# Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes

Jonathan Gruber[a,*], Botond Kőszegi[b]

[a] *Department of Economics, MIT and NBER, E52-355, 50 Memorial Drive, Cambridge, MA 02142-1347, USA*
[b] *UC Berkeley, USA*

## Abstract

One of the most cogent criticisms of excise taxes is their regressivity, with lower income groups spending a much larger share of their income on goods such as cigarettes than do higher income groups. We argue that traditional quantity-based measures of incidence are only appropriate under a very restrictive "time-consistent" model of consumption of sin goods. A model that is much more consistent with existing evidence on smoking decisions is a time-inconsistent formulation where excise taxes on cigarettes serve a self-control function that is valued by smokers who would like to quit but cannot. This self-control function benefits lower income groups more, since they have a significantly higher price sensitivity of smoking. Calibrations show that, as a result, cigarette taxes are much less regressive than previously assumed, and are even progressive for a wide variety of parameter values.
© 2003 Elsevier B.V. All rights reserved.

*Keywords:* Cigarettes; Excise tax; Self-control function

## 1. Introduction

One of the most cogent criticisms of excise taxes is their regressivity. Lower income groups consume the types of "sin" goods to which excise taxes apply at an equal or greater rate than do higher income groups. As a result, there is a strong inverse relationship between income and the income shares devoted to consumption of goods such as cigarettes, alcohol, and gasoline. This relationship is weakened, but not reversed, when more permanent

---

\* Corresponding author. Tel.: +1-617-253-8892; fax: +1-617-253-1330.
*E-mail address:* gruberj@mit.edu (J. Gruber).

measures of resources such as consumption are used to determine incidence (Poterba, 1989).

The presumption behind such regressivity arguments is that they reflect the degree to which each income group is "hurt" by different tax policies. For an economist, the appropriate measure for this analysis is utility. For example, when we say that an increase in the tax on gasoline would be borne most heavily by low-income consumers, we (should) mean that if the tax was instituted, their utility would be most detrimentally affected.

Incidence is traditionally computed in terms of quantities consumed, however, because in neoclassical economics, this is an accurate measure for the utility effect of taxation. 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, we can measure the "harm" of a US$1 tax in direct proportion to the amount of that good consumed by a given income group.

This same analysis can be applied to addictive goods if consumers are "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 decision-making process as is consumption of all other goods. Consumers trade off the utility gains from consuming the good against the costs of doing so. 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 such a model, incidence analysis is effectively the same for addictive as for non-addictive goods.

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 trade-offs, relative to the discount rate that they use in evaluating similar trade-offs between today and the future. 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 et al., 1998; O'Donoghue and Rabin, 1999b), retirement decisions (Diamond and Kőszegi, 1998), and growth (Barro, 1999).

The goal of this paper is to explore the implications of applying quasi-hyperbolic discounting to incidence analysis. 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 Section 1, with some background on smoking, addiction modeling, and time inconsistency, and Section 3 introduces the model. We then turn to the focus of this paper, incidence analysis. Standard incidence analysis, which applies to the Becker–Murphy model, 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.

Section 4 then presents a calibration exercise that shows the quantitative importance of this point. A key parameter that determines incidence in our model is the price sensitivity of smoking; a higher price sensitivity raises the value of the self-control adjustment to incidence. As we document empirically using data on cigarette expenditures from the Consumer Expenditure Survey (CEX), the price sensitivity of smoking is much larger for the lowest income groups than for their higher income counterparts, with an elasticity of − 1 for the bottom income quartile. As a result, for many parameter values, taxes are much less regressive or even progressive when the benefits of self-control are included. Section 6 concludes with the policy implications of our findings.

## 2. Background

### 2.1. 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.

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 et al. (1999a,b). Estimates vary, but cluster around 40¢ per pack, which is well below the average level of state and federal excise taxation (76¢ per pack, according to Gruber, 2001).

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, as discussed in Evans et al. (1999a,b) and updated in Gruber and Kőszegi (2002). Cigarette expenditures as a share of income are 3.2% in the bottom quartile of the income distribution, but are only 0.4% of income in the top quartile. This regressivity persists when incidence is computed on a lifetime basis, as shown in Poterba (1989); regressivity is lessened, but expenditures in the bottom quartile of the consumption distribution are still much higher as a percentage of consumption than in the top consumption quartile.

## 2.2. The case for time inconsistency in smoking

There are four types of evidence for time inconsistency, two of which apply to behavior in general, and two of which apply to smoking in particular. The first is laboratory experiments, which 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.

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 high illiquid wealth holdings (e.g. housing) combined with high levels of borrowing at high interest rates (on credit cards) is consistent with the self-control problems inherent in time-inconsistent models. And Della Vigna and Malmendier (2001) argue 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.

Fourth, there is evidence on two key features that distinguish time-consistent agents from time-inconsistent ones. 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 decision-makers 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 decision-making. 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 et al., 1982). Various punishment and self-control strategies for quitting are also 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% say that they will not be smoking 5 years later, but only 31% of them have in fact quit 5 years hence. Moreover, among those who smoke more than one pack/day, the smoking rate 5 years later among those who stated that they would be smoking (72%) is actually *lower* than the smoking rate among those who stated that they would not be smoking (74%) (U.S. Department of Health and Human Services, 1994).

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 *in any domain*. 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.[1]

---

[1] There is a large empirical literature which has tested and generally supported a 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. But, as Gruber and Kőszegi (2001) 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 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 a second key premise, time consistency.

Much of this evidence is consistent with at least two other recent models of individual decision-making as well (Bernheim and Rangel, 2001; Gul and Pesendorfer, 2001; Pesendorfer and Gul, 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.

## 3. 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 Kőszegi, 2001). This model marries what are in our opinion the two most important aspects of the consumption of addictive goods. First, the instantaneous utility from consuming a good such as cigarettes depends in specific ways on past consumption of the same good. We will be focusing on two kinds of intertemporal linkages: reinforcement—the tendency of past consumption of cigarettes to increase the "craving" for (formally, the marginal utility of) a cigarette today—and health costs. Both of these properties of addictive substances can be captured in an instantaneous utility function of the form

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

$a_t$ and $c_t$ are the levels of consumption of the addictive and ordinary goods, respectively. Reinforcement and health costs 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

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

where $0 < d < 1$ is the depreciation rate of the stock. Reinforcement means simply that $v_{aS} (a_t, S_t) > 0$, while adverse health effects are formalized by assuming $v_S (a_t, S_t) < 0$.

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

---

[2] Bernheim and Rangel (2001) assume that in the "visceral" state when the agent overconsumes the drug, she is not price-sensitive. Gul and Pesendorfer (2001) and Pesendorfer and Gul (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.

[3] For 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.

a *T*-period model. A *time-consistent*, exponential discounter agent makes decisions at time *t* according to the discounted utility function

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

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

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

$\beta$ and $\delta$ are usually assumed to be between zero and one. This formulation is intended to capture the idea (for which we argued above) that decision-makers 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 ($\delta$) 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 self-control problem arises from the fact that future selves will again prefer to be impatient in the short run, creating a conflict between the current self and future ones. There are two extreme assumptions one can make about how this conflict plays out. Under naiveté, the decision-maker 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 as she chooses to be more impatient than she expected. 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 *decrease* the incidence of taxes.[4] But we believe borrowing to finance consumption is not an important factor in cigarette smoking. In addition, the agent's use of savings and addictive goods consumption as a self-control device considerably complicates the analysis,[5] and we have

---

[4]  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, and this extra effect would be beneficial for the agent.

[5]  In 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.

not solved for the dynamics of our model with savings. For some very expensive addictive goods, however, it may be important to extend our model in this way.[6]

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

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

As in our earlier paper (Gruber and Kőszegi, 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

$$
v(a_t, S_t) = \alpha_a a_t + \alpha_s S_t + \frac{\alpha_{aa}}{2} a_t^2 + \alpha_{as} a_t S_t + \frac{\alpha_{ss}}{2} S_t^2
$$

$$
u(c_t) = \alpha_c c_t \tag{6}
$$

where $\alpha_a$, $\alpha_{as}$, and $\alpha_c$ are positive and $\alpha_s$, $\alpha_{aa}$, and $\alpha_{ss}$ 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 (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 $\lambda_t$ and $\mu_t$ are constants. Gruber and Kőszegi (2001) prove that for $\beta \geq (1/2)$, $\lambda_{T-j}$ converges to a constant $\lambda^{*s}$ as $j \to \infty$, so that marginal responsiveness to past consumption is approximately stationary far from the end of the horizon. Although this stationarity is not crucial for any of our qualitative points, we will take advantage of it in our calibrations.

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 $\lambda^{*s}$ (Gruber and Kőszegi, 2001).

As in any model where different socially relevant actors have different tastes, a discussion of 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. We follow the literature and take the agent's long-run preferences as those relevant for social welfare maximization.

---

[6] For a more thorough introduction of this model, see Gruber and Kőszegi (2001); for a contrast of sophistication and naiveté, see O'Donoghue and Rabin (1999a).

Since the discount factor $\beta$ applies to all future periods, this is also the appropriate welfare measure if a sophisticated representative agent were to vote in a tax change today that is instituted starting in the next period.[7,8] More generally, as emphasized by O'Donoghue and Rabin (2003), the extra discount factor $\beta$ that each self applies to the future is only relevant for that single self, so ignoring it when evaluating the agent's overall welfare may be justified. Even more importantly, since we will be focusing on a period length of 1 month, any reasonable $\delta$ in our model is very close to 1. Thus, since the agent cares a lot about future discounted utility, our results would be almost identical if we considered her current preferences at the time of the tax change as the appropriate welfare measure.[9]

## 4. 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. 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

$$\max q_1, \ldots, q_n \quad U(q_1, \ldots, q_n)$$
$$\text{s.t.} \quad p_1 q_1 + \ldots + p_n q_n \leq W \tag{7}$$

We know that the derivative of the maximum in this problem with respect to $p_1$ is $-\alpha_c q_1$, where $\alpha_c$ is the Lagrange multiplier on the constraint. That is, for a maximizing consumer,

---

[7] To see this, note simply that a price change that starts tomorrow has no negative income effect today, and the response of today's consumption to future taxes has no effect on discounted utility due to the envelope theorem. Thus, the effect on discounted utility comes exclusively from future periods, which are discounted exponentially.

[8] Interestingly, 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 in Gruber and Kőszegi (2000).

[9] In Gruber and Kőszegi (2002), we discuss the implications of this model for the optimal level of taxes on cigarettes. We prove that, so long as $\beta < 1$, it will be optimal to tax cigarettes, even in the absence of externalities. And we show through calibrations similar to those below that the optimal tax on cigarettes is US$1−2 if consumers are only slightly time-inconsistent, or US$5−10 if consumers are considerably time-inconsistent. Those results, as well as the incidence results in this paper, presumes that there are not fully effective private self-control devices available, but this seems a reasonable assumption. 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 important enforcement problems.

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.

The same is 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

$$-\alpha_c a_t - \alpha_c \delta a_{t+1} - \alpha_c \delta^2 a_{t+2} - \ldots - \alpha_c \delta^{T-t} a_T. \tag{8}$$

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.

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. We focus here on the more complicated case of sophisticates, and mention only briefly how our results differ for naifs.

To arrive at the correct measure of incidence, we take the derivative of the agent's long-run utility (the welfare measure for which we have argued above) with respect to $p$:

$$\frac{d}{dp}\left(\sum_{j=1}^{T}\delta^{j-1}(v(a_j, S_j) + \alpha_c(I_j - pa_j))\right). \tag{9}$$

Here, each $a_j$ is a function of $S_j$ and $p$, which we will have to take into account in evaluating the total derivative with respect to $p$. Let

$$U_j(S_j, p) = \sum_{i=j}^{T}\delta^{i-j}(v(a_i, S_i) + \alpha_c(I_i - pa_i)).$$

Thus we are interested in

$$\begin{aligned}
\frac{d}{dp}U_1(S_1, p) &= \frac{d}{dp}\left(v(a_1, S_1) + \alpha_c(I_1 - pa_1) + \delta U_2(S_2, p)\right)\\
&= -\alpha_c a_1 + \frac{\partial a_1}{\partial p}\left(v_a(a_1, S_1) - p\alpha_c + \delta(1-d)\frac{\partial}{\partial S_2}U_2(S_2, p)\right)\\
&\quad + \delta\frac{\partial}{\partial p}U_2(S_2, p).
\end{aligned}$$

Using that self 1 solves $v_a(a_1, S_1) - p\alpha_c + \beta\delta(1-d)\,(\partial/\partial S_2)\,U_2(S_2, p) = 0$, this can be rewritten as

$$-\alpha_c a_1 - \frac{1-\beta}{\beta}\frac{\partial a_1}{\partial p}(v_a(a_1, S_1) - p\alpha_c) + \delta\frac{\partial}{\partial p}U_2(S_2, p).$$

Iterating this procedure, we get that the derivative of the agent's discounted utility with respect to $p$ is

$$-\alpha_c \underbrace{\left( \sum_{j=1}^{T} \delta^{j-1} a_j \right)}_{\text{standard term}} - \underbrace{(1-\beta) \sum_{j=1}^{T} \delta^{j-1} \frac{\partial a_j}{\partial p} \frac{v_a(a_j, S_j) - p\alpha_c}{\beta}}_{\text{self-control adjustment}}. \tag{10}$$

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 attached 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:

$$\begin{aligned}
\frac{v_a(a_j, S_j) - p\alpha_c}{\beta} &= -(1-d)\delta v_S(a_{j+1}, S_{j+1}) - (1-d)^2 \delta^2 (1 + (1-\beta)\lambda_{j+1}) v_S \\
&\quad \times (a_{j+2}, S_{j+2}) - (1-d)^3 \delta^3 (1 + (1-\beta)\lambda_{j+1})(1 + (1-\beta)\lambda_{j+2}) v_S \\
&\quad \times (a_{j+3}, S_{j+3}) - \ldots
\end{aligned} \tag{11}$$

We will use expressions (10) and (11) to study the self-control adjusted incidence of cigarette taxation. But even though our analysis is motivated by cigarette consumption, the methods are applicable to other goods as well.

We will need the following preliminary lemma for our analysis:

**Lemma 1.**

1. $\lambda_t$ *is decreasing in $\beta$ for each t.*
2. $\lambda_t$ *is decreasing in $\delta$ for each t.*
3. $\lambda_t$ *is increasing in t.*

**Proof.** Appendix A. □

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.

Next, notice that the right-hand side of Eq. (11) depends on the constants $\lambda_i^s$ 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 exacerbates the future self-control problem as well by increasing the short-run desire to consume. 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.[10] The same logic holds for naive consumers, so the conclusion that the true burden of taxation is lower for addictive harmful goods is true for them as well.

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 of the 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 (10) 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 tool 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 $\beta$. In this case, the 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 $\beta$ is an increase in the multiplier $1 - \beta$ in the second term in expression (10). The intuition is obvious: if $\beta$ 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, Eq. (11) 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 $\beta$ is lower, since in that case the future selves are making a worse decision. Exacerbating the problem is that $\lambda_i^s$ is greater for agents with a lower $\beta$; 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 $\beta$), but a lower long-run discount factor $\delta$. In fact, if the poor have a lower $\delta$, then, all else equal, the burden of taxes falls *more heavily* on them. Lowering $\delta$ has the direct effect of increasing discounting in

---

[10] This feature is all the more striking given the result in our earlier paper that addictiveness decreases the optimal tax, since future selves exert more self-control to reduce an addictive behavior. 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.

Eq. (10), and decreasing the right-hand side of Eq. (11), 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.[11] Though the standard incidence measure is also lower when $\delta$ is lower, the former effect is proportionally larger.

Lowering $\delta$ 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. The conclusion might be surprising: if the poor's higher smoking rate comes from a smaller *short-term* discount factor ($\beta$), then a cigarette tax is not as regressive as currently believed, while if it comes from a lower *long-term* discount factor ($\delta$), 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 $\alpha_s$, which reduces the marginal harm from smoking. A change in $\alpha_s$ does not affect $\lambda_t^s$ for any $t$, so from Eq. (11), 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.[12]

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 $\delta$ 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.[13]

## 5. What difference does it make?

### 5.1. Calibrating tax incidence—theory

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 decision-makers start off from a steady-state consumption level: $a_t$ does not depend on $t$. This is not too restrictive, since smokers tend to reach their steady-state consumption level by their early twenties. Second, we ignore end-of-life effects operating through $\lambda_t$ and the price elasticity of

---

[11] An effect acting in the opposite direction arises from the addictiveness of the good: lowering $\delta$ increases $\lambda_j^s$. 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 $\delta$ always outweighs the latter (indirect) effect. The proof is omitted, but is available upon request.

[12] As 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.

[13] Note that this formulation differs from the standard interpersonal externality setup, from which it may seem that the corrective effect of a tax should depend only on the price responsiveness and $1 - \beta$ times the health cost of smoking. The reason is that there are intertemporal linkages in our framework, so changing the behavior of one self also changes the behavior of other selves.

consumption. Third, we assume that the disutility associated with smoking, $v_S(a_t, S_t)$, is a constant $v_S$.[14]

Setting $\lambda_t = \lambda^{*s}$ and $v_S(a_t, s_t) = v_S$, expression (11) becomes

$$\frac{v_a(a_j, S_j) - p\alpha_c}{\beta} = -(1-d)\delta v_S \frac{1 - [(1-d)\delta(1 + (1-\beta)\lambda^{*s})]^{T-j}}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})}. \tag{12}$$

Using the above and setting $a_t = a$, expression (10) for the utility-based measure of incidence becomes (for self 1)

$$-\alpha_c \frac{1-\delta^T}{1-\delta} a - (1-\beta) \sum_{j=1}^{T} \delta^{j-1} \frac{\partial a}{\partial p}$$

$$\times \left( -(1-d)\delta v_S \frac{1 - [(1-d)\delta(1 + (1-\beta)\lambda^{*s})]^{T-j}}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})} \right), \tag{13}$$

which can be rewritten as

$$-\alpha_c \frac{1-\delta^T}{1-\delta} a - (1-\beta)(1-d)(-v_s) \frac{\partial a}{\partial p} \frac{1}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})}$$

$$\times \left( \frac{1-\delta^T}{1-\delta}\delta + \delta^T \frac{1 - [(1-d)(1 + (1-\beta)\lambda^{*s})]^T}{1 - (1-d)(1 + (1-\beta)\lambda^{*s})} \right). \tag{14}$$

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

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

$$- (1-\beta)(1-d)\frac{-v_S}{\alpha_c p} \underbrace{\left( -\frac{\partial a}{\partial p}\frac{p}{a} \right)}_{\eta} \underbrace{\frac{\delta^T(1-\delta)}{1-\delta^T} \frac{1 - [(1-d)(1 + (1-\beta)\lambda^{*s})]^T}{1 - (1-d)(1 + (1-\beta)\lambda^{*s})} \frac{1}{1 - (1-d)\delta(1 + (1-\beta)\lambda^{*s})}}_{\text{length-of-horizon effect}}. \tag{15}$$

The quantity $-\alpha_c((1-\delta^T)/(1-\delta))a$ 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 (15) 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. 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.

---

[14] When interpreted literally, this assumption is inappropriate in our framework, since it contradicts our assumption of reinforcement. But more naturally interpreted, the assumption means that we ignore the *variation* in $v_S$ over time, and replace it with an average. Since, for the purposes of welfare analysis, the health consequences of smoking are much more important than the effect of smoking on the future pleasure of smoking, this does not seem to be too much of a stretch.

The latter term in expression (15) 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 the length-of-horizon effect is less than $((\delta^{T-1}(1-\delta))/((1-\delta^T)(1-(1-d)(1+(1-\beta)\lambda^{*s}))))$ 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 $\delta$'s, the above term will be small. For example, for $\delta = 0.9$, and assuming $\beta = 0.6$, $d = 0.6$, and $\lambda^{*s} = 0.7$—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 5.6% 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.

In our calibration, we will use the money equivalent of the *combined* discounted damage of a cigarette in all future periods, $H_S = \frac{(1-d)\delta}{1-(1-d)\delta} \frac{v_S}{\alpha_c}$. This gives the final expression for the adjustment factor:

$$1 - (1-\beta)\frac{-H_S}{p}\underbrace{\left(-\frac{\partial a}{\partial p}\frac{p}{a}\right)}_{\eta}\frac{1-(1-d)\delta}{1-(1-d)\delta(1+(1-\beta)\lambda^{*s})}. \tag{16}$$

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 Appendix B. 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 (Theorem 1 in Appendix A).[15]

## 5.2. Calibrating tax incidence—parameters

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 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

---

[15] Since we have made $v_S(a_t, S_t)$ constant, our calculations in the appendix do ignore the effect of naifs' optimism about future consumption on current consumption. By the intertemporal complementarity of consumption, naifs' optimism about the future restrains current overconsumption. This makes a tax less beneficial. Although we have no way of estimating this "optimism effect," we can prove that incidence is unambiguously lower for naifs than for sophisticates even if it is accounted for.

expressing it in current dollars gives a figure of US$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. Viscusi and Aldy (2003) show that studies of the value of life by age imply discount rates in the wide range of 1–17%; these estimated discount rates, of course, reflect both short and long-term discounting. We therefore consider long-term annual discount rates of 3% and 10% for our estimates.

   We use the fact that smokers die on average roughly 6 years earlier (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.[16]

   At these figures, the cost in terms of life years lost per pack of cigarettes is US$35.64 when the annual discount rate is 3%. Of course, this figure will vary with the agent's $\delta$ that we are considering; since costs are at the end of life, they will fall as $\delta$ falls. But it is an enormous figure for any reasonable $\delta$, 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. The evidence in Viscusi and Hersch (2001) suggests that this is true, in that the compensating differentials that smokers require for risky jobs are about half those of the differentials required by nonsmokers. 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 market behavior will in general not reflect true long-run valuations (which is what we need for our calibrations). 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.

   Another crucial parameter for our calibrations is the short-term discount factor $\beta$. $\beta$ is not only hard to estimate due to the lack of evidence on impatience in smoking, but also because it depends on the period length chosen for our analysis. We think of one period in

---

[16] Age 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.

our model as 1 month. In laboratory experiments, monthly $\beta$'s tend to lie between 0.6 and 0.9, so we calibrate incidence for both of these extremes (although most estimates are closer to 0.6).

A central variable for computing incidence is the price elasticity of demand for cigarettes, and 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 et al. (1999a,b) 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 using the Consumer Expenditure Survey (CEX), a nationally representative survey which provides the highest quality and most comprehensive micro-data on consumption in the US. We match 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 1. The coefficient estimates show the impact of a US$ 1 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¢, 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).

We estimate the impact of prices on consumption across income quartiles, consumption quartiles, and education categories (high school dropout, high school graduate with no college, some college but no bachelor's degree, and bachelor's degree or greater); the latter two are proxies for permanent income. 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 $-1$. 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.

---

[17] Data 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.

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

Table 1
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 graduates | 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 graduates | 0.341 (0.116) [− 0.387] | 0.215 (0.137) [− 0.642] | 0.175 (0.082) [− 0.400] |

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.

While we discussed above our computation of the health damage from a pack of cigarettes, we 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 evidence on the former is reviewed in Viscusi and Aldy (2003), and they suggest a consensus income elasticity of the value of life of 0.5, which we use here.

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 time preference parameters across income groups. We therefore initially assume that preference parameters are the same, although we discuss below the implications of varying these parameters across groups. Given the evidence in Gruber and Kőszegi (2001), 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.

## 5.3. Calibrating tax incidence—basic results

Putting together the results of the last two subsections, Table 2 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 2 by group. The figures in the table represent the incidence of a US\$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, as computed from the 1997–1998 CEX in Gruber

Table 2
Incidence calibrations

| Group | (1) | (4) | (7) | (10) | (13) |
|---|---|---|---|---|---|
| | Ex ante | $\beta = 0.9,$ $\delta = 0.97$ | $\beta = 0.6,$ $\delta = 0.97$ | $\beta = 0.9,$ $\delta = 0.9$ | $\beta = 0.6,$ $\delta = 0.9$ |
| *Current income distribution* | | | | | |
| Low | 1.69 | 0.56 | − 3.60 | 1.12 | − 0.98 |
| Second | 0.71 | 0.21 | − 1.62 | 0.46 | − 0.47 |
| Third | 0.47 | 0.13 | − 1.11 | 0.30 | − 0.32 |
| Top | 0.18 | 0.04 | − 0.48 | 0.11 | − 0.15 |
| *Expenditure distribution* | | | | | |
| Low | 0.83 | 0.18 | − 2.22 | 0.50 | − 0.71 |
| Second | 0.66 | 0.14 | − 1.80 | 0.40 | − 0.58 |
| Third | 0.49 | 0.29 | − 0.43 | 0.39 | 0.03 |
| Top | 0.22 | − 0.05 | − 1.05 | 0.08 | − 0.42 |
| *Education categories* | | | | | |
| HSD | 0.75 | − 0.07 | − 3.06 | 0.34 | − 1.17 |
| HSG | 0.66 | − 0.04 | − 2.62 | 0.31 | − 1.00 |
| SCL | 0.4 | 0.35 | 0.14 | 0.37 | 0.27 |
| CLG | 0.13 | 0.05 | − 0.25 | 0.09 | − 0.06 |

Table shows ex ante (column 1) and adjusted (columns 2–4) incidence of a US$1/pack rise in the tax on cigarettes. All estimates are for $d = 0.6$ and $\lambda^{*s} = 0.7$; corresponding $\beta$ and $\delta$ for each estimate is shown in top row. 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 Eq. (17) in text.

and Kőszegi (2002). The remainder of the results show incidence after applying our time inconsistency adjustment.

There are two important points for interpreting these results. 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 are so large 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 2, taxes are the most regressive ex ante, with the burden on the lowest income group almost 10 times that on the highest. In the first set of columns, we show the adjusted incidence measure for $\beta = 0.9$ and an annual discount rate of 3%. This has only a modest effect on the ratio measure of regressivity, but narrows the absolute gap considerably, with the gap between the highest and lowest income groups falling by two-thirds in absolute terms.

Varying $\beta$ has an enormous effect. For $\beta = 0.6$, taxes are now beneficial for every group, and much more so for the lowest income group. That is, taxes are now highly progressive, with very large benefits for the lowest income group, and much smaller benefits for the highest income group. This largely reflects the much higher price sensitivity of this bottom group.

The next two columns vary $\delta$. A lower value for $\delta$ significantly mitigates the self-control adjustment, since it serves to reduce the benefit of lengthening life (since those later years of life are more highly discounted).[19] For $\beta = 0.9$, there is little change in regressivity. But for a lower $\beta = 0.6$, taxes are once again progressive.

The conclusions are much stronger when we use more permanent measures of income in panels 2 and 3 of Table 2. Panel 2 uses the expenditure distribution rather than the income distribution. Here, we find a significant narrowing for $\beta = 0.9$, with the gap in incidence between the top and bottom groups shrinking from 0.61 in the first column to 0.23. With $\beta = 0.6$, once again, the tax becomes highly progressive using this distributional metric. For lower $\delta$, there is once again relatively little effect when we have $\beta = 0.9$, but the taxes once again become (moderately) progressive with $\beta = 0.6$.

Perhaps the strongest findings are obtained with a different proxy for permanent income: education. The results for this proxy for permanent income are much more modest, as shown in panel 3. This is because the income differences are much more modest by education group, so that there are much more uniform life values by group than the other two panels. In this case, taxes are progressive in three of the four cases, and modestly regressive in the case where $\beta = 0.9$ and an annual 10% discount rate.

Thus, while these calibrations are sensitive to parameter values and assumptions, the overall message is clear: cigarette excise taxes are much less regressive, and potentially progressive, when we account for the self-control benefits of taxation. 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.

As discussed earlier, a key sensitivity for these calculations is to the source of higher smoking rates, and lower values of life, for the poor. If the poor smoke more because they have higher long-term discount factor $\delta$, for example, then this would increase the regressivity of the results relative to Table 2. On the other hand, if the poor smoke more because they have a higher short-term discount factor, $\beta$, then taxes are even more progressive than shown in Table 2. While we have no evidence on the relative values of the short- and long-term discount factors across income groups, it is worth noting that the effect of varying $\beta$ is much stronger than the effect of varying $\delta$. We can illustrate this by varying these two parameters, holding all else constant, in a manner consistent with the revealed differences in values of life across the rich and the poor. For example, if the long-term annual discount rate for the poor is 33% when the discount rate for the rich is only 10%—which is consistent with the differences in value of life implied by an elasticity of 0.5—then the regressivity of taxation is very similar to that shown in our adjusted figures in Table 2. If the annual discount rate of the poor is 15% when the discount rate for the rich is 3%—which is also consistent with these value of life differences—then the burden on the poor is twice as large as that shown in these Table 2

---

[19] We vary delta by choosing a higher value of a life-year so that the discounted value of life still equals the US$6.7 billion figure. But, even with these higher life-year values, the benefit of increasing life by 6 years at the end of life is greatly reduced with a higher discount rate.

adjusted figures (so that the burden on the rich is zero, and the burden on the poor is 1.1% of income).

On the other hand, if $\beta = 0.75$ for the poor and $\beta = 0.9$ for the rich—which accounts for only about one-tenth of the difference in the values of life—taxes become highly progressive, with a burden on the rich of 0.03% of income, and a burden on the poor of negative 2.3% of income. And if $\beta = 0.6$ for the poor and $\beta = 0.9$ for the rich—which accounts for about one-quarter of the difference in the values of life—then the burden on the poor rises to negative 7% of income. 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.

The calibrations presented in Table 1 also impose a constant effect of taxes by age. In fact, the incidence of taxes will vary with age. In particular, the self-control benefits of taxation will become larger for older smokers, for whom the damage of smoking is nearest. One particularly interesting group to focus on is teenagers, who are considering starting smoking. We have not included this group in our calibrations, since they clearly violate one of our important simplifying assumptions, that the agent starts from a steady-state level of consumption. But we can comment qualitatively on how overall incidence, and the regressivity of taxes, would be affected by the self-control adjustment. On the one hand, the damage of smoking is most distant for this group, so that the self-control adjustment will be the smallest, and taxes will be the most regressive. On the other hand, those who are initiating smoking have been shown to be the most price sensitive, with estimated elasticities on the order of 25–50% higher than the average across all ages (Gruber and Zinman, 2001). Moreover, if the tax prevents teenagers from smoking altogether, it has no standard incidence at all. Thus, all of its effect comes from the self-control adjustment, which is likely to be higher for lower income groups.

A natural comparison for thinking about the broader range of public policy is between tax increases and regulations, such as clean air laws that restrict smoking in public places. As we highlight in Gruber and Kőszegi (2000), our model suggests that clean air laws are justified not simply as a means of controlling second-hand smoke, but as another instrument of self-control. This implies that they have lower incidence across all groups than would be computed from a traditional analysis. How much lower depends on the elasticity of smoking with regards to these regulations; evidence on this point is mixed, with Gruber and Zinman (2001) finding little effect for youths, but Evans et al. (1999a,b) finding large effects of specific workplace bans on smoking. In addition, modeling the incidence of regulations would require information both on their effects throughout the income distribution (e.g. do they impact places where particularly low income groups work or congregate?), and on the relative sensitivity of smoking to regulations throughout the income distribution. If regulations impose their largest costs in terms of time lost to compliance (e.g. due to moving outside for a cigarette break), and if lower income groups have the lowest value of their time, then it is possible that the regulations could have the smallest behavioral impact on low income groups, in contrast to taxes. In this case, regulations might be much more regressive than are tax increases. This is clearly an important point for future research.

## 6. 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 5 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 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, lowering the incidence of the tax. 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 long-run 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 dwarfs 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 US$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 individual decision-making 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.

## Acknowledgements

## Appendix A. Proofs

**Lemma 1.**

*1. $\lambda_t$ is decreasing in $\beta$ for each t.*
*2. $\lambda_t$ is decreasing in $\delta$ for each t.*
*3. $\lambda_t$ is increasing in t.*

**Proof.** Since the utility function is quadratic, we have

$$S_{t+1} = (1-d)(S_t + a_t) = (1-d)(S_t + \lambda_t S_t + \mu_t) \tag{17}$$

$$a_{t+1} = \lambda_{t+1} S_{t+1} + \mu_{t+1} = \lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1} \tag{18}$$

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

$$
\begin{aligned}
\alpha_a + \alpha_{aa}(\lambda_t S_t + \mu_t) + \alpha_{as} S_t - p\alpha_c &= (1-d)\delta[(1 + (1-\beta)\lambda_{t+1}) \\
&\quad \times [\alpha_a + \alpha_{aa}(\lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1}) \\
&\quad + \alpha_{as}(1-d)(S_t + \lambda_t S_t + \mu_t) - p\alpha_c] \\
&\quad - \beta(\alpha_s + \alpha_{as}(\lambda_{t+1}(1-d)(S_t + \lambda_t S_t + \mu_t) + \mu_{t+1}) \\
&\quad + \alpha_{ss}(1-d)(S_t + \lambda_t S_t + \mu_t))]
\end{aligned}
\tag{19}
$$

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

$$\lambda_t = -1 + \frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \delta(1-d)^2[(1 + (1-\beta)\lambda_{t+1})(\alpha_{aa}\lambda_{t+1} + \alpha_{as}) - \beta\alpha_{as}\lambda_{t+1} - \beta\alpha_{ss}]} \tag{20}$$

Define the function $f_s(\lambda)$ according to Eq. (20). 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 $\beta$ and $\delta$.

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 $\lambda^2$. 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

$$f_s(-1) = -1 + \frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \beta\delta(1-d)^2[-\alpha_{aa} + 2\alpha_{as} - \alpha_{ss}]} > -1 \tag{21}$$

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

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

$$= \frac{\alpha_{as} - \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]}{-\alpha_{aa} + \delta(1-d)^2[-\beta\alpha_{as}\lambda_T - \beta\alpha_{ss}]}. \tag{22}$$

This being $< \lambda_T = (\alpha_{as}/-\alpha_{aa})$ 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,

$$f_s'(\lambda) = -\frac{(\alpha_{as} - \alpha_{aa})(2(1-\beta)\lambda\alpha_{aa} + \alpha_{aa} + (1-2\beta)\alpha_{as})}{[-\alpha_{aa} + \delta(1-d)^2[(1+(1-\beta)\lambda)(\alpha_{aa}\lambda + \alpha_{as}) - \beta\alpha_{as}\lambda - \beta\alpha_{ss}]]^2}, \tag{23}$$

which gives

$$f_s'(-1) = -(1 - 2\beta)$$

$$\times \left[\frac{\alpha_{as} - \alpha_{aa}}{-\alpha_{aa} + \delta(1-d)^2[(1+(1-\beta)\lambda)(\alpha_{aa}\lambda + \alpha_{as}) - \beta\alpha_{as}\lambda - \beta\alpha_{ss}]}\right]^2 \geq 0. \tag{24}$$

---

[20] The second derivative of the reciprocal of a quadratic $q$ is $(-(q^2 q'' - q(q')^2)/q^4)$, which is positive as long as $q$ is positive and concave.

Finally, we prove that $f_s(\lambda_t)$ is decreasing in $\beta$ and $\delta$. For $\delta$, this is quite simple: the term multiplying $\delta$ in the denominator is positive since $\lambda_t \leq \lambda_T$ and the utility function is strictly concave. It takes slightly more work to prove that $f_s(\lambda_t)$ is decreasing in $\beta$. The derivative of the denominator with respect to $\beta$ is a positive constant times

$$-[\alpha_{aa}\lambda_{t+1}^2 + 2\alpha_{as}\lambda_{t+1} + \alpha_{ss}]. \tag{25}$$

Notice that this is equal to

$$-(\lambda_{t+1}\, 1)\begin{pmatrix} \alpha_{aa} & \alpha_{as} \\ \alpha_{as} & \alpha_{aa} \end{pmatrix}\begin{pmatrix} \lambda_{t+1} \\ 1 \end{pmatrix}, \tag{26}$$

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

   This completes the proof.   □

**Theorem 1.** $\lambda^{*n} \geq \lambda^{*s} \geq \lambda^{*TC}$.

**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

$$(1-\beta)[\lambda^2(\alpha_{aa} + p^2\alpha_{cc}) + 2\alpha_{as}\lambda + \alpha_{ss}]$$

$$= (1-\beta)(\lambda\, 1\, p\lambda)\begin{pmatrix} \alpha_{aa} & \alpha_{as} & 0 \\ \alpha_{as} & \alpha_{ss} & 0 \\ 0 & 0 & \alpha_{cc} \end{pmatrix}\begin{pmatrix} \lambda \\ 1 \\ p\lambda \end{pmatrix} \leq 0 \tag{27}$$

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:

$$V^s(S_{t+1}) = \sum_{i=1}^{T-t} \delta^i U(a_{t+i}^s, c_{t+i}^s, S_{t+i}^s)$$

$$V^{TC}(S_{t+1}) = \sum_{i=1}^{T-t} \delta^i U(a_{t+i}^{TC}, c_{t+i}^{TC}, S_{t+i}^{TC}), \tag{28}$$

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]

---

[21] The 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 $S_{t+1}$.

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

$$\max_{a_t} v(a_t, S_t) + u(I_t - pa_t) + \beta V^i((1 - d)(S_t + a_t)), \tag{29}$$

leading to the first-order condition

$$v_a(a_t, S_t) - pu'(c_t) + \beta(1 - d)V_S^i(S_{t+1}) = 0. \tag{30}$$

Differentiating this totally with respect to $S_t$ gives

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

$$= -1 + \frac{v_{as}(a_t, S_t) - v_{aa}(a_t, S_t) - p^2 u''(c_t)}{-v_{aa}(a_t, S_t) - p^2 u''(c_t) - \beta\delta(1 - d)^2 V_{SS}^i(S_{t+1})}. \tag{31}$$

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

## Appendix 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

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

$$\left. \times [(v_a(a_{t+k}, S_{t+k}) - p\alpha_c)\lambda^{*n} + v_s(a_{t+k}, S_{t+k})] \right]. \tag{32}$$

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

$$v_a(a_t, S_t) - p\alpha_c = -\beta \frac{\delta(1-d)}{1-\delta(1-d)} v_S. \tag{33}$$

Combining this with the rest of the above expression yields

$$-\alpha_c \sum_{t=1}^{\infty} \delta^t a_t + \sum_{t=1}^{\infty} \delta^t \frac{\partial a_t}{\partial p} \left[ \frac{\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})} v_S - \beta \frac{\delta(1-d)}{1-\delta(1-d)} v_S \right.$$
$$\left. - \beta \frac{\delta(1-d)}{1-\delta(1-d)} \frac{\delta(1-d)\lambda^{*n}}{1-\delta(1-d)(1+\lambda^{*n})} v_S \right]. \tag{34}$$

Rearranging gives

$$-\alpha_c \sum_{t=1}^{\infty} \delta^t a_t + \sum_{t=1}^{\infty} \delta^t \frac{\partial a_t}{\partial p} (1-\beta) \frac{\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})} v_S. \tag{35}$$

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

$$1 - (1-\beta) \frac{-\partial a}{\partial p} \frac{p}{a} \frac{-h_S}{p} \frac{\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})}. \tag{36}$$

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

$$1 - (1-\beta) \frac{-\partial a}{\partial p} \frac{p}{a} \frac{-H_S}{p} \frac{1-\delta(1-d)}{1-\delta(1-d)(1+\lambda^{*n})}. \tag{37}$$

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, $\lambda^{*n}$. 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 1), 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 Univ. Press, Cambridge, UK.

Ainslie, G., Haslam, N., 1992. Hyperbolic discounting. In: Loewenstein, G., Elster, J. (Eds.), Choice Over Time. Russell Sage Foundation, New York, NY, pp. 57–92. Chap. 3.

Angeletos, G.-M., Laibson, D.I., Repetto, A., Tobacman, J., Weinberg, S., 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.S., Murphy, K.M., 1988. A theory of rational addiction. Journal of Political Economy 96 (4), 675–700.

Bernheim, B.D., Rangel, A., 2001. Addiction, conditioning, and the visceral brain. Mimeo, Stanford University.

Bernstein, D.A., 1970. The modification of smoking behavior: an evaluative review. In: Hunt, W.A. (Ed.), Learning Mechanisms in Smoking. Aldine Publishing, Chicago, pp. 3–41.

Burns, J., 1992. Looking to the future. In: Burns, J. (Ed.), Special Report: Business and Health. Medical Economics Publishing, Washington, D.C, pp. 21–22.

CDC (various years). You Can Quit Smoking, Tobacco Information and Prevention Source, Webpage: http://www.cdc.gov/tobacco.

Chaloupka, F.J., Warner, K.E., 2000. The Economics of Smoking. In: Newhouse, J., Culyer, A. (Eds.), Handbook of Health Economics 16, 1539–1628.

Cutler, D., Gruber, J., Hartman, R., Landrum, M.B., Newhouse, J.P., Rosenthal, M., 2001. The economic impacts of the tobacco settlement. Journal of Policy Analysis and Management 21 (1), 1–19.

Della Vigna, S., Malmendier, U., 2001. Self-control in the market: evidence from the Health Club Industry. Mimeo, Harvard University.

Diamond, P., Kőszegi, B., 2002. Quasi-hyperbolic discounting and retirement. Journal of Public Economics (in press, available on line).

Evans, W., Farelly, M., Montgomery, E., 1999a. Do workplace smoking bans reduce smoking? American Economic Review 89, 728–747.

Evans, W., Ringel, J., Stech, D., 1999b. Tobacco taxes and public policy to discourage smoking. In: Poterba, J. (Ed.), Tax Policy and the Economy. MIT Press, Cambridge, MA, pp. 1–56.

Grabowski, J., Hall, S.M., 1985. Tobacco use, treatment strategies, and pharmacological adjuncts: an overview. In: Grabowski, J., Hall, S.M. (Eds.), Pharmacological Adjuncts in Smoking Cessation. National Institute on Drug Abuse Monograph, vol. 53. United States Govt. Printing Office, Rockville MD, Washington D.C, pp. 1–13.

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., Kőszegi, B., 2000. Is addiction 'rational?' Theory and evidence. NBER Working Paper 7507.

Gruber, J., Kőszegi, B., 2001. Is addiction 'rational?' Theory and evidence. Quarterly Journal of Economics 116 (4), 1261–1305.

Gruber, J., Kőszegi, B., 2002. A Theory of Government Regulation of Addictive Bads: Optimal Tax Levels and Tax Incidence for Cigarette Excise Taxation. NBER working paper 8777.

Gruber, J., Mullainathan, S., 2001. Do cigarette taxes make smokers happier? Mimeo, Massachusetts Institute of Technology.

Gruber, J., Zinman, J., 2001. Youth smoking in the U.S.: evidence and implications. In: Gruber, J. (Ed.), Risky Behavior Among Youth: An Economic Analysis. University of Chicago Press, Chicago, pp. 69–120.

Gul, F., Pesendorfer, W., 2001. Temptation and self-control. Econometrica 69 (6), 1403–1435.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Quarterly Journal of Economics 112, 443–477.

Laibson, D.I., Repetto, A., Tobacman, J., 1998. Self-control and saving for retirement. In: Brainard, W.C., Perry, G.L. (Eds.), Brookings Papers on Economic Activity, vol. 1. The Brookings Institution, Washington D.C, pp. 91–196.

Manning, W.G., et al., 1989. The taxes of sin. The Journal of the American Medical Association 261 (11), 1064–1069.

Manning, W.G., et al., 1991. The Costs of Poor Health Habits. Harvard University Press, Cambridge MA.

O'Donoghue, T., Rabin, M., 1999a. Doing it now or later. American Economic Review 89 (1), 103–124.

O'Donoghue, T., Rabin, M., 1999b. Procrastination in preparing for retirement. In: Aaron, H.J. (Ed.), Behavioral Dimensions of Retirement Economics. Brookings Institution Press, Washington D.C., pp. 125–156.

O'Donoghue, T., Rabin, M., 2003. Studying optimal paternalism, illustrated by a model of sin taxes. Mimeo, UC Berkeley.

Pesendorfer, W., Gul, F., 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., Crimi, P., Lapsanski, D., Martel, L., Reid, P., 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., 1995. Cigarette taxation and the social consequences of smoking. In: Poterba's, J. (Ed.), Tax Policy and the Economy. MIT Press for the National Bureau of Economic Research, Cambridge, pp. 51–101.

Viscusi, K., Aldy, J., 2003. The value of a statistical life: a critical review of market estimates from throughout the world. National Bureach of Economic Research Working Paper No. 9487.

Viscusi, W.K., Hersch, J., 2001. Cigarette smokers as job risk takers. Review of Economics and Statistics 83, 269–280.