I. LAIBSON, A. REPETTO, J. TOBACMAN, AND S. WEINBERG (2001): "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," Mimeo, Harvard University.

- BARRO, R. (1999): “Ramsey Meets Laibson in the Neoclassical Growth Model,” *Quarterly Journal of Economics*, 114, 1125–1152.
- BECKER, G., M. GROSSMAN, AND K. M. MURPHY (1994): “An Empirical Analysis of Cigarette Addiction,” *American Economic Review*, 84(3), 396–418.
- BECKER, G. S., AND K. M. MURPHY (1988): “A Theory of Rational Addiction,” *Journal of Political Economy*, 96(4), 675–700.
- BERNHEIM, B. D., AND A. RANGEL (2001): “Addiction, Conditioning, and the Visceral Brain,” Mimeo, Stanford University.
- BERNSTEIN, D. A. (1970): “The Modification of Smoking Behavior: An Evaluative Review,” in *Learning Mechanisms in Smoking*, ed. by W. A. Hunt, pp. 3–41. Aldine Publishing Company, Chicago.
- BURNS, J. (1992): “Looking to the Future,” in *Special Report: Business and Health*, ed. by J. Burns, pp. 21–22. Medical Economics Publishing.
- CDC (various years): “You Can Quit Smoking,” Tobacco Information and Prevention Source, Webpage: www.cdc.gov/tobacco.
- CHALOUPKA, F. J. (1991): “Rational Addictive Behavior and Cigarette Smoking,” *Journal of Political Economy*, 99(4), 722–742.
- CHALOUPKA, F. J., AND K. E. WARNER (1998): “The Economics of Smoking,” in *The Handbook of Health Economics*, ed. by J. Newhouse, and A. Culyer. forthcoming.
- DELLA VIGNA, S., AND U. MALMENDIER (2001): “Self-control in the Market: Evidence from the Health Club Industry,” Mimeo, Harvard University.
- DIAMOND, P., AND B. KŐSZEGI (1998): “Hyperbolic Discounting and Retirement,” Revise and resubmit, *Journal of Public Economics*.
- EVANS, W., J. RINGEL, AND D. STECH (1999): “Tobacco Taxes and Public Policy to Discourage Smoking,” in *Tax Policy and the Economy*, ed. by J. Poterba, pp. 1–56. Cambridge, MA: MIT Press.

- GRABOWSKI, J., AND S. M. HALL (1985): "Tobacco Use, Treatment Strategies, and Pharmacological Adjuncts: An Overview," in *Pharmacological Adjuncts in Smoking Cessation*, ed. by J. Grabowski, and S. M. Hall, pp. 1–13. National Institute on Drug Abuse Monograph 53.
- GRUBER, J. (2001): "Tobacco at the Crossroads: The Past and Future of Smoking Regulation in the U.S.," *Journal of Economic Perspectives*, 15(2), 193–212.
- GRUBER, J., AND B. KŐSZEGI (2000): "Is Addiction 'Rational?' Theory and Evidence," *NBER Working Paper 7507*.
- (2001): "Is Addiction 'Rational?' Theory and Evidence," *Quarterly Journal of Economics*, 116(4), 1261–1305.
- GRUBER, J., AND S. MULLAINATHAN (2001): "Do Cigarette Taxes Make Smokers Happier?," Mimeo, Massachusetts Institute of Technology.
- GUL, F., AND W. PESENDORFER (forthcoming): "Temptation and Self-Control," *Econometrica*.
- KIRBY, K. N., AND R. J. HERRNSTEIN (1995): "Preference Reversals Due to Myopic Discounting of Delayed Reward," *Psychological Science*, 6(2), 83–89.
- LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–477.
- LAIBSON, D. I., A. REPETTO, AND J. TOBACMAN (1998): "Self-Control and Saving for Retirement," in *Brookings Papers on Economic Activity*, ed. by W. C. Brainard, and G. L. Perry, vol. 1, pp. 91–196. The Brookings Institution.
- MILLER, J. L. (1978): "Self-Control in the Elimination of Cigarette Smoking: Case Histories Using a Changing Criterion Design," Master's thesis, Western Michigan University.
- MURRAY, R. G., AND S. A. HOBBS (1981): "Effects of Self-Reinforcement and Self-Punishment in Smoking Reduction: Implications for Broad-Spectrum Behavioral Approaches," *Addictive Behaviors*, 6(1), 63–67.
- O'DONOGHUE, T., AND M. RABIN (1999a): "Doing It Now Or Later," *American Economic Review*, 89(1), 103–124.

- (1999b): “Procrastination in Preparing for Retirement,” in *Behavioral Dimensions of Retirement Economics*, ed. by H. J. Aaron, pp. 125–156. Brookings Institution Press.
- (2000): “Addiction and Present-Biased Preferences,” Mimeo, Cornell University.
- PESENDORFER, W., AND F. GUL (2000): “A Theory of Addiction,” Mimeo.
- POTERBA, J. M. (1989): “Lifetime Incidence and the Distributional Burden of Excise Taxes,” *American Economic Review*, 79(2), 325–330.
- PROCHASKA, J. O., P. CRIMI, D. LAPSANSKI, L. MARTEL, AND P. REID (1982): “Self-Change Processes, Self-Efficacy and Self-Concept in Relapse and Maintenance of Cessation and Smoking,” *Psychological Reports*, 51, 983–990.
- THALER, R. (1981): “Some Empirical Evidence on Dynamic Inconsistency,” *Economics Letters*, 8, 201–207.
- U.S. DEPARTMENT OF HEALTH AND HUMAN SERVICES (1994): *Preventing Tobacco Use Among Young People: A Report of the Surgeon General*. National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health.
- VISCUSI, W. K. (1993): “The Value of Risks to Life and Health,” *Journal of Economic Literature*, 31(4), 1912–1946.
- VISCUSI, W. K., AND W. N. EVANS (1990): “Utility Functions That Depend on Health Status: Estimates and Economic Implications,” *American Economic Review*, 80(3), 353–374.
- VISCUSI, W. K., AND J. HERSCH (2001): “Cigarette Smokers as Job Risk Takers,” *Review of Economics and Statistics*, 83, 269–280.

Table 1: Distribution of Smoking Expenditures

Quartile or Group	Income Quartiles	Consumption Quartiles	Education Groups
Bottom Quartile / HS Dropouts	0.032	0.015	0.014
Second Quartile / HS Grads	0.014	0.013	0.013
Third Quartile / Some College	0.009	0.010	0.008
Top Quartile / College Grads	0.004	0.004	0.003

Note: Table shows the share of income (in columns 1 & 3) or expenditure (in column 2) on cigarettes. Data on consumption from CEX; data on income from CPS.

Table 2: Optimal Tax Calibrations for Sophisticates

$\beta = 0.9, \delta = 0.97$					
	$d = 0.5$	$d = 0.6$	$d = 0.7$	$d = 0.8$	$d = 0.9$
$\alpha^s = 0.9$.41	1.53	2.28	2.82	3.24
$\alpha^s = 0.7$	1.16	2.00	2.57	2.99	3.31
$\alpha^s = 0.5$	1.88	2.46	2.86	3.16	3.39
$\beta = 0.6, \delta = 0.97$					
	$d = 0.5$	$d = 0.6$	$d = 0.7$	$d = 0.8$	$d = 0.9$
$\alpha^s = 0.9$	2.32	7.54	10.37	12.15	13.37
$\alpha^s = 0.7$	6.00	9.37	11.35	12.65	13.57
$\alpha^s = 0.5$	8.94	10.98	12.26	13.13	13.77
$\beta = 0.9, \delta = 0.9$					
	$d = 0.5$	$d = 0.6$	$d = 0.7$	$d = 0.8$	$d = 0.9$
$\alpha^s = 0.9$	0.23	0.78	1.16	1.43	1.63
$\alpha^s = 0.7$	0.60	1.02	1.30	1.51	1.67
$\alpha^s = 0.5$	0.96	1.25	1.45	1.59	1.71
$\beta = 0.6, \delta = 0.9$					
	$d = 0.5$	$d = 0.6$	$d = 0.7$	$d = 0.8$	$d = 0.9$
$\alpha^s = 0.9$	1.29	3.85	5.25	6.13	6.74
$\alpha^s = 0.7$	3.09	4.76	5.74	6.39	6.85
$\alpha^s = 0.5$	4.55	5.56	6.19	6.63	6.95

Note: Table presents optimal internalty taxes (in dollars) for various combinations of the rate of depreciation (d), the impact of the past stock on current smoking for sophisticated hyperbolic agents (α^s), the long run discount factor (β), and the short run discount factor (δ), based on authors calculations from equation (17) in the text.

Table 3: Price Sensitivity of Cigarette Expenditure

Quartile or Group	Income Quartiles	Consumption Quartiles	Education Groups
Overall		0.168 (0.049) [-0.661]	
Bottom Quartile / HS Dropouts	-0.032 (0.078) [-1.086]	-0.016 (0.059) [-1.050]	-0.044 (0.115) [-1.080]
Second Quartile / HS Grads	0.144 (0.087) [-0.699]	0.113 (0.083) [-0.770]	0.045 (0.096) [-0.927]
Third Quartile / Some College	0.273 (0.113) [-0.534]	0.414 (0.104) [-0.311]	0.417 (0.100) [-0.107]
Top Quartile / College Grads	0.341 (0.116) [-0.387]	0.215 (0.137) [-0.642]	0.175 (0.082) [-0.400]

Note: Table shows estimates of impact of cigarette prices on cigarette consumption levels. Estimated impact on consumption is first number in each cell; standard error of that estimate in parentheses below. Figure in square brackets is implied elasticity of quantity of cigarettes consumed. First row shows full sample estimates; remaining rows show estimates by income/expenditure/education categories. Estimates from CEX data over 1980-1998 period.

Table 4: Incidence Calibrations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Group	Ex-Ante		$= 0.9, \quad = 0.97$		$= 0.6, \quad = 0.97$			$= 0.9, \quad = 0.9$			$= 0.6, \quad = 0.9$		
		Low Value	Mid Value	Equal Value	Low Value	Mid Value	Equal Value	Low Value	Mid Value	Equal Value	Low Value	Mid Value	Equal Value
Current Income Distribution													
Low	1.69	1.21	0.20	-0.58	-0.56	-5.30	-8.96	1.45	0.94	0.54	0.56	-1.82	-3.67
2nd	0.71	0.37	0.17	0.09	-0.90	-1.82	-2.19	0.54	0.44	0.40	-0.10	-0.56	-0.75
3rd	0.47	0.16	0.11	0.16	-0.98	-1.23	-1.00	0.31	0.29	.31	-0.26	-0.38	-0.27
Top	0.18	-0.01	0.06	0.09	-0.73	-0.35	-0.23	0.08	0.12	0.14	-0.28	-0.09	-0.02
Expenditure Distribution													
Low	0.83	0.48	0.12	-0.26	-0.80	-2.51	-4.26	0.65	0.47	0.28	0.01	-0.85	-1.73
2nd	0.66	0.27	0.11	0.03	-1.16	-1.93	-2.31	0.46	0.38	0.34	-0.25	-0.65	-0.83
3rd	0.49	0.31	0.27	0.30	-0.36	-0.54	-0.40	0.40	0.38	0.39	0.06	-0.03	0.04
Top	0.22	-0.15	-0.01	0.04	-1.52	-0.86	-0.60	0.03	0.10	0.13	-0.65	-0.32	-0.19
Education Categories													
HSD	0.75	0.11	0.09	-0.26	-2.24	-2.35	-3.98	0.43	0.42	0.24	-0.75	-0.81	-1.63
HSG	0.66	0.03	-0.01	-0.10	-2.27	-2.46	-2.91	0.34	0.32	0.28	-0.82	-0.91	-1.14
SCL	0.4	0.35	0.34	0.35	0.15	0.11	0.15	0.37	0.37	0.37	0.27	0.25	0.27
CLG	0.13	0.03	0.05	0.07	-0.33	-0.27	-0.17	0.08	0.09	0.10	-0.10	-0.07	-0.02

Note: Table shows ex ante (column 1) and adjusted (columns 2-13) incidence of a \$1/pack rise in the tax on cigarettes. All estimates are for $d = 0.6$ and $s = 0.7$; corresponding and for each estimate is shown in top row. Within each set of columns, first column shows incidence with values of life dictated by elasticity of life value = 1; second column shows case where value of life of lowest income group is half that of highest income group; and final column shows case of equal value of life. First panel shows incidence calculations according to annual income quartiles; second panel shows calculations according to consumption quartiles; final panel shows incidence calculations according to education groups. Table results from evaluation of equation (22) in text.