SWEDISH ECONOMIC POLICY REVIEW 12 (2005) 7-39

Optimal taxes for sin goods

Ted O'Donoghue and Matthew Rabin*

Summary

■ Sin goods, such as potato chips, generate immediate consumption benefits but future health costs. Because self-control problems—in the form of a preference for immediate gratification—can lead people to consume more of such goods than they themselves prefer, there might be scope for regulation to counteract such over-consumption. We investigate one particular form of regulation: taxes. Imposing simple per-unit taxes (and returning the proceeds to consumers) can generally improve social surplus. Moreover, in some instances, it can be better to use more sophisticated schemes, such as sin licenses, that induce people to make prospective choices about future consumption. ■

JEL classification: H21, D11, D91.
Key words: Hyperbolic discounting, present bias, time-inconsistent preferences, optimal taxation.

* *Ted O'Donoghue is Associate Professor in the Department of Economics at Cornell University. Matthew Rabin is Professor in the Department of Economics at University of California, Berkeley.*

SWEDISH ECONOMIC POLICY REVIEW 12 (2005) 7-39

Optimal taxes for sin goods

Ted O'Donoghue and Matthew Rabin*

As obesity rates have risen in recent years and health worries have ensued, there has been more and more talk of regulating unhealthy consumption. The question arises whether such regulation is called for. If people want to consume large quantities of potato chips and gain weight as a result, who are we to say they shouldn't?

The standard economic approach to such regulation provides one possible answer: If part of the burden of obesity is borne by society and not just the individual, then there is a standard externality justification for regulation. We believe, however, that a large part of the worry about obesity among both the populace and policymakers is not driven by such externalities, but rather by a concern that the people themselves may not be happy with their own consumption of unhealthy goods. Because the standard economic approach *a priori* assumes that people always behave in their best interests, it provides no scope to address such "over-consumption". But recent research in behavioral economics suggests that people might consume more than they themselves prefer, and, if so, regulation might be justified as a means to counteract over-consumption of unhealthy goods.

In this paper, we investigate one particular form of regulation—taxes on the purchase of unhealthy goods—and how it might combat over-consumption due to self-control problems. We first analyze the potential benefits of a simple per-unit tax, and then investigate more sophisticated schemes such as selling licenses or vouchers that permit (or require) people to make choices about the prices they will face in the future.[1]

* *We thank Per Krusell and other conference participants for useful comments. This paper draws from ideas in two related papers (O'Donoghue and Rabin 2003, 2005).*

[1] Two other recent papers also study the welfare effects of sin taxes. Using survey data from Canada, Gruber and Mullainathan (2005) provide empirical evidence that higher local cigarette taxes lead to increased happiness. Gruber and Koszegi (2004) also study cigarette taxation in the presence of self-control problems, where they calibrate tax incidence for different income groups. There is also an older literature that studies "merit goods" for which the government has a different notion of indi-

In Section 1, we briefly discuss some general issues that must be addressed when pursuing "paternalistic" regulation. In Section 2, we introduce the specific error that we will study: over-consumption due to self-control problems—more precisely, due to people having a preference for immediate gratification that they disagree with at every other moment of their lives. If eating potato chips generates immediate consumption value and future health costs, a person should consume potato chips whenever the consumption value is sufficiently large relative to the health consequences. Consumers with zero self-control problems—as typically assumed in economics—weigh these costs and benefits appropriately, and so they consume optimally. From a prior perspective, people with non-zero self-control problems also weigh these costs and benefits appropriately and so they want—and might even naively expect—to consume optimally. In the moment, however, their preference for immediate gratification kicks in, leading them to over-weight the immediate consumption value relative to the future health costs, and hence they might over-consume potato chips.

In Section 3, we investigate a simple regulation designed to combat such over-consumption: impose a per-unit tax on potato chips, and return the tax proceeds to consumers—perhaps by lowering income taxes or taxes on some non-sin commodity, such as socks. If consumers are homogeneous, or if we can observe each consumer's preferences and impose an individual-specific tax, then such a per-unit tax can implement optimal behavior for everyone. Our primary interest, however, is the more realistic case in which there is unobservable population heterogeneity—both in the degree of self-control problems and in tastes for potato chips. In this case, a simple per-unit tax cannot implement optimal behavior for everyone. We argue, however, that a per-unit tax can be quite effective at increasing social surplus relative to no tax, because the tax creates large (first-order) benefits for people with self-control problems, while at the same time it creates small (second-order) costs for fully self-controlled individuals. We show, moreover, that such taxes may sometimes in fact help *everybody* relative to no taxes. Intuitively, because self-control problems lead to increased potato-chip consumption, people with self-control problems consume more potato chips than fully self-controlled peo-

viduals' optimal consumption than the individuals themselves have (see Musgrave, 1959; and Besley, 1988).

ple. If the proceeds from potato-chip taxes are returned evenly to everyone in the population, the result will be that income is naturally redistributed from the people with self-control problems to fully self-controlled individuals, and so fully self-controlled individuals may in fact benefit from potato-chip taxes as well.

In Section 4, we investigate more sophisticated schemes. Simple per-unit taxes merely alter in-the-moment incentives. But since people with self-control problems would like to behave optimally in the future, it might be useful to induce *prospective* choices wherein people make choices now that influence their future in-the-moment incentives. We investigate sin licenses as a means to induce such choices. Specifically, we consider policies where there is a presumptive per-unit tax, but people can buy a license that exempts them from the tax in the future. We demonstrate that such licenses can in principle be quite useful—indeed, for the special case where everyone knows exactly their future consumption values, we can use a 1 cent license to implement optimal future behavior for everyone, because only those who should be consuming see any value in the license. We also discuss how such licenses might work more generally.

Finally, in Section 5 we conclude by discussing some broader issues with respect to our proposed policies.

1. Optimal paternalism

Recent research in behavioral economics suggests a variety of errors that can lead people not to behave in their own best interests. It seems natural for public-policy analysis to take such errors into account—to design policy with an eye toward how that policy might help people to make better choices. Before we engage in such "paternalistic" regulation, however, we presumably want to be confident that this regulation will do more good than harm.

Two major worries about paternalism are often put forth. First, in trying to help people who make errors, how can we be sure that we are not causing significant harm to people who do not make errors? There are surely many people who are choosing optimally (or nearly optimally), and we want to be wary of limiting their ability to do so. Second, given the private information that heterogeneous individuals have about their own preferences, how can we know what people ought to be doing? Most individuals know better than we do what's in

their best interests, and so we certainly shouldn't tell people what to do.

These concerns have been much discussed in the behavioral-economics literature, and, in response, initial research on paternalism has taken a very cautious approach: Identify unobtrusive ways to combat errors while (arguably) imposing little or no harm on fully rational people and minimal implementation costs. This basic approach has been put forth by various authors under various labels: "cautious paternalism" (O'Donoghue and Rabin, 1999a), "asymmetric paternalism" (Camerer et al., 2003), "libertarian paternalism" (Sunstein and Thaler, 2003a, b), and "benign paternalism" (Choi et al., 2003). Some examples of cautiously paternalistic policy are rules for easy-to-change default outcomes and rules for the way information is framed.

While this cautious approach is a useful place to start, additional insights can be gleaned from more rigorous investigations. In O'Donoghue and Rabin (2003), we emphasize the importance of studying "optimal paternalism", by which we mean formally analyzing optimal policy as a function of our beliefs about the degree of and prevalence of errors in the population. In other words, write down assumptions about the distribution of rational and irrational types of agents, about the available policy instruments, and about the government's information about agents, and then investigate which policies achieve the "best" outcomes. Our analysis in this paper represents a simplified version of this approach.

A careful investigation of optimal paternalism will impose discipline on both sides of the debate. On one hand, it reveals that paternalism is not synonymous with restrictions on choices or telling people what to do, as is sometimes claimed. Given reasonable informational limits on policymakers, it will virtually never be optimal to tell people what to do, and rarely optimal to merely ban certain types of consumption. In fact, the best policies might not even take the form of restricting choice sets at all. For instance, the simple taxes that we discuss in Section 3 do not reduce choice sets. Rather, they merely change relative prices so as to alter incentives—eliminating some options, but also making new options available. Moreover, the more sophisticated schemes that we discuss in Section 4 involve the *expansion* of choice sets, illustrating how in some instances the best way to help consumers make better choices is to make new options available.

On the other side of the debate, a careful investigation of optimal paternalism will reveal that paternalism is not a recipe for "anything

goes" wherein people propose paternalistic policies that merely reflect their own personal prejudices. Careful thought will force people to assess policies in terms of exactly what circumstances they will be useful vs. not useful. Moreover, careful thought will also force proponents of paternalism to recognize a third concern about paternalism: regulation designed to combat one error could end up exacerbating another error. Even if we are correct that some people are not behaving in their own best interests, if we are incorrect about the source of this misbehavior, we may end up exacerbating the problem. We shall return to this issue in Section 4.

2. Over-consumption of sin goods

Consider a person's decision whether to consume a sin good such as potato chips. The essential feature of such goods is that consumption generates immediate enjoyment but future health costs or other future negative consequences. The person's optimal consumption trades off the consumption value vs. the health costs; but due to self-control problems, the person may weigh this trade-off incorrectly and end up over-consuming potato chips.

To fix ideas, consider a simple model of potato-chip consumption. A person decides each week whether to purchase a bag of potato chips at a price p. The person enjoys eating potato chips: the immediate pleasure from eating a bag of potato chips is $v > 0$. But the person also recognizes that eating potato chips generates a future cost of $c > 0$. Hence, the person should optimally consume potato chips when $v > c + p$.

Evidence suggests, however, that people often do not make optimal trade offs between immediate and future payoffs. Rather, people tend to pursue immediate gratification in a way that they themselves disapprove of from a long-run perspective. From a long-run perspective, people make even-handed trade offs between payoffs in different periods. But in the moment, they tend to give too little weight to future payoffs relative to immediate payoffs.[2] In the present context,

[2] For evidence, see Ainslie (1991, 1992), Ainslie and Haslam (1992a, b), Loewenstein and Prelec (1992), Thaler (1991), and Thaler and Loewenstein (1992). For a recent overview, see Frederick et al. (2002). This tendency is often referred to as "hyperbolic discounting". In recent years, economists have explored the implications of such preferences in various environments: savings decisions (Laibson 1997, 1998; Angeletos et al., 2001), procrastination (O'Donoghue and Rabin, 1999b,

such a preference for immediate gratification creates a propensity to over-consume potato chips. From a prior perspective, the person would like to trade off the immediate eating pleasures against the future health consequences in an even-handed manner. But when the time comes to make her decision, the person gives too little weight to the future health consequences, and hence is over-prone to buy and eat potato chips.

We can incorporate this idea into our model by assuming that there is a present-bias factor $\beta \leq 1$ that reflects the extent to which the person underweights future costs.[3] Hence, it is optimal to purchase potato chips—and from a prior perspective the person would like to purchase potato chips—if $v > c + p$. In the moment, however, the person actually purchases potato chips if $v > \beta(c + p)$. Because $\beta < 1$ implies $c + p > \beta(c + p)$, the person might choose to buy potato chips when in fact it would be optimal not to buy. In particular, if $\beta(c + p) < v < c + p$, the person buys when she shouldn't, as illustrated in Figure 1.[4]

Figure 1

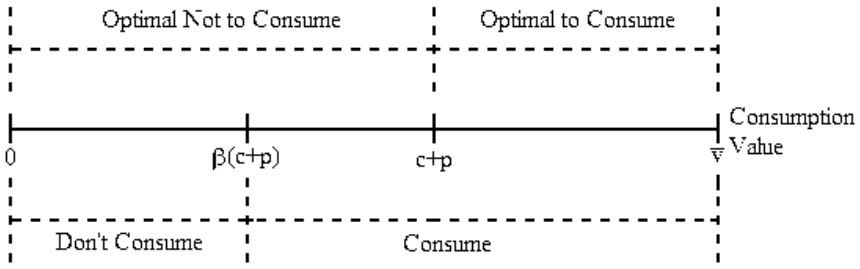

2001; Fischer, 1999); and information acquisition (Carrillo and Mariotti, 2000; Benabou, and Tirole 2002).

[3] This formulation represents a simplified version of the usual β,δ model that has been used widely since Laibson (1997) to model a time-inconsistent preference for immediate gratification. In principle, some people might have a future bias—that is, $\beta > 1$—in which case they would consume too few potato chips relative to what they would like. But since the evidence seems to suggest that the more prevalent error is a present bias, we shall assume for simplicity that everyone has $\beta \leq 1$.

[4] Implicit in this formulation is that the person treats the price paid as a future cost. While we feel that this is an appropriate assumption—because the price reflects forgone future consumption of other goods—our conclusions do not depend on this assumption.

In principle, the behavior of people with such preferences depends on whether they are aware vs. unaware of their future preference for immediate gratification—on whether they are "sophisticated" or "naïve".[5] Suppose for instance that a person is given the ability to commit now, for a small price, to next week's potato-chip consumption. A *sophisticate* who is aware of her future preference for immediate gratification might find such a commitment quite valuable. In our model, if a sophisticate has $\beta(c + p) < v < c + p$, she'd be willing to pay for a commitment to no potato-chip consumption next week. In contrast, a *naïf* who is unaware of her future preference for immediate gratification believes that she will behave optimally next week and hence is not willing to pay anything for such a commitment.

This discussion reveals two key intuitions. First, the distinction between sophistication and naiveté matters only when people are making prospective decisions that alter their future consumption outcomes. Hence, this distinction will be irrelevant in our discussion in Section 3 of simple linear taxes.[6] It becomes more relevant for our discussion in Section 4 of sin licenses, because these schemes revolve around getting people to make choices now that alter future incentives. This brings us to the second key intuition: when people with a preference for immediate gratification are asked to make prospective decisions, they focus on their long-run preferences. Hence, the prospective aspect of decisions to buy licenses might prove useful in improving outcomes, which will be the focus of Section 4.

Another important feature of many sin goods is that they are repeatedly purchased in small quantities for short-term consumption. People do not make a single trip to the store to purchase a lifetime supply of potato chips. Rather, they make numerous purchases over the course of their lives, each time buying only for the current day, week, or perhaps month. Of course, this distinction between making infrequent long-term consumption decisions vs. making frequent short-term consumption decisions is (mostly) irrelevant for fully self-controlled people. For people with self-control problems, however, this distinction can be crucial. When making a long-term consump-

[5] This distinction—and the labels—was first identified by Strotz (1956) and Pollak (1968), and it has been investigated recently by O'Donoghue and Rabin (1999b, 2001).

[6] If there were habit formation, then the distinction can become relevant again because even current consumption can alter future consumption outcomes. We discuss habit formation more in Section 5.

tion decision, such as how many potato chips to consume over the next year, a preference for immediate gratification creates relatively little distortion—you might choose to over-consume today, but you won't choose now to over-consume on days months into the future. But when people make a series of short-term consumption decisions, such as choosing each day how many potato chips to consume today, even a small preference for immediate gratification can create large distortions—because you might choose every day to over-consume on that day.[7] We will return to this issue in Section 4, where we discuss how sin licenses might be a useful way to induce people to make infrequent long-term consumption decisions for sin goods.

While self-control problems represent a natural source of over-consumption, other errors might influence potato-chip consumption as well.[8] One possibility is that people might have irrationally optimistic beliefs about the health costs associated with eating potato chips. Indeed, we could reinterpret our model in these terms, where now β reflects the degree of over-optimism about health costs. There is one important difference, however, which is that people with irrational beliefs would feel the same way when making in-the-moment decisions and prospective decisions, and so the prospective policies that we suggest in Section 4 may not be helpful for such types.

A more problematic error is that, if people's consumption values fluctuate, people might have a propensity to falsely project their current preferences onto the future. Specifically, on days when they have high valuations, people might expect future consumption values to be larger than they actually will be; and on days when they have low valuations, people might expect future consumption values to be smaller than they actually will be.[9] For decisions that concern only current consumption, these biased predictions for future tastes are

[7] To the extent that people are sophisticated, it may be possible to mitigate this problem by limiting short-term access to the good. But since for many sin goods it is hard to limit one's access—most people are always within a few minutes of a store that sells potato chips—sophisticates may be no better situated than naifs.

[8] We also emphasize that, more generally, a preference for immediate gratification is not the source of all over-consumption. For instance, while a preference for immediate gratification would not generate over-consumption of socks, a person might still over-consume socks due to a persistent irrational (incorrect) belief that she needs socks, or perhaps even due to an excessive visceral enjoyment from the act of buying socks.

[9] This error is an example of the more general concept of "projection bias", as modeled in Loewenstein et al. (2003).

irrelevant. These biased predictions could, however, lead to bad prospective decisions, such as whether to purchase licenses. We will address this issue in more detail after we discuss licenses in Section 4.

Our analysis shall assume no externalities of the traditional sort wherein a person's consumption of potato chips creates direct harm to others. Such externalities are likely quite important. In particular, the future health costs generated by potato-chip consumption may end up being borne by society and not the individual. If so, then even if the individual behaves optimally from her own perspective, she may end up over-consuming potato chips from a societal perspective. While a more thorough analysis would need to incorporate such effects, we abstract away from such externalities to better highlight that sin taxes can be useful even in their absence.

Finally, our analysis shall assume that, in the market, potato chips are priced at cost, and so we abstract away from distortions driven by monopoly pricing or by cost subsidies to specific industries. Such distortions might be important. Indeed, for sin goods such as potato chips, monopoly pricing might improve social surplus relative to competition, although only with a redistribution from potato-chip eaters to firm owners. Perhaps more important, subsidies for particular goods might lead to many sin goods being priced below cost—e.g., corn subsidies in the United States have certainly reduced the price of corn-syrupy goods such as soda. When combined with self-control problems, such subsidies might lead to vast over-consumption from a social perspective. Once again, a more thorough analysis would need to incorporate such effects.

3. A simple solution: Taxes

In the previous section, we described how self-control problems can lead people to over-consume sin goods such as potato chips. Although this over-consumption is not driven by standard externalities, it can be given a modified externality intuition. Because potato chips generate future health consequences, current consumption imposes a negative "externality" on one's future self.[10] While conventional economic analysis assumes that people fully internalize this externality, people with self-control problems only do so partially, and hence they over-consume. This externality intuition suggests a simple solution:

[10] Herrnstein et al. (1993) labeled such within-person externalities "negative internalities".

impose a per-unit tax on potato chips to induce people to consume less, and return the proceeds to consumers via a lump-sum transfer—or perhaps by lowering income taxes or taxes on some non-sin commodity, such as socks. In this section, we investigate this solution in more detail.

By imposing a per-unit tax on potato chips, we increase the effective market price of potato chips, and hence reduce consumers' demand for potato chips. In our model, if we impose a per-unit tax t on potato chips, then the consumer must pay an effective price of $p + t$. Hence, she will purchase potato chips when $v > \beta(c + p + t)$, and so the larger the tax the less likely she is to buy.

If we know enough about a consumer, then we can implement that person's first-best outcome. Specifically, we counteract the person's self-control problem by setting the tax so that the net costs from potato chips are correct. In our model, the person should consume when $v > c + p$. Given tax t and self-control problem β, she actually consumes when $v > \beta(c + p + t)$. If we set the tax $t = ((1 - \beta)/\beta)(c + p)$, the actual consumption decision exactly corresponds to the optimal consumption decision.[11] Hence, if there is a homogeneous population, or if we can observe each consumers type and impose an individual-specific tax, then we can implement first-best optimal behavior for everyone.

But clearly there is unobservable heterogeneity in the population. People vary in how much they enjoy potato chips (there is population heterogeneity in v); and people vary in their propensity for immediate gratification (there is population heterogeneity in β).[12] In this case, we cannot implement optimal behavior for everyone, because doing so requires individual-specific taxes that depend on unobservable, individual-specific characteristics. What is the best we can do with uniform taxes and lump-sum transfers?

The answer to this question depends on our definition of "best". As a first cut, we investigate optimal taxes given a social-welfare function that puts equal weight on all people. For such a social-welfare function, the monetary redistributions associated with taxes and lump-sum transfers have no direct effect on social welfare. Income is

[11] Note that $\beta = 1$ implies no tax is necessary, which reflects that consumers with zero self-control problems consume optimally in the absence of a tax.

[12] In principle, there is also heterogeneity in the health cost, but for simplicity we shall ignore it.

redistributed from potato-chip consumers to non-consumers, but since our social-welfare function weights everyone equally, these redistributions cancel out. Hence, the only concern in terms of setting optimal taxes is to minimize the distortions in potato-chip consumption. In our model, distortions occur when a person either buys despite having $v < c + p$ or does not buy despite having $v > c + p$.

To build intuition, consider the implications of imposing a small potato-chip tax in our model—specifically, let's compare $t = 0$ to $t = \varepsilon$. Figure 2 illustrates the implications of this tax for fully self-controlled people and for people with a specific $\beta < 1$.

Figure 2

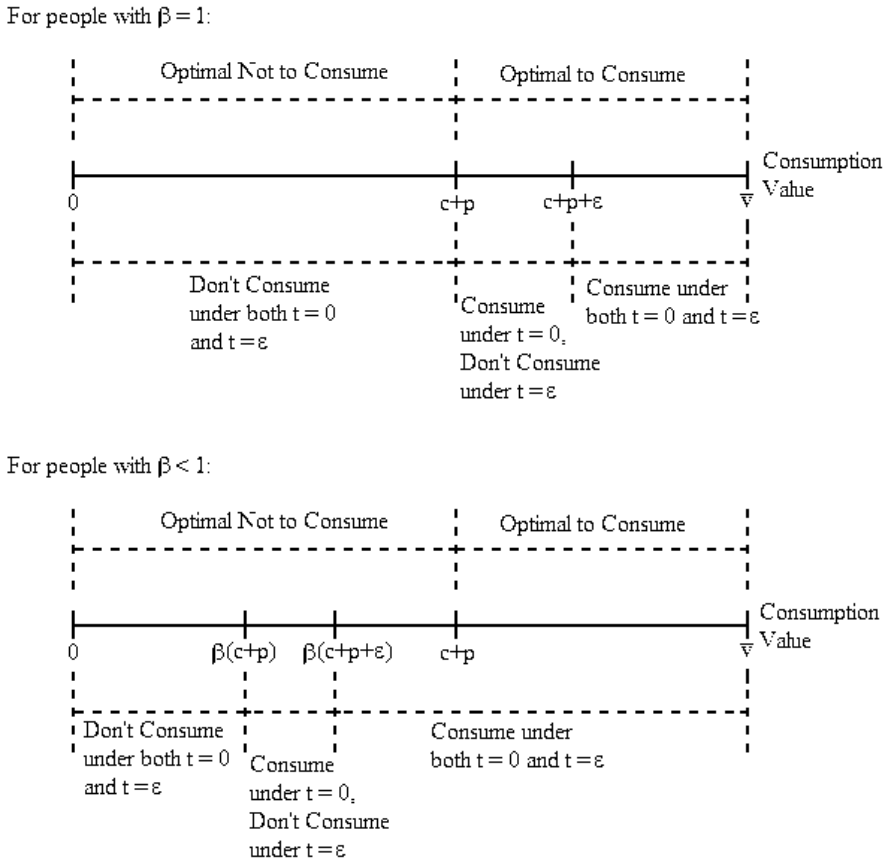

Amongst fully self-controlled people, the only people whose behavior is influenced by the tax are those with a consumption value

between $c + p$ and $c + p + \varepsilon$. For these types, the tax leads to increased consumption distortions, because they optimally consume without the tax while they suboptimally do not consume with the tax. But the only reason that the tax changes their behavior is that their net benefits from potato chips are small—indeed, the maximum lost utility occurs for those with $v = c + p + \varepsilon$, and the utility loss is merely ε. This result reflects that, while potato-chip taxes create consumption distortions for fully rational agents, these consumption distortions will be small (second-order) when the tax is not large.[13]

Amongst people with self-control problems $\beta < 1$, the people whose behavior is influenced by the tax are those with a consumption value between $\beta(c + p)$ and $\beta(c + p + \varepsilon)$. For these types, the tax leads to decreased consumption distortions, because they suboptimally consume without the tax while they optimally do not consume with the tax. Moreover, the people affected are those with the largest initial distortion under no tax—the utility gain for those with $v = \beta(c + p)$ is $(1 - \beta)(c + p)$. This result reflects that potato-chip taxes reduce consumption distortions for people with self-control problems, and these improvements can be significant (first-order).

This intuition that a potato-chip tax creates large (first-order) benefits for people with self-control problems and small (second-order) costs for fully self-controlled individuals is quite robust, and so a per-unit tax on sin goods can be quite effective at increasing social surplus relative to no tax. Indeed, in our model, one can show that for any distribution of consumption values and self-control problems, if we weight everyone equally, it is optimal to impose a per-unit tax on potato chips (see Appendix Box 1).

While the intuition above argues for positive taxes on sin goods, it is not clear whether these taxes would be large enough to be of any economic significance. It's beyond the scope of our analysis to provide a serious empirical calibration, but some simple back-of-the-envelope numbers from our model illustrate that there should be no

[13] If we are completely wrong, and people make no errors, how much harm will we have done? The intuition in the text suggests that the answer may be not much. In other words, the same rational-choice theory that says it is optimal not to impose distortionary taxes also says that the harm from small taxes is negligible. It's an open empirical question, but we conjecture that the harm from even seemingly large potato-chip taxes—of even 20-30 percent—would be small to fully rational people.

presumption that optimal taxes will be small. For instance, if the consumption values are distributed uniformly, and if proportion ϕ of the population has $\beta = \underline{\beta}$ while proportion $1 - \phi$ has $\beta = 1$, one can derive that the optimal tax is

$$t^* = \frac{\phi\underline{\beta}(1-\underline{\beta})}{\phi\underline{\beta}^2 + (1-\phi)}(c+p)$$

(again see Appendix Box 1). Table 1 uses this formula to provide a few numerical examples of the optimal tax.[14]

Table 1. Optimal tax as a percentage of the price

Future Health Cost	Proportion of Population with $\beta < 1$	Their β	Optimal Tax (%)
2p	1.0	0.99	3.0
2p	0.5	0.99	1.5
2p	1.0	0.95	15.8
2p	0.5	0.95	7.5
2p	0.5	0.90	14.9
2p	0.2	0.90	5.6
3p	1.0	0.99	4.0
3p	0.5	0.99	2.0
3p	1.0	0.95	21.1
3p	0.5	0.95	10.0
3p	0.5	0.90	19.9
3p	0.2	0.90	7.5

Thus far, our discussion of sin taxes has been framed in terms of how we might help people who make errors to the detriment of those who do not. While this trade-off holds in terms of distortions in potato-chip consumption, we have not yet incorporated the taxes and lump-sum transfers into our analysis. As we have discussed, any income redistributions cancel out when we average the entire population. If, however, we look at specific sub-groups of the population, there might be an income redistribution from one group to another. In particular, in order to assess the extent to which we are helping people who make errors to the detriment of those who do not, we

[14] Table 1 reflects that the optimal tax depends on the magnitude of the future health costs relative to the purchase price, and in particular the larger are the health costs, the larger is the optimal tax.

can group people by their degree of self-control problem—that is, by their β's. We can then investigate by how much people with $\beta < 1$ are on average helped, and by how much people with $\beta = 1$ are on average hurt.

When one does this, it turns out that a per-unit tax on sin goods may not even on average hurt the fully self-controlled individuals—it could be that the tax creates a *quasi-Pareto improvement* in the sense that people with self-control problems and fully self-controlled people are both on average better off with the tax. Intuitively, because people with self-control problems are on average more prone to consume potato chips, the tax and lump-sum transfers naturally redistribute income from people with self-control problems to fully self-controlled people. Hence, people with self-control problems are on average better off due to reduced consumption distortions, while at the same time fully self-controlled people are on average better off because they receive a net income transfer.[15]

Of course, such quasi-Pareto improvements are not always possible. Because a potato-chip tax creates only small (second-order) distortions for fully self-controlled people while the positive net income transfer is large (first-order), the ability to make fully self-controlled people on average better off is quite general. It is the people with self-control problems who might be on average worse off, because the reductions in consumption distortions may be outweighed by the negative net income transfer. Which effect dominates depends on the marginal effectiveness of the tax at reducing suboptimal consumption, which in turn depends on the magnitude of self-control problems. For instance, in our model, when consumption values are distributed uniformly and proportion ϕ of the population has $\beta = \underline{\beta}$ while proportion $1 - \phi$ has $\beta = 1$, one can show that a positive tax can create a quasi-Pareto improvement as long as $\underline{\beta} > 1 - \phi$ (see Appendix Box 2).

Our discussion above interprets a quasi-Pareto improvement in terms of group averages: Among people with a particular β, some are better off, some are worse off, but on average they're better off. If, however, we reinterpret the population heterogeneity in consumption values as identical people facing fluctuations across time—e.g., each

[15] In O'Donoghue and Rabin (2005), we demonstrate in a model of continuous consumption that it can even be the case that everyone is better off with a sin tax, so the tax and lump-sum transfers create a full Pareto improvement.

day you draw a consumption value from some distribution—then quasi-Pareto improvements in fact represent full Pareto improvements in the sense that each individual is better off from an ex ante expected utility perspective.

4. A more sophisticated solution: Sin licenses

Our discussion in Section 3 assumes a kind of static, "one-shot" world in which people make isolated decisions about whether to purchase potato chips for current consumption, and in which the only available policy instruments are contemporaneous taxes and lump-sum transfers. In that context, we can do no better than the simple linear taxes that we studied in Section 3. More generally, however, we might do better by taking advantage of the intertemporal nature of these decisions. In particular, as we discussed in Section 2, it might be useful to induce people with self-control problems to make *prospective choices* wherein people make choices now that influence their future in-the-moment incentives. In this section, we explore schemes designed to accomplish this goal.

Our general theme will be the usefulness of giving people the option now to choose whether to be subject to high vs. low prices in the future. While there are multiple ways to implement such a choice, we focus on one natural way to do so: impose a high presumptive tax, and then sell licenses (or vouchers) that permit people to buy the good tax-free (or at a reduced tax) in the future. Of course, whether people with self-control problems buy such licenses will depend on whether they are sophisticated vs. naïve about future self-control problems. Hence, we will need to assess the different responses for these two types.

We begin with a very stark case. Suppose that people only consume in the future, and they know exactly what their future tastes will be—that is, they know exactly what their consumption values will be, and they know exactly what their health costs will be. In this case, we can implement the first-best outcome with a very simple sin license: One can buy a license for 1 cent that permits you to buy potato chips tax-free in the future, where you can still buy potato chips without this license, but doing so requires paying some presumptive tax t. If the presumptive tax is sufficiently large, people will pay the 1 cent to buy the license if and only if they currently would like themselves to consume potato chips in the future. And because from a prior per-

spective both fully self-controlled people and people with self-control problems want to maximize their long-run preferences, everyone consumes (nearly) optimally.

To more fully illustrate this intuition, consider such a policy in our model. For fully self-controlled people, anyone who has $v > c + p + 1¢$ will purchase the license and later consume potato chips, while anyone who has $v < c + p + 1¢$ will not purchase the license and not consume potato chips. Because it is optimal to consume when $v > c + p$, fully self-controlled people consume (nearly) optimally.

For people with self-control problems who are naïve, they want to consume optimally in the future—that is, to consume only when they have $v > c + p$. Moreover, because they are naïve, they expect to behave optimally in the future, and thus they, like fully self-controlled people, purchase the license when they have $v > c + p + 1¢$. When it comes time to consume and they have an increased desire to consume, everyone who bought a license will consume. What about those who did not buy the license? They will choose to consume only if $v > \beta(c + p + t)$, and so as long as we set the tax to be large enough—specifically, $t \geq ((1 - \beta)/\beta)(c+p)$—all naïfs who did not buy the license will end up not consuming. Hence, naïfs also end up consuming optimally.

Finally, consider people with self-control problems who are sophisticated. Like naïfs, they want to consume in the future when they have $v > c + p$, and so everyone who has such tastes will buy the license. Sophisticates with $v < c + p$ would like to not buy the license and not consume. As long as we set the tax to be large enough—again, $t \geq ((1 - \beta)/\beta)(c + p)$—they'll correctly predict that not buying the license will lead to not consuming, and so they'll not buy the license. Hence, sophisticates end up consuming optimally as well.

The essence of this example is that by setting the tax that must be paid without a license to be sufficiently high, the license effectively creates a "free" means to commit oneself to future consumption or non-consumption. For the case in which people only consume in the future and know exactly what their future tastes will be, fully self-controlled people are not hurt by being forced to make such a commitment, and people with self-control problems are in fact helped by being forced to make such a commitment. But this case is clearly very special and unrealistic. We next consider several more realistic complications.

An important way in which the environment above may be misleading is that we have assumed that everyone who will want a license for the future will indeed have it when the future arrives. If, however, there are new customers—or perhaps customers who buy potato chips infrequently—or if people forget their license at home, or lose it, or didn't have time to buy one in the first place, then a problem arises, because the large presumptive tax may create suboptimal nonconsumption among this group. Even so, the basic conclusion still holds, except that we might set a smaller presumptive tax, accepting some suboptimal consumption among people with self-control problems in order to reduce distortions among people who don't have a license.

To illustrate, let's reconsider the environment above when a portion of the population arrives without having faced the prospective license decision. Using our logic above, for any tax $t > 0$, fully self-controlled people who get to make the prospective license decision will consume optimally. For those who did not make the prospective license decision—and hence have no license—they will consume when $v > c + p + t$. As in Section 3 (recall Figure 2), those with $c + p < v < c + p + t$ don't consume when they should. The larger is t, the larger is this distortion.

Among people with self-control problems who are naïve, again they will buy the license if and only if it is optimal to consume—if and only if they have $v > c + p$. In the moment, all naïfs who bought the license clearly choose to consume. Among those who must pay the tax—either because they chose not to buy the license or because they never had the opportunity—they will buy if they have $v > \beta(c + p + t)$. If the presumptive tax is sufficiently small that $\beta(c + p + t) < c + p$, then naïfs end up consuming whenever they have $v > \beta(c + p + t)$. Hence, as in Section 3 (again recall Figure 2), naïfs with $\beta(c + p + t) < v < c + p$ consume when they should not, and the smaller is t, the larger is this distortion.

Among people with self-control problems who are sophisticated, from a prior perspective they perceive that, whether or not they purchase a license, they will buy potato chips whenever $v > \beta(c + p + t)$. Because they have no desire to pay the tax, they buy the license whenever $v > \beta(c + p + t)$. When it comes time to consume, all sophisticates who have a license obviously consume (as planned). Among sophisticates who never faced the prospective license decision, they also

will consume if they have $v > \beta(c + p + t)$. Hence, like naïfs, sophisticates end up consuming whenever they have $v > \beta(c + p + t)$, and hence the distortions for sophisticates are identical to the distortions for naïfs.[16]

In this situation, the optimal presumptive tax will thus trade off distortions among fully self-controlled people who didn't get to make the prospective license decision against distortions among people with self-control problems who are now less deterred from consumption even when paying the tax. For instance, if the consumption values are distributed uniformly, if proportion ϕ of the population has $\beta = \underline{\beta}$ while proportion $1 - \phi$ has $\beta = 1$, and if proportion θ of the population doesn't get to make the prospective license decision, one can derive that the optimal tax is

$$t = \frac{\phi\underline{\beta}(1-\underline{\beta})}{\phi\underline{\beta}^2 + (1-\phi)\theta}(c + p)$$

(see Appendix Box 3). Notice that the optimal presumptive tax decreases as θ gets larger, reflecting that as a larger proportion of the population is unable to secure a license in advance (or loses their license), the presumptive tax creates more distortions among this group, and hence it is optimal to reduce it.

As we discuss in Section 2, an important feature of many sin goods is that they are repeatedly purchased in small quantities for short-term consumption, and as a result a preference for immediate gratification may be particularly problematic. Hence, in addition to inducing people with self-control problems to make prospective choices, it is also useful to induce them to make less frequent longer-term consumption decisions. In fact, it may be possible to use sin licenses to accomplish this goal.

[16] In terms of economy-wide efficiency of potato-chip consumption, naiveté vs. sophistication is irrelevant. Notice, however, that naïfs are (on average) worse off than sophisticates because they more often consume without a license. Specifically, among people who make the prospective license decision and also have $\beta(c + p + t) < v < c + p$, naïfs and sophisticates both end up consuming, but only sophisticates ex ante choose to purchase a license. (For economy-wide efficiency, this difference is irrelevant because it merely creates a redistribution from naïfs to sophisticates and fully self-controlled people.)

Specifically, suppose that when people arrive at the store, we offer them two options. They can pay the presumptive tax t on today's purchases. Alternatively, they can buy a license for a fee $f > t$ that exempts them from the tax for one year starting now. If the presumptive tax is non-negligible, and as long as the license fee is not too large, such a scheme will induce people not to focus merely on whether to consume potato chips today, but rather to decide whether they want to buy potato chips frequently over the course of the upcoming year. Ideally, people for whom potato-chip consumption is optimal will purchase the license and then consume tax-free, while people for whom potato-chip consumption is not optimal will both not purchase the license and be deterred by the presumptive tax from purchasing any potato chips at all.

To illustrate this intuition, we again return to our simple model. In particular, consider the case where there is a presumptive tax t, but for a fee $f = t + 1¢$, a person can buy a license that exempts her from this tax for n periods beginning now. Given the stationarity of the environment, fully self-controlled people clearly would only consider long-term consumption (with the license) or no consumption, and they prefer the former when $v > c + p + t/n$. Notice the advantage of intertemporal bundling from the perspective of fully self-controlled individuals: Because the license fee can be "amortized" over many purchases, even a large presumptive tax may not deter too much consumption.

Among people with self-control problems, by setting a sufficiently large presumptive tax, we can minimize the set of people who suboptimally consume. Specifically, people who have $\beta(c + p + t) < v < c + p$ will suboptimally consume, and the larger is t, the fewer people will fall in this category. At the same time, because the license applies for many periods, people who should consume will choose to buy the license and indeed consume unless their net benefits from potato-chip consumption are small. In particular, anyone with $v > \beta(c + p + t) > c + p$ will clearly buy the license, and anyone with $\beta(c + p + t) > v > c + p$ will buy the license as long as $v - (c+p)$ is not too small.

Again, the optimal version of this policy of course depends on the details of the environment. For instance, if the consumption values are distributed uniformly, and if proportion ϕ of the population has $\beta = \underline{\beta}$ while proportion $1 - \phi$ has $\beta = 1$, one can derive that the optimal presumptive tax as a function of n is

$$t = \frac{\phi\underline{\beta}(1-\underline{\beta})}{\phi\underline{\beta}^2 + (1-\phi)/n^2}(c+p)$$

(see Appendix Box 4). This tax permits some suboptimal consumption among people with self-control problems in order to limit the amount of suboptimal non-consumption among fully self-controlled people. But the longer the license lasts, the closer we come to inducing the first-best outcome. In particular, as n becomes large, the presumptive tax approaches $((1-\underline{\beta})/\underline{\beta})(c+p)$, which is the tax that induces first-best behavior among people with self-control problems. At the same time, because as this tax is amortized over many periods, we also approach first-best behavior among fully self-controlled people as well.

Up to now, we have focused on the use of a 1 cent license. The reader may find this odd, and it is. A 1 cent license works above because of our strong assumption that people know exactly what their future tastes will be. If so, then the license has positive value for those whom we want to consume in the future, and no value or even negative value for everyone else. Hence, the 1 cent cost induces the right people to buy the license. More realistically, people face some uncertainty about their future tastes, and if so there is some option value to obtaining the license. In this case, things become considerably more complicated.

To build intuition, consider first the case in which people consume only in the future, and they have no idea at all what their future consumption values will be—e.g., future consumption values will be randomly drawn from some distribution. At first glance, because there is no way to screen on consumption values, it might seem that the best we can do is to implement the second-best optimal tax for the static case t^* (from Section 3). If the world is populated by fully self-controlled people and people with self-control problems who are naïve, then indeed we can do no better. If, however, some people with self-control problems are sophisticated, then there may be scope for improvement. Sophisticates recognize that having a lower tax tomorrow might increase the likelihood of future over-consumption. Hence, it might be optimal to set a high tax designed for sophisticates, while at the same time offering 1 cent license to avoid the tax designed for fully self-controlled individuals and naïfs. As long as the

benefits of better future consumption decisions outweigh the costs of possibly paying the higher tax, such a plan would make sophisticates better off. More generally, this case highlights how, for the case where future tastes are uncertain, sophisticates might benefit from the existence of sin licenses that they do not buy.

Next consider the case where there is limited persistence in consumption values. For instance, it might be that, with some probability μ, a person's future consumption value is identical to her current consumption value, but otherwise her future consumption value takes on some new random value. Could licenses improve upon the static tax t^* in this case? The logic above implies that a 1 cent license could possibly be useful if there are sophisticates in the population, because they can choose not to buy the license as a commitment device. But in fact, a more expensive license might help improve the outcomes for fully self-controlled individuals and naïfs. Specifically, consider an expensive license that permits tax-free purchases in the future, while the tax without a license is the static optimal tax t^*. Naïfs and fully self-controlled individuals will buy the license only if their current consumption value is large enough, and by setting the license fee appropriately, we can implement any cutoff current valuation $\underline{v}$ that we want. The benefit of this policy is that people with current valuations between $\underline{v}$ and $c + p + t^*$ will optimally consume in the future when their tastes persist, whereas they would not consume if we merely charged everyone the static tax t^*. The cost of this policy is that, among naifs with high current consumption values whose tastes do not persist, they might end up over-consuming due to facing the lower tax. As long as the persistence of consumption values is strong enough—μ is large enough—the benefits will outweigh the costs.

More generally, there may be scope for even more sophisticated schemes. In the simple environment that we have studied here, because the consumption choice is binary, there is limited scope for screening. In a more general model with continuous demands for potato chips, it may be possible to better screen consumers by offering a menu of licenses that differ in terms of the license fee and tax reduction—i.e., for a small fee you'll receive a small tax reduction in the future, and for a large fee you'll receive a large tax reduction in the future. Similarly, in a more general model with additional periods, it may be possible to better screen consumers by offering a menu of licenses that differ in terms of when the license is valid—i.e., you must choose between a license that is good only for the rest of this

month vs. a license that is valid starting next month and remains valid for the 6 months that follow. Another dimension that could be varied is whether licenses are specified in terms of time—when it is valid—or quantities—how many units it is valid for. We hope to address these possibilities further in future research.

As we have emphasized, the schemes we propose in this section rely on people making good prospective decisions. Hence, we need to be wary that people might be making other errors that undermine their ability to make good prospective decisions. If we are confident that, whatever the error, the problem is that people are prone to over-consume, then simple taxes of the form we discuss in Section 3 might be a safe policy that merely reduces over-consumption from any source and is unlikely to create further distortions. While we have demonstrated that more sophisticated schemes can in principle lead to even better outcomes if over-consumption is driven by self-control problems, these schemes might backfire if people make other errors. For instance, if over-consumption is driven by an irrational optimism about the health costs, then making choices prospectively doesn't change people's decisions. And if we incorrectly assume that people make good prospective decisions and offer, say, the 1 cent license, people who over-consume due to irrational optimism will be worse off than under a simple per-unit tax.

If people have projection bias of the form discussed in Section 2, then inducing prospective commitments could be even more problematic. In particular, if tastes fluctuate—such as in the way we discuss above—then people with projection bias will over-estimate future consumption values on days when their consumption value is high, and hence they'll over-value the license. Analogously, they will underestimate future consumption values on days when their consumption value is low, and hence they'll under-value the license. In other words, people with projection bias are prone to make bad prospective decisions. If such people are otherwise fully rational, then such errors may not be too problematic, and will mostly take the form of having to pay the high tax because they suboptimally chose not to buy a license. In contrast, if people are characterized both by projection bias when predicting future tastes and by in-the-moment preference for immediate gratification, then errors due to projection bias could be compounded with errors due to self-control problems. For instance, over-estimates of future consumption values may lead a person to buy a license, and then the person might end up consuming

despite low consumption values due to her preference for immediate gratification. Such issues are potentially quite important and require further investigation before implementing our proposed policies.

5. Discussion

We conclude by discussing some broader issues with respect to our proposed policies.

Throughout our analysis, we have assumed that any tax proceeds are returned to consumers via lump-sum transfers. While we have suggested that such lump-sum transfers might actually take the form of reduced income taxes or reduced taxes on other commodities, we have not addressed the possibility that the tax proceeds might lead to increased or decreased distortions. In principle, either might happen. On the positive side, to the extent that income taxes or other commodity taxes are distortionary, and using revenue from potato-chip taxes to reduce these taxes in fact reduces such distortions, the benefits from potato-chip taxes would be even larger than our analysis suggests. On the negative side, however, the tax proceeds might, in fact, not be returned to consumers, and instead used for some new wasteful spending. Indeed, there is a real worry that when governments find new sources of revenue, they use it to spend more rather than reduce other taxes.

In addition to worries about the tax proceeds being used for new wasteful spending, there are a number of other practical issues that might make one wary of pursuing taxes on sin goods. Because commodity taxes are often imposed at the state or even local level, there might be worries of smuggling when adjacent areas have different tax rates. There might also be worries that taxes on sin goods are inherently regressive. Of course, one must recognize that they are regressive only in monetary terms, since the whole point is that the people paying the tax lose income but gain in terms of more efficient consumption. Moreover, any regressivity in sin taxes could, in principle, be counter-balanced with increased progressivity in other taxes.

Two additional worries are (i) that any attempt to apply the ideas in this paper would mean an overload of new regulations in terms of which goods are taxed and which goods are not, and (ii) that such taxes interfere with some people's libertarian views that the government should not be influencing people's decisions. While both of these worries may be valid, we do not believe they provide a blanket

reason not to proceed. Rather, we should incorporate these concerns into the calculus of deciding whether to proceed. Indeed, we suspect people will be less concerned with these worries once they recognize the arbitrariness of existing taxes—in the United States, for instance, in many states, socks and underwear are taxed, while soda and potato chips are not. Hence, we would not be moving from a world of no taxes to a world of taxes on specific goods, or even from a world of uniform taxes across all goods to a world of taxes on specific goods, but rather from a world of taxes on one set of goods to a world of taxes on a different set of goods.

A major feature of preferences from which our analysis has abstracted away is habit formation. Many goods are habit-forming in the sense that current consumption influences future tastes for the good. Addictive products such as cigarettes and alcohol are obvious examples, but one could even argue that people develop a taste for potato chips or similar goods. We have abstracted away from habit formation because doing so makes the analysis considerably more tractable and transparent. We recognize, however, the importance of assessing how habit formation might change our conclusions.[17]

Finally, we note an ancillary contribution of our analysis. Many of the more sophisticated schemes that we discuss are designed to provoke different reactions from different types of people. As such, if we were to implement these policies and catalog people's responses, we would effectively be running an experiment that could test for the existence of and nature of self-control problems in the population. This approach might prove useful in the future.

[17] If, due to addictive properties, the people prone to over-consume are also likely to have price-insensitive demand, it could conceivably be optimal to subsidize the good (see Bernheim and Rangel, 2004).

Appendix

Box 1

Suppose that consumption values v are distributed according to $F(v)$ with support $[0, \bar{v}]$ with $\bar{v}$ large, suppose that self-control problems are distributed according to $G(\beta)$ with support $[\underline{\beta}, 1]$ with $\underline{\beta} < 1$, and suppose that these distributions are independent.

The net benefits from potato-chip consumption (excluding taxes and lump-sum transfers, which cancel out) are $v - c - p$ if a person consumes and 0 if a person does not consume. Amongst people with a specific β, as a function of t, they consume when $v > \beta(c + p + t)$, and so their average net benefits from potato-chip consumption are

$$w(t;\beta) \equiv \int_{\beta(c+p+t)}^{\bar{v}} (v - c - p)\,dF(v)\,.$$

Note that

$$\frac{dw(t;\beta)}{dt} = \beta\left[(1-\beta)(c+p) - \beta t\right] f\left(\beta(c+p+t)\right).$$

Weighting everyone equally, the social-welfare function is

$$\Omega(t) = \int_{\underline{\beta}}^{1} w(t;\beta)\,dG(\beta)\,,$$

and so

$$\frac{d\Omega(t)}{dt} = \int_{\underline{\beta}}^{1} \frac{dw(t;\beta)}{dt}\,dG(\beta)\,.$$

Because $dw(t;\beta)/dt$ evaluated at $t = 0$ is positive for people with $\beta < 1$ and zero for people with $\beta = 1$, $d\Omega(t)/dt$ evaluated at $t = 0$ is positive, and so the optimal potato-chip tax $t^* > 0$.

We can derive an optimal tax if we impose more structure on the distributions of v and β. For instance, suppose that consumption values are distributed uniformly on $[0, \bar{v}]$, and suppose that proportion ϕ of the population has $\beta = \underline{\beta}$ while proportion $1 - \phi$ has $\beta = 1$. Then:

$$\frac{d\Omega(t)}{dt} = \phi\left(\underline{\beta}\left[(1-\underline{\beta})(c+p) - \underline{\beta}t\right]\right)\frac{1}{\bar{v}} + (1-\phi)(-t)\frac{1}{\bar{v}},$$

and therefore the optimal tax is

$$t^* = \frac{\phi\underline{\beta}(1-\underline{\beta})}{\phi\underline{\beta}^2 + (1-\phi)}(c+p).$$

Box 2

We build off the analysis in Box 1 for the case in which consumption values are distributed uniformly on $[0, \bar{v}]$ and proportion ϕ of the population has $\beta = \underline{\beta}$ while proportion $1 - \phi$ has $\beta = 1$. Let L denote the lump-sum transfer that everyone in the population receives, and assume that all tax proceeds are returned to consumers. Amongst people with $\beta = \underline{\beta}$, as a function of t, they consume when $v > \underline{\beta}(c + p + t)$, and so the proportion that consumes potato chips is $(\bar{v} - \underline{\beta}(c + p + t))/\bar{v}$. Amongst people with $\beta = 1$, as a function of t, they consume when $v > c + p + t$, and so the proportion that consumes potato chips is $(\bar{v} - (c + p + t))/\bar{v}$. Because the lump-sum transfer is just equal to the tax proceeds, $L = t^*[\phi(\bar{v} - \underline{\beta}(c + p + t))/\bar{v} + (1 - \phi)(\bar{v} - (c + p + t))/\bar{v}]$.

Hence, the average net income transfer amongst people with $\beta = \underline{\beta}$ is $z = L - t^*(\bar{v} - \underline{\beta}(c + p + t))/\bar{v} = -t^*[(1 - \phi)(1 - \underline{\beta})(c + p + t)/\bar{v}]$, and the average net income transfer amongst people with $\beta = 1$ is $z = L - t^*(\bar{v} - (c + p + t))/\bar{v} = t^*[\phi(1 - \underline{\beta})(c + p + t)/\bar{v}]$.

Average welfare amongst people with a specific β is $\hat{w}(t;\beta) \equiv w(t;\beta) + z$. Because

$$\frac{d\hat{w}(t;\underline{\beta})}{dt} = \frac{1}{\bar{v}}\left[(1-\underline{\beta})(\underline{\beta}-(1-\phi))(c+p) - (\underline{\beta}^2 + 2(1-\phi)(1-\underline{\beta})t\right]$$

and

$$\frac{d\hat{w}(t;1)}{dt} = \frac{1}{\bar{v}}\left[(1-\underline{\beta})\phi(c+p) - (1-2\phi(1-\underline{\beta}))t\right],$$

and because, as long as $\underline{\beta} > (1-\phi)$, both of these are positive at $t = 0$, positive sin taxes can make both groups on average better off.

Box 3

We proceed as in Box 1, except that we must re-derive the average net benefits from potato-chip consumption when proportion θ of the population doesn't get to make the licensing decision.

Among fully self-controlled people, those who make the prospective license decision end up consuming when $v > c + p$, and those who did not make the prospective license decision end up consuming when $v > c + p + t$. Hence, the average net benefits from potato-chip consumption among people with $\beta = 1$ are

$$w(t;1) = (1-\theta)\int_{c+p}^{\bar{v}} (v-c-p)\frac{1}{\bar{v}}dv + \theta\int_{c+p+t}^{\bar{v}} (v-c-p)\frac{1}{\bar{v}}dv.$$

Note that $\dfrac{dw(t;1)}{dt} = -\theta t / \bar{v}$.

Among people with self-control problem $\underline{\beta}$, regardless of whether they make the prospective license decision, and regardless of whether they are sophisticated or naïve, they end up consuming when $v > \underline{\beta}(c + p + t)$. Hence, the average net benefits from potato-chip consumption among people with $\beta = \underline{\beta}$ are

$$w(t;\underline{\beta}) = \int_{\underline{\beta}(c+p+t)}^{\bar{v}} (v-c-p)\frac{1}{\bar{v}}dv.$$

Note that $\frac{dw(t;\underline{\beta})}{dt} = \underline{\beta}\left[(1-\underline{\beta})(c+p) - \underline{\beta}\,t\right]/\bar{v}$.

Because $\Omega(t) = \phi\, w(t;\underline{\beta}) + (1-\phi)\, w(t;1)$,

$$\frac{d\Omega(t)}{dt} = \phi\left(\underline{\beta}\left[(1-\underline{\beta})(c+p) - \underline{\beta}t\right]\right)\frac{1}{\bar{v}} + (1-\phi)\left(-\theta\, t\right)\frac{1}{\bar{v}},$$

and therefore the optimal tax is $t = \frac{\phi\underline{\beta}(1-\underline{\beta})}{\phi\underline{\beta}^2 + (1-\phi)\theta}(c+p)$.

Box 4

Again, we proceed as in Box 1, except that we must re-derive the average net benefits from potato-chip consumption. Also, note that while the license lasts for n periods, for all types behavior is the same in every period, and so we can focus on maximizing per-period social surplus.

Fully self-controlled people end up consuming when $v > c + p + t/n$. Hence, their average net benefits from potato-chip consumption are

$$w(t;1) = \int_{c+p+t/n}^{\bar{v}} (v-c-p)\frac{1}{\bar{v}}dv.$$

Note that $\frac{dw(t;1)}{dt} = -t/(n^2\bar{v})$.

To understand people with self-control problem $\underline{\beta}$, first note that it would never be optimal to set $t > ((1-\underline{\beta})/\underline{\beta})(c+p)$, because otherwise we would create distortions for all types. For any $t \le ((1-\underline{\beta})/\underline{\beta})(c+p)$, $\underline{\beta}(c+p+t) < c+p$, and so people end up consuming whenever $v > \underline{\beta}(c+p+t)$, regardless of whether they are sophisticated or naïve. Hence, their average net benefits from potato-chip consumption are

$$w(t;\underline{\beta}) = \int_{\underline{\beta}(c+p+t)}^{\bar{v}} (v-c-p)\frac{1}{\bar{v}}dv.$$

Note that

$$\frac{dw(t;\underline{\beta})}{dt} = \underline{\beta}\left[(1-\underline{\beta})(c+p) - \underline{\beta}\,t\right]/\bar{v}\,.$$

Because $\Omega(t) = \phi\, w(t;\underline{\beta}) + (1-\phi)\, w(t;1)$,

$$\frac{d\Omega(t)}{dt} = \phi\left(\underline{\beta}\left[(1-\underline{\beta})(c+p) - \underline{\beta}t\right]\right)\frac{1}{\bar{v}} + (1-\phi)\left(-t/n^2\right)\frac{1}{\bar{v}},$$

and therefore the optimal tax is $t = \dfrac{\phi\underline{\beta}(1-\underline{\beta})}{\phi\underline{\beta}^2 + (1-\phi)/n^2}(c+p)$.

References

Ainslie, G. (1991), Derivation of rational economic behavior from hyperbolic discount curves, American Economic Review 81, 334-340.

Ainslie, G. (1992), Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person, Cambridge University Press, New York.

Ainslie, G. and Haslam, N. (1992a), Self- control, in G. Loewenstein and J. Elster (eds.), Choice Over Time, Russell Sage Foundation, New York.

Ainslie, G. and Haslam, N. (1992b), Hyperbolic discounting, in G. Loewenstein and J. Elster, (eds.), Choice Over Time, Russell Sage Foundation, New York.

Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J. and Weinberg, S. (2001), The hyperbolic buffer stock model: Calibration, simulation, and empirical evaluation, Journal of Economic Perspectives 15, 47-68.

Benabou, R. and Tirole, J. (2002), Self-confidence and personal motivation, Quarterly Journal of Economics 117, 871-915.

Bernheim, B.D. and Rangel, A. (2004), Addiction and cue-triggered decision processes, American Economic Review 94, 1558-1590.

Besley, T. (1988), A simple model for merit good arguments, Journal of Public Economics 35, 371-383.

Camerer, C., Issacharoff, S., Loewenstein, G., O'Donoghue, T. and Rabin, M. (2003), Regulation for conservatives: Behavioral economics and the case for "asymmetric paternalism", University of Pennsylvania Law Review 151, 1211-1254.

Carrillo, J. and Mariotti T. (2000), Strategic ignorance as a self-disciplining device, Review of Economic Studies 67, 529-544.

Choi, J., Laibson, D., Madrian, B. and Metrick, A. (2003), Optimal defaults, American Economic Review 93, 180-185.

Fischer, C. (1999), Read this paper even later: Procrastination with time-inconsistent preferences, Resources for the Future Discussion Paper 99-20, Resources for the Future, Washington, D.C.

Frederick, S., Loewenstein, G. and O'Donoghue, T. (2002), Time discounting and time preference: A critical review, Journal of Economic Literature 40, 351-401.

Gruber, J. and Koszegi, B. (2004), Tax incidence when individuals are time-inconsistent: The case of cigarette excise taxes, Journal of Public Economics 88, 1959-1987.

Gruber, J. and Mullainathan, S. (2005), Do cigarette taxes make smokers happier?, B.E. Journals: Advances in Economic Analysis & Policy 5, Article 4.

Herrnstein, R., Loewenstein, G., Prelec, D. and Vaughan, W. Jr. (1993), Utility maximization and melioration: Internalities in individual choice, Journal of Behavioral Decision Making 6, 149-185.

Laibson, D. (1997), Hyperbolic discounting and golden eggs, Quarterly Journal of Economics 112, 443-477.

Laibson, D. (1998), Life-cycle consumption and hyperbolic discount functions, European Economic Review 42, 861-871.

Loewenstein, G. and Prelec, D. (1992), Anomalies in intertemporal choice: Evidence and an interpretation, Quarterly Journal of Economics 107, 573-597.

Loewenstein, G., O'Donoghue, T. and Rabin, M. (2003), Projection bias in predicting future utility, Quarterly Journal of Economics 118, 1209-1248.

Musgrave, R.A. (1959), The Theory of Public Finance, McGraw-Hill, New York.

O'Donoghue, T. and Rabin, M. (1999a), Procrastination in preparing for retirement, in H. Aaron (ed.), Behavioral Dimensions of Retirement Economics, Brookings Institution Press and Russell Sage Foundation, Washington, D.C. and New York.

O'Donoghue, T. and Rabin, M. (1999b), Doing it now or later, American Economic Review 89, 103-124.

O'Donoghue, T. and Rabin, M. (2001), Choice and procrastination, Quarterly Journal of Economics 116, 121-160.

O'Donoghue, T. and Rabin, M. (2003), Studying optimal paternalism, illustrated by a model of sin taxes, American Economic Review 93, 186-191.

O'Donoghue, T. and Rabin, M. (2005), Optimal sin taxes, Mimeo, Cornell University.

Pollak, R.A. (1968), Consistent planning, Review of Economic Studies 35, 201-208.

Sunstein, C. and Thaler, R.H. (2003a), Libertarian paternalism, American Economic Review 93, 175-179.

Sunstein, C. and Thaler, R.H. (2003b), Libertarian paternalism is not an oxymoron, University of Chicago Law Review 70, 1159-1202.

Strotz, R.H. (1956), Myopia and inconsistency in dynamic utility maximization, Review of Economic Studies 23, 165-180.

Thaler, R.H. (1991), Some empirical evidence on dynamic inconsistency, in R.H. Thaler (ed.), Quasi Rational Economics, Russell Sage Foundation, New York, 127-133.

Thaler, R.H, and Loewenstein, G. (1992), Intertemporal choice, in R.H. Thaler (ed.), The Winner's Curse: Paradoxes and Anomalies of Economic Life, Free Press, New York.