WORKING
PAPER

Time-Inconsistency and Welfare

JAY BHATTACHARYA, DARIUS LAKDAWALLA

WR-158

April 2004

HEALTH

NBER WORKING PAPER SERIES

TIME-INCONSISTENCY AND WELFARE

Jay Bhattacharya
Darius Lakdawalla

Working Paper 10345
http://www.nber.org/papers/w10345

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2004

We have benefited from discussions with Steve Garber, Michael Grossman, Jim Hosek, David Loughran, John Romley, Neeraj Sood, Eric Talley, and participants at the Spring 2003 NBER Health Economics Program Meeting. Lakdawalla is grateful for the financial support of RAND's Institute for Civil Justice. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

Time-Inconsistency and Welfare
Jay Bhattacharya and Darius Lakdawalla
NBER Working Paper No. 10345
March 2004
JEL No. I1, H2, D6

ABSTRACT

Self-control devices, such as rehabilitation programs, group commitment, and informal fines, can make time-inconsistent smokers better off. Health economists have used this result to argue in favor of cigarette taxes that restrain smoking. However, taxes alone are not Pareto-improving overall, because they benefit today's smoker at the expense of her future selves, who have less demand for self-control. We suggest an alternative class of taxation policies that provide self-control and benefit a smoker at every point in life. Smokers could be allowed to purchase "smoking licenses" when they start to smoke, and in exchange commit their future selves to face compensated cigarette taxes. We show that this scheme – which could be made voluntary – improves the welfare of current and future smokers, generates positive revenue for the government, and can be made incentive-compatible. Similar schemes can also be envisioned to address problems of time-inconsistency in other contexts.

Jay Bhattacharya
Stanford University
Center for Primary Care and Outcomes Research
117 Encina Commons
Stanford, CA 94305-6019
and NBER
jay@stanford.edu

Darius Lakdawalla
RAND Corporation
1700 Main Street
Santa Monica, CA 90407
and NBER
darius@rand.org

1 Introduction

Traditional intertemporal economic models feature agents that fully account for the effects of current decisions on future outcomes and constraints. Such agents make time-consistent decisions in the sense that today's plans for future consumption will not change tomorrow. Government intervention aimed at changing intertemporal prices can only harm them.

In recent years, this view has been challenged by economists who argue that people make time-inconsistent decisions, in various intertemporal contexts including savings, health investment, and smoking.[1] While the analysis of time-inconsistency has a long history in economics, Laibson (1994, 1997) rekindled interest in it as a means of understanding anomalies in intertemporal decisionmaking.[2] Since time-inconsistent agents do not make intertemporally optimal decisions, government manipulation of relative prices could potentially improve welfare. At a minimum, time-inconsistency challenges the received doctrine of laissez-faire in competitive markets.

However, this theoretical result does not so easily translate into policy prescriptions. Since a time-inconsistent individual has different preference orderings at different points in her life, she also has different preferences for policy. The Paretian approach to this problem adopts only those policies that make an individual better off at all points in time; we will call such

[1] See Lowenstein and Prelec (1992) for a start.

[2] See Strotz (1955) and Phelps and Pollak (1968) for early work by economists on myopic behavior.

policies *Pareto self-improving*.[3] While nearly all economists would agree that a Pareto self-improving policy ought to be adopted, it is extremely difficult to find a tax or subsidy policy that is Pareto self-improving for a time-inconsistent agent. The value of a tax is as a commitment device, but the value of such a device decreases with age. For example, an agent at the end of life has little or no interest in future commitment, even though she may have clamored for it 40 years earlier. This vexing result obtains even though time-inconsistent agents make demonstrably inefficient decisions.

In the face of this problem, economists have chosen to depart from the Pareto criterion in various ways. For instance, in their analysis of smoking and cigarette taxation Gruber and Koszegi (2004, 2001) have advocated a "dictatorship of the present" that privileges the preferences of the current period self over all future selves; the same approach is suggested by Cropper and Laibson (1998) in the context of policy evaluation. Alternatively, O'Donoghue and Rabin (1999) choose to privilege the "long-run self" who is constituted by equally weighting the preferences of all temporal selves. These or other approaches are reasonable and defensible, but they lack the incontrovertibility of the Pareto criterion, because the researcher must ultimately make a judgement about which self or selves to favor.

A more attractive approach is to characterize policies that are Pareto

[3]To our knowledge, Phelps and Pollak (1968) are the first authors to adopt a Pareto criterion in an analysis of time-consistency, although they focus on an intergenerational version. Goldman (1979) implicitly assumes such a criterion for time-inconsistent agents, while Laibson (1994) discusses it in the context of a time-inconsistent savings model.

self-improving for time-inconsistent individuals; this is the strategy we adopt in this paper. In particular, we propose a scheme of "consumption licenses" that we show to be Pareto self-improving. These can be used to improve the welfare of time-inconsistent smokers and addicts, to encourage more efficient savings behavior by time-inconsistent individuals, or to ameliorate other problems of time-inconsistent decisionmaking.

Taxes or subsides alone are too blunt to be Pareto self-improving. Since the demand for taxes and other self-control devices declines as time-horizons shorten, it is inevitable that selves close enough to the end of their lives will be hurt by taxes, even though younger selves may benefit. Cross-sectionally, any given level of taxation is likely to benefit younger consumers at the expense of older consumers. The solution to this problem is to shift the cost of taxation entirely onto the people with the highest willingness to pay, namely the youngest selves, while continuing to target the effect of taxation toward the future selves.

Consider the example of cigarette taxation for time-inconsistent smokers. Suppose an individual were allowed to purchase a smoking "license" early in life if she chooses to be a smoker. This license binds her to face a *compensated* cigarette tax in the future. While the future self is forced to face a tax that she did not choose, she is exactly compensated for it by means of a lump-sum transfer that is fixed in advance. This leaves the future self indifferent between having the license and not having it. Moreover, the current

self values taxation as a self-control device and is willing to pay for it. She is better off with the license, even though she paid for it. This scheme improves the welfare of a time-inconsistent individual at every point in time and generates revenue that can be used for administrative costs, general revenue, or for further welfare improvement. The government is creating a market in which the current self can pay to restrain the future self. The creation of this market generates social surplus. As discussed in our conclusions, similar schemes can be envisioned for encouraging retirement savings or mitigating other problems of time-inconsistency.

An important advantage of this system is that any time-inconsistent smoker would voluntarily opt in for some positive license fee and tax. As a result, it can be selectively applied to smokers with a self-control problem, without affecting clear-sighted rational addicts. Moreover, smokers with more severe self-control problems can choose to pay the highest license fees and tax rates. The costs of these Pareto self-improving schemes lie in difficulties of enforcement and compliance; in Section 5 we discuss the construction of incentive-compatible licenses and mechanisms for enforcement. As we discuss there, cheating in our scheme is relatively easy to observe *ex post*, and this affords the government some leverage in deterring such behavior.

2 Time-Inconsistency and Welfare

Loosely speaking, time-inconsistency under-weights the future in favor of the present. A myopic individual lacks intertemporal integrity in the sense that her preferences today conflict with her preferences tomorrow. But what does it mean to make a time-inconsistent individual better off when her different temporal incarnations disagree about what is better or best?

It is problematic that the different incarnations of a myopic individual can express meaningful preferences about only current and future consumption.[4] Different temporal selves thus make decisions over different commodity spaces. The difficulty with defining a welfare criterion in this context is well described by Goldman (1979):

> The question of Pareto efficiency is especially vexing... If we were to identify the players...by both calendar time and the history of actions prior to the times of their decisions, then Pareto comparisons under alternative histories become impossible since the set of players changes.

There is an obvious way around this difficulty that several authors, including Goldman (1979), have adopted. We can define a *Pareto self-improving* intervention as an intervention that makes every temporal self at least as well off as before the intervention. Just as in the traditional inter-personal Pareto

[4]Preferences are "meaningful" here if they can be acted upon. While some may wish to change the past and may sincerely regret errors made, past consumption decisions cannot be undone.

criterion, the Pareto self-improvement criterion avoids the problem identified by Goldman (1979) by evaluating each self on her own terms.

Different authors have argued for other welfare criteria, all of which necessarily privilege different temporal selves over others. Cropper and Laibson (1998) assume that "the goal of the government at time t is to maximize the well-being of self t." This approach is also adopted by Gruber and Koszegi (2001). Caplin and Leahy (2000) call this welfare criterion the "dictatorship of the present," since it calls for government policy to support the interests of the current self even if future selves disagree.

O'Donoghue and Rabin (1999), on the other hand, advocate maximizing the welfare of the "long-run" self.[5] This involves positing the existence of a fictitious period zero self, who makes no decisions but weighs the utility of all the other myopic selves equally. O'Donoghue and Rabin argue that because myopic individuals have present-biased preferences, their welfare ought to be evaluated in a future-oriented way. The long-run preference criterion accomplishes this goal, but like the "dictatorship of the present" welfare criterion, it may condone an intervention even if some selves are made worse off. However, as O'Donoghue and Rabin point out: "When applied to intertemporal choice, the Pareto criterion often refuses to rank two strategies even when one is much preferred by virtually all period selves, while the other is preferred by only one period self."

[5]Gruber and Koszegi (2001) consider both dictatorship of the present and long-run welfare criteria.

Clearly, adopting a Pareto self-improvement criterion limits the set of interventions that can be said to improve outcomes for the time-inconsistent relative to the more paternalistic criteria. Nonetheless, a major theme of our paper is that even with such a restrictive welfare criterion, there are still interventions that can be said to improve outcomes for the time-inconsistent. Moreover, a Paretian criterion has well-known advantages worth restating. It avoids privileging the welfare of one self over another, and a Pareto self-improving intervention has the virtue that it could be voluntarily chosen. Finally, the policies developed under the Pareto self-improvement criterion would be adopted under the "dictatorship of the present" and long-run welfare criteria as well. This suggests a viable research strategy: search first for interventions that satisfy Pareto self-improvement, and relax welfare criteria accordingly if this fails to be fruitful.

3 A Model of Time-Inconsistency

In traditional intertemporal economic models, future returns are discounted exponentially; if tomorrow's return is discounted at a rate $\beta \in (0, 1)$, then the day after tomorrow's return is discounted at a rate β^2, and so on. Adapting an approach to intergenerational discounting from Phelps and Pollak (1968), Laibson (1994) incorporates time-inconsistency into this framework by introducing a second discounting parameter—$\delta \in (0, 1)$—that applies to all future returns, except those in the immediate future. A time-inconsistent

agent discounts tomorrow at β, the next day at $\beta\delta$, and the next day at $\beta^2\delta$. Heuristically, myopic agents discount the future "too much," at the rate δ.

We lay out our "consumption license" proposal in the context of smoking behavior, or more generally, addiction. However, we discuss in our conclusion how our results for smoking licenses have natural applications to savings and other problems of time-inconsistency. We analyze a finite-period model to bring into relief the problem that individuals at the end of life have very different demands for self-control than their earlier incarnations. For purposes of clarity and tractability, we limit this model to three periods, the minimum length necessary to illustrate time-inconsistency. In each period, individuals allocate their (constant) income to cigarettes, other non-addictive consumption goods, and saving. We allow both borrowing and lending.

We model time-inconsistency as hyperbolic discounting: next period's utility is discounted by some factor $\beta\delta$, while the following period's utility is discounted by $\beta^2\delta$. While $\beta < 1$ for all individuals, $\delta < 1$ for the time-inconsistent, and $\delta = 1$ for the time-consistent. Let the time-inconsistent smoker's period utility at time t be given by u_t. Since there are three periods, there are three selves, each with a different utility function, U_t:

$$U_3 = u_3 \tag{3.1}$$

$$U_2 = u_2 + \delta\beta u_3 \tag{3.2}$$

$$U_1 = u_1 + \delta\beta u_2 + \delta\beta^2 u_3 \tag{3.3}$$

We consider a general period utility function that satisfies the conventional assumptions of addiction. Period utility u_t depends on consumption c_t, the stock of addiction S_{t-1}, and cigarettes smoked a_t. We make the usual assumptions that addiction lowers the level of utility, so that $\frac{\partial u_t}{\partial S_{t-1}} < 0$, but that it raises the marginal utility of smoking, so that $\frac{\partial^2 u_t}{\partial S_{t-1}\partial a_t} > 0$. Each period, individuals receive income I and choose their net savings v_t, which can be negative if they choose to borrow. All debts must be cleared in the final period. The stock of addiction depreciates at the rate γ and evolves according to:

$$S_t = (1-\gamma)S_{t-1} + a_t. \tag{3.4}$$

In the terminology of Laibson (1997), there are (at least) two fundamentally different types of time-inconsistent individuals, sophisticates and naifs, distinguished by their extent of self-awareness. A sophisticated agent realizes that she will betray herself in the future, and thus undertakes actions now which restrict future behavior. A naive agent betrays her old selves each

period, but makes plans and choices blissfully unaware that she will betray herself again.

In this paper, we focus our attention on sophisticates because there are no Pareto self-improving interventions that naifs would condone *ex ante*. Time-inconsistent individuals who do not understand their self-control problem never have positive demand for self-control devices; therefore, taxes and other such devices never improve their utility, at least according to the myopic preferences they themselves possess.

Matters are a bit more complicated than this, however, because the lifetime utility of naifs can be unambiguously increased by assigning them the consumption path chosen by sophisticates. This is because their actual preferences match those of sophisticates, but naifs themselves fail to account for this. Nonetheless, *ex ante*, naifs would reject an offer consisting of a sophisticate's consumption path, because they misunderstand their own preferences. The welfare of a naif compelled to accept the sophisticate's consumption bundle would be *retrospectively* Pareto self-improved in the sense that she could observe at the end of her life that she had been made better off. However, *prospectively* (i.e., when offered the bundle) she would not feel as if her welfare would be improved by the new bundle. The Paretian analysis in neoclassical economics is concerned with prospective, or forward-looking, welfare.[6] As a result, it is not possible to Pareto self-improve the welfare of

[6]Caplin and Leahy (2000) take issue with the focus of neoclassical theory on forward-looking behavior, and argue that preferences at any point in time are best specified over

a naif by assigning her the sophisticate's consumption bundle, or by providing a self-control device that would be valued by a sophisticate. The policy implications of naive time-inconsistency are discussed further in Section 6.

Sophisticated agents who understand their self-control problem take steps to combat it, although their measures are imperfect, and their utility can be enhanced by the provision of self-control devices. Optimal life-cycle decisions for sophisticates represent a subgame-perfect equilibrium, which can be derived via backwards induction. Accordingly, we begin our analysis of the problem in the third and final period, where the individual takes as given the stock of addiction S_2 and net savings v_2, chosen in the previous period. She solves:

$$\max_{a_3} u_3(v_2 + I - pa_3, S_2, a_3) \tag{3.5}$$

Optimal smoking in the final period satisfies:

$$-p\frac{\partial u_3}{\partial c} + \frac{\partial u_3}{\partial a} = 0 \tag{3.6}$$

This defines a policy function $a_3(S_2, v_2)$ that the period 2 agent takes as given, according to the following problem:

---

an individual's entire life, including the past.

$$\max_{v_2,a_2} u_2(v_1 + I - v_2 - pa_2, S_1, a_2) + \beta\delta u_3(I + v_2 - pa_3(S_2, v_2), S_2, a_3(S_2, v_2))$$
$$s.t.\, S_2 = (1-\gamma)S_1 + a_2 \tag{3.7}$$

Exploiting the first order condition that characterizes the policy function a_3, the first order conditions for period 2 smoking and savings simplify to:

$$\begin{aligned} -p\frac{\partial u_2}{\partial c} + \frac{\partial u_2}{\partial a} + \beta\delta\frac{\partial u_3}{\partial S} &= 0 \\ \frac{\partial u_2}{\partial c} &= \beta\delta\frac{\partial u_3}{\partial c} \end{aligned} \tag{3.8}$$

So far, none of these first order conditions departs from the time-consistent conditions, because this is essentially a two-period model.

With the solutions to these later problems in hand, we can characterize the behavior of the self in period 1, and observe the effects of time-inconsistency on smoking and savings. The period 1 self takes as given the

policy functions of the later selves and solves:

$$\begin{aligned}
&\max_{a_1,v_1} u_1(I - pa_1 - v_1, S_0, a_1) + \\
&\quad \beta\delta u_2(I + v_1 - v_2(S_1,v_1) - pa_2(S_1, v_1), S_1, a_2(S_1, v_1)) + \\
&\quad \beta^2\delta u_3(I - pa_3(S_1, v_1) + v_2(S_1, v_1), (1-\gamma)S_1 + a_2(S_1, v_1), a_3(S_1, v_1)) \\
&\quad\quad s.t.\ S_1 = (1-\gamma)S_0 + a_1
\end{aligned} \tag{3.9}$$

The optimal savings and smoking conditions illustrate the nature of time-inconsistency in this framework.

$$\begin{aligned}
&\overbrace{-p\frac{\partial u_1}{\partial c} + \frac{\partial u_1}{\partial a} + \beta\delta\frac{\partial u_2}{\partial S} + \beta^2\delta^2(1-\gamma)\frac{\partial u_3}{\partial S}}^{Time-Consistent\ FOC} + \\
&\overbrace{\beta^2\delta(1-\delta)\left(\frac{\partial u_3}{\partial S}\frac{dS_2}{dS_1} + \frac{\partial u_3}{\partial c}\frac{\partial v_2}{\partial S_1}\right)}^{Self-Discipline\ Incentive} = 0
\end{aligned} \tag{3.10}$$

$$\overbrace{\frac{\partial u_1}{\partial c} - \beta\delta\frac{\partial u_2}{\partial c}}^{Time-Consistent\ FOC} = \overbrace{\beta^2\delta(1-\delta)\frac{\partial u_3}{\partial S}\frac{\partial a_2}{\partial v_1}}^{Under-Saving\ Incentive} + \overbrace{\beta^2\delta(1-\delta)\frac{\partial u_3}{\partial c}\frac{\partial v_2}{\partial v_1}}^{Over-Saving\ Incentive} \tag{3.11}$$

Equation 3.10 characterizes the optimal smoking decision, while equation 3.11 characterizes optimal savings. The first four terms in equation 3.10 represent the lifetime marginal utility of smoking in a standard time-consistent problem: this includes the current net marginal utility of smoking, plus the

future marginal disutility of building up an addictive stock. The last term, which is zero for a time-consistent individual with $\delta = 1$, represents the attempt of a time-inconsistent agent to discipline his future selves.

Self 1 recognizes the negative impact of self 2's smoking decisions on the capital stocks inherited by self 3; the negative impact is summarized by the term in large parentheses.[7] Self 1 would like self 2 to pass more capital goods onto self 3; instead, self 2 chooses to pass along fewer capital goods (savings) and more capital bads (addiction). Since self 1 knows this, she disciplines self 2 by cutting back on smoking today and thus giving self 2 less incentive to smoke. This shifts some of self 2's investments out of smoking and into savings and consumption.

While the sophisticated agent cuts back smoking to discipline her future self, the net effect of time-inconsistency on cigarette consumption is ambiguous because of its effect on savings. For a rational individual, the right-hand side of 3.11 would be zero, but the sign for a time-inconsistent individual is ambiguous. There are two offsetting effects, which appear on the right-hand side. On the one hand, the period 1 self would like to encourage more savings in period 2 by saving more herself. This is the "Over-Saving Incentive," by which the time-inconsistent agent saves more than the time-consistent agent. On the other hand, saving more today will allow the period 2 self to smoke more; this "Under-Saving Incentive" gives the period 1 self incentives

[7]The term $\frac{dS_2}{dS_1}$ (which equals $(1-\gamma) + \frac{\partial a_2}{\partial S_1} > 0$) represents the total derivative of S_2 with respect to S_1.

to withhold savings. This incentive to withhold savings translates into an incentive to smoke more for the myopic agent. Therefore, it is unclear whether the myopic agent smokes more or less than the rational agent. It is likewise unclear whether a myopic addict saves more or less than a rational one.

4 Smoking Licenses

Absent any government intervention, the period 1 self undertakes costly activities—namely cutting back consumption—to promote self-control in future periods. This is why she is willing to pay for a licensing scheme. In particular, suppose that the period 1 self can purchase a smoking license, at some fee ϕ. This commits the period 2 self to a cigarette tax τ, but also entitles her to receive a lump-sum transfer of $q = \tau a_2$. Finally, to ensure that the scheme is strictly welfare-improving for the periods 2 and 3 selves—as opposed to just welfare-neutral—we allow a small positive payment ϵ made in period 3 to participants.[8] The period 1 self is willing to give up resources in exchange for a higher relative price of cigarettes in period 2, even though

[8]This scheme is somewhat similar to one envisioned by O'Donoghue and Rabin (2003). To solve the problem of heterogeneous preferences *across* people, they propose that the government charge a one-time fee in exchange for the right to purchase *either* the addictive or non-addictive commodity, alongside a slightly higher one-time fee in exchange for the right to purchase both commodities. Time-consistent consumers would choose the more expansive license, while time-inconsistent consumers, wanting a self-control device, would choose to buy the license for the non-addictive good only. However, such licenses are not generally Pareto self-improving for time-inconsistent agents, except in particular circumstances. From the point of view of later selves, binding restrictions are imposed without any offsetting compensation.

self 2's welfare is unchanged.

In equilibrium $q = \tau a_2$, but q is set in advance of period 2 and is thus taken as given by the period 2 self, who faces the following problem:

$$\begin{aligned} &\max_{v_2,a_2} u_2(I + v_1 - v_2 - p(1+\tau)a_2 + q(\tau), S_1, a_2) + \\ &\quad \beta\delta u_3(I - pa_3(S_2, v_2, \tau) + v_2 + \epsilon, S_2, a_3(S_2, v_2, \tau)) \\ &\quad s.t.\ S_2 = (1-\gamma)S_1 + a_2 \end{aligned} \tag{4.1}$$

Since compensated (Hicksian) demands are always downward-sloping, a compensated cigarette tax will lower smoking, so that $\frac{da_2}{d\tau} < 0$. In addition, a compensated cigarette tax will promote savings. Formally, the first order conditions for this problem are:

$$\begin{aligned} -p(1+\tau)\frac{\partial u_2}{\partial c} + \frac{\partial u_2}{\partial a} + \beta\delta\frac{\partial u_3}{\partial S} &= 0 \\ -\frac{\partial u_2}{\partial c} + \beta\delta\frac{\partial u_3}{\partial c} &= 0 \end{aligned} \tag{4.2}$$

A compensated tax change has no first-order impact on consumption, because $q(\tau) - p(1+\tau)a_2 \equiv -pa_2$, regardless of τ. Using this result simplifies the comparative statics and allows us to derive the Hicksian elasticities:

$$\frac{da}{d\tau} = \frac{p\frac{\partial u_2}{\partial c}(\frac{\partial^2 u_2}{\partial c^2} + \beta\delta\frac{\partial^2 u_2}{\partial c^2})}{D} < 0 \tag{4.3}$$

$$\frac{dv}{d\tau} = \frac{-\frac{\partial u_2}{\partial c}\frac{\partial^2 u_2}{\partial c^2}p^2(1+\tau)}{D} > 0 \tag{4.4}$$

D is the positive Hessian term associated with the maximization problem. Note that we are not considering the effect of the period 3 payment ϵ, because we are assuming it to be sufficiently small.

Even though the compensated tax alters the allocation of resources, it does not affect self 2's utility: the first order impact of the compensated tax on self 2's consumption is zero. The small transfer ϵ thus makes the agent strictly better off from both the self 2 and self 3 perspectives. The period 1 self will also be better off with some nonzero tax, as long as the agent is time-inconsistent. The impact of the compensated tax on this self's utility can be derived by differentiating the following with respect to τ:[9]

[9]In the objective function below, we have substituted the compensating transfer $q = p\tau a_2(S_1, v_1, \tau)$.

$$
\begin{aligned}
&\max_{a_1,v_1} u_1(I - pa_1 - v_1 - \phi, S_0, a_1) + \\
&\quad \beta\delta u_2(I + v_1 - v_2(S_1, v_1, \tau) - pa_2(S_1, v_1, \tau), S_1, a_2(S_1, v_1, \tau)) + \\
&\quad \beta^2\delta u_3(I - pa_3(S_1, v_1, \tau) + v_2(S_1, v_1, \tau) + \epsilon, S_2, a_3(S_1, v_1, \tau)) \qquad (4.5) \\
&s.t.\ S_1 = (1-\gamma)S_0 + a_1 \\
&\quad S_2 = (1-\gamma)S_1 + a_2(S_1, v_1, \tau)
\end{aligned}
$$

From the period 1 self's perspective, the marginal utility of the compensated tax is:

$$
\overbrace{\beta^2\delta(1-\delta)\frac{\partial u_3}{\partial S}\frac{\partial a_2}{\partial \tau}}^{Smoking\ Prevention} + \overbrace{\beta^2\delta(1-\delta)\frac{\partial u_3}{\partial c}\frac{\partial v_2}{\partial \tau}}^{Savings\ Promotion} \geq 0 \qquad (4.6)
$$

The marginal utility is exactly zero for a time-consistent individual whose $\delta = 1$, but strictly positive for a time-inconsistent individual with $\delta < 1$. The compensated tax shifts period 2 resources out of smoking and into savings. Therefore, the marginal utility of licensing has two components. The "Smoking Prevention" effect is the benefit to the period 3 agent of less smoking in period 2. The "Savings Promotion" effect is the benefit of more period 2 savings.

Since the scheme generates surplus for a time-inconsistent individual, the government can sell licenses to such agents. The proceeds can be used to cover the administrative costs of the program, and/or to generate public

revenue for any number of purposes. The license fee willingly paid by a time-inconsistent agent satisfies:

$$\phi < \frac{\overbrace{\beta^2\delta(1-\delta)\frac{\partial u_3}{\partial S}\frac{\partial a_2}{\partial \tau}}^{Smoking\ Prevention} + \overbrace{\beta^2\delta(1-\delta)\frac{\partial u_3}{\partial c}\frac{\partial v_2}{\partial \tau}}^{Savings\ Promotion}}{\frac{\partial u_1}{\partial c}}, \tag{4.7}$$

which is simply the dollar equivalent of the utility derived from the compensated tax. The dollar-value of the compensated tax rises with the degree of time-inconsistency (δ), the harmful future effects of smoking ($\frac{\partial u_3}{\partial S}$), the tax-responsiveness of smoking and savings ($\frac{\partial a_2}{\partial \tau}$ and $\frac{\partial v_2}{\partial \tau}$), and the extent of under-saving, which affects the ratio $\frac{\frac{\partial u_3}{\partial c}}{\frac{\partial u_1}{\partial c}}$.

5 Commitment and Enforcement

An important feature of the licensing scheme we propose is that it would be voluntarily chosen by time-inconsistent agents, but refused by time-consistent ones. In this way, the government could target commitment devices toward those who need them and could avoid inflicting distortionary taxation on others.

The problem with any voluntary commitment device is that it affords several opportunities to cheat. In this section, we discuss potential safeguards. Naturally, many of these problems disappear or become attenuated in a mandatory licensing environment. Voluntary licensing schemes that of-

fer some choice over different licensing schemes, on the other hand, will be subject to potential cheating problems. For the period 2 self, there are incentives to avoid the tax chosen by the period 1 self entirely, while pocketing the transfer. For the period 1 self facing an array of license/tax combinations, there are incentives to choose a tax and transfer scheme such that the tax collected from the period 2 self does not equal the transfer. In both of these situations, cheaters are difficult to observe *ex ante*, but considerably easier to observe *ex post*. Therefore, *ex post* fines against cheaters can encourage compliance. We characterize the size of the necessary fines.

5.1 Malfeasance in Tax Compliance

Even though self 1 may prefer the licensing and taxation scheme, self 2 still has incentives to skip out on the taxes altogether. However, since the government knows the transfer amounts and can potentially keep track of taxes paid, it can easily determine *ex post* who has paid less tax than promised.[10] An individual who has paid less tax than he initially committed to can be fined an amount in period 3. In our three-period model, a fine equal to the total period 2 tax bill exactly offsets the gain from avoiding taxation. Therefore, any negligible probability of additional sanctions will induce compliance.

From self 2's perspective, the marginal utility of evading taxes altogether

[10]This could be either because the individual evaded taxes or because she misreported her future consumption.

is simply $\frac{\partial u_2}{\partial c}(\tau p a_2)$. There are no indirect effects of taxation, because from self 2's perspective, self 3's decisionmaking is optimal; in other words, with just a two-period horizon, there is no time-inconsistency. Similarly, the marginal disutility associated with the fine of $\pi = \tau p a_2$ in period 3 is given by: $\beta\delta\frac{\partial u_3}{\partial c}(\tau p a_2)$. Optimal decisionmaking implies that $\frac{\partial u_2}{\partial c} = \beta\delta\frac{\partial u_3}{\partial c}$. Therefore, from self 2's perspective, the marginal disutility of the lost payment in period 3 is exactly equal to the marginal utility of avoiding all taxes. Consequently, any negligible probability of further sanctions is enough to induce compliance.

The results for this case are clean in part because period 3 is both the end of the smoking license and the end of life. If, instead, the agent chooses a license that expires before the end of life, the results would be complicated by the fact that the fine in period 3 could serve a self-discipline function. In this case, the results would depend on the extent of under-saving (which influences the marginal utility of consumption in period 3 relative to period 2), and the disciplinary value of restricting income in period 3. One could still rely on the useful fact that cheating is observable *ex post*; the difference would be that the necessary fine would not always be exactly equal to the monetary gain from tax evasion.

5.2 Malfeasance in Taxation Choice

When individuals are heterogeneous in their demands, they can also cheat by choosing the "wrong" tax and transfer bundle. In particular, suppose there are light smokers and heavy smokers, where light smokers smoke less in period 2 than heavy smokers: $a_2^l < a_2^h$. For each type, there is an optimal tax and transfer bundle. Define (q_l, τ_l) as the bundle for light smokers and (q_h, τ_h) as the bundle for heavy smokers. We restrict the licensing fee ϕ to be the same across types.[11] Suppose further that, while these two types are known to exist, no single individual can be identified as belonging to one or the other type. Therefore, a light smoker can freely choose the heavy smoker's optimal bundle, and vice-versa.

With time-inconsistent agents, either type of smoker can have incentives to misrepresent herself. A heavy smoker may misrepresent himself to secure a lower transfer to his future self. This reduces income and welfare in period 2, but this may be attractive to the period 1 self since it may restrain future smoking. On the other hand, a light smoker may misrepresent himself in order to secure a larger transfer in period 2 and thus generate a positive net transfer for his future self. The positive income effect might outweigh the disciplinary effects. Misrepresentation can have ill effects either by destroying budget-balance or by allowing the period 1 self to lower the welfare of the period 2 self. In either case, the result can be a policy that is not Pareto

[11] Allowing different licensing fees would create more difficult problems of enforcement, because it may not be possible to discern willingness to pay for the license even *ex post*.

self-improving.

To design an incentive-compatible licensing scheme, we will take advantage both of the welfare-improving aspects of the optimal policy, and of the fact that it is easy to observe, *ex post*, smokers who have chosen the wrong bundle, because they will violate budget-balance. These ideas are similar to those we use to construct an incentive-compatible licensing scheme in the previous section.

We need two safeguards. First, we will allow smokers in period 2 to opt-out of the licensing scheme, as long as they forfeit a tiny "processing fee" from their period 3 pay-out. Period 2 selves would never opt-out of their optimal scheme, but they would opt-out of any scheme in which tax payments exceed the bond posted. Faced with the possibility of period 2 opt-out, the period 1 self will never choose a scheme under which tax payments exceed the bond. On the other hand, to eliminate the choice of a scheme where the bond payment is less than taxes, the government could adopt a policy of fining agents who violate budget-balance in period 2. Under a few conditions discussed below, this penalty is sufficient to force each type to choose its proper bundle.

Imposing a bit more structure on the model, we can show that the threat of a period 3 fine, coupled with allowing opt-out in period 2, are enough to induce the period 1 self to choose the correct bundle. For heavy smokers, cheating means choosing the bundle (q_l, τ_l). Under this bundle, the transfer

does not adequately compensate the period 2 self, because $\tau_l a_h > \tau_l a_l = q_l$. Therefore, even if the period 1 self chooses it, the period 2 self will choose to opt-out and forego a small processing fee, because this has a discrete negative impact on the latter's utility.

For the light smoker, cheating means choosing the bundle (q_h, τ_h), which involves a transfer that exceeds the individual's tax payment. While self 2 would not choose to opt out of this scheme on her own, the threat of a fine could cause her to do so. For self 2, the first-order benefit of remaining in the fraudulent scheme is the extra consumption earned, $\tau_h(a_h - a_l)$. Since self 2 is time-consistent, it is straightforward to show that a fine of $\pi = \tau_h(a_h - a_l)$ is exactly enough to offset the value of cheating.[12] Any negligible probability of further sanctions (like imprisonment, for example) would be enough to encourage self 2 to opt out rather than face the fine. Knowing this, self 1 would never willingly choose the incorrect bundle.

The analysis above was aided by the time-consistency of self 2, which is an artifact of a three-period problem. We conjecture that extending the time horizon will not fundamentally change the mechanisms that guarantee truthful revelation of type. With three periods, Pareto self-improving licensing schemes all have the same structure: a license fee paid by self 1, a tax on the addictive good imposed on self 2, and a fixed bond returned to self 2 that makes her indifferent to the tax. The main complication of a longer time

[12]This amount plus interest could be imposed on self 3 with the same effect.

horizon is that it increases the dimensionality of possible licensing schemes. Self 1 would still pay a single licensing fee, but there would be a tax and corresponding bond for each self up to the second to last self. Though formalization is more complicated, the nature of the licensing scheme would still be the same. In exchange for self 1 purchasing the licensing fee, the tax receipts imposed on each self after self 1 would exactly equal the value of the bond paid by the government.

Assume still that there are still only two types of smokers—heavy and light. We conjecture that a mechanism like the one we outline above, except extended to multiple periods, would prevent self 1 from misrepresenting herself. Each self beyond the first would be provided with an opt-out incentive from the license, and the government would have to check each period that the bond did not exceed taxes collected. In a sense, extending the time horizon makes the government enforcement job easier, since it provides for more punishment opportunities. In addition to satisfying the usual condition that the monetary value of the fine has to be equal to the monetary value of cheating, it must also be true that cheating overall does not act as a self-disciplinary mechanism. If it does, the monetary value of the fine has to exceed the monetary value of cheating. As in the three-period case, it seems likely to us that mechanisms like this would lead in equilibrium to truthful revelation of type by self 1, though a formalization of this argument is beyond the scope of this paper.

6 Practical Considerations

We have presented a view of the licensing scheme in a three-period model, but the length of each period could vary arbitrarily. Its key feature is the presence of an initial period with a licensing fee but no tax, a second period with taxes and a lump-sum transfer, and a final period where the original licensing fee is recovered as a lump-sum. Moreover, the assumption that individuals buy the license in the initial period is also not necessary. People could choose to opt in at any time. This may be especially useful if agents realize over time that they have a self-control problem.

Many state and local governments use cigarette taxation as a source of revenue in itself. While the scheme we considered above is designed to be revenue-neutral, rents could in principle be extracted from time-inconsistent smokers, since they have a positive willingness to pay for a license. Another possibility is the introduction of licenses as a voluntary supplement to existing cigarette taxation laws. Sophisticated time-inconsistent smokers would willingly choose to purchase a license that binds them to a higher future price of cigarettes. Time-consistent or naif time-inconsistent smokers are indifferent between purchasing a license or not. Small costs like a processing fee for those who wish to be taxed would thus discourage them from obtaining one. As such, a voluntary licensing scheme has the further advantage of differentiating between time-inconsistent smokers with a demand for self-control, and others without such a demand. Blanket taxation is not able to make this

distinction.

One of the most difficult parts of implementing any optimal tax system—and this is no exception—is calculating the optimal tax schedule. The usual approach (cf., Gruber and Koszegi, 2001) is to calibrate utility functions and calculate optimal tax schedules. A similar approach could be taken here, although it is clear that such calibrations are imprecise at best. Unlike calibration or estimation exercises, taxes never have standard errors. In the case of smoking licenses though, the standard calibration approach could be complemented by a revealed preference approach that could be used to refine the optimal fee schedule over time. If the government offers a schedule of taxes and fees and finds that few smokers are taking up licenses, it can (assuming that smokers are time-inconsistent) infer that the license fees are set too high relative to the tax rates. To produce information about fees that are set too low, it could offer smokers the chance to buy out of their licenses in the second period, for some small but nontrivial amount of money. If fees and taxes are set optimally, no smoker would willingly give up resources to opt out of the scheme. Therefore, if many smokers are opting out, it provides evidence that fees are set too low relative to taxes. A voluntary program would also have other advantages. If some agents are time-consistent, they are better off without a licensing scheme, or regulation of any kind.

While the exact nature of the enforcement is probably best determined elsewhere, it is worth discussing some implementation issues. License and tax

information could be embedded—visually or electronically—in state drivers' licenses and identification cards. Merchants would then be required to request identification from everyone, rather than just suspected under-age smokers. Even though the scheme is Pareto-improving at every point in time, enforcement and monitoring will be necessary, as the preceding section discussed.

Finally, it is worth returning to the problem of naif smokers. Licensing does not (prospectively) Pareto-improve their utility, and they would also refuse to opt into a voluntary scheme. This is one argument in favor of mandatory taxes rather than voluntary licenses. However, even mandatory taxes are blunt instruments for improving the welfare of naifs, because they do not Pareto-improve the utility of any smoker from a lifetime perspective. Another approach is to regulate directly the smoking of groups we believe to be naifs, such as juveniles. A licensing policy is not mutually exclusive with these kinds of direct regulations. It is important to emphasize, however, that the justifications for these policies are by nature paternalistic rather than Paretian, because society is determining that individual preferences are inappropriate and misinformed.

7 Conclusions

Time-inconsistency destabilizes the conventional view that government regulation cannot improve individual welfare when markets are competitive.

However, conventional solutions like taxes often fail to meet economists' usual Pareto criterion. A tax that an individual finds optimal early in life may be welfare-reducing later in life, when the demand for self-control is lower. One way to solve the problem is to collect the costs imposed by taxation and shift them entirely onto the individuals with the highest willingness to pay, namely the youngest individuals.

The analytics of our discussion were focused on addiction and smoking, but similar schemes can be used to deal with time-inconsistency in other contexts, like savings. For example, an individual could pay for a future "savings subsidy" that would be coupled with a lump-sum surcharge on an individual's income taxes. In practice, the government could allow tax-deductible contributions to a retirement account of $X, and levy an exactly offsetting surcharge on an individual's income taxes. This raises the relative return to savings in the future, but maintains budget-balance and welfare-neutrality for the future self. Moreover, the current period self would be willing to pay for the savings subsidy and thus generate additional revenue, either for the administrative costs of the program or for general state revenues. The general point is that government policies to improve the welfare of time-inconsistent individuals ought to shift their costs onto the early selves with the highest willingness to pay. An early payment coupled with future compensated taxes is one way to accomplish this objective, and to do so in a manner consistent with the Pareto criterion.

While we have focused on individual welfare, our results could be interpreted primarily in terms of public finance as well. Time-inconsistent agents represent a type of arbitrage opportunity that governments can exploit. Our analysis suggests the possibility of welfare-neutral revenue-raising schemes in the presence of time-inconsistent agents. Put bluntly, time-inconsistent agents can literally "pump money" into an economy, because they are willing to pay for the creation of a market that permits enforceable trades across time.

Future research is needed to calculate the appropriate tax and fee schedule for a licensing scheme. We discussed above the possibility of taking a revealed preference approach to calculating this schedule. Indeed, the policy implementation for smoking could be preceded by a small-scale social experiment designed to elicit the appropriate schedule from a representative population of smokers. One of the advantages of licensing over taxes alone is that it permits a revealed preference approach, precisely because it makes smokers better off at all points during the life-cycle, rather than simply at those points where the willingness to pay for self-control is at its peak.

References

Caplin, A., and J. Leahy (2000). "The Social Discount Rate." *Working Paper 7983*. Cambridge, MA: National Bureau of Economic Research.

Cropper, M., and D. Laibson (1998). "The Implications of Hyperbolic Discounting for Project Evaluation." *Working Paper 1943*. World Bank.

Goldman, S. M. (1979). "Intertemporally Inconsistent Preferences and the Rate of Consumption." *Econometrica* 47(3):621–626.

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

Gruber, J., and B. Koszegi (2004). "Tax Incidence When Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes." *Journal of Public Economics* Forthcoming.

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

Laibson, D. I. (1994). *Hyperbolic Discounting and Consumption.* Ph.D. thesis, Massachusetts Institute of Technology.

Lowenstein, G., and D. Prelec (1992). "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics* 57:573–598.

O'Donoghue, T., and M. Rabin (1999). "Doing It Now or Later?" *American Economic Review* 89(1):103–124.

O'Donoghue, T., and M. Rabin (2003). "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes." *American Economic Review* 93(2):186–191.

Phelps, E. S., and R. A. Pollak (1968). "On Second-Best National Saving and Game-Equilibrium Growth." *Review of Economic Studies* 35:185–199.

Strotz, R. H. (1955). "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23:165–180.