SWEDISH ECONOMIC POLICY REVIEW 12 (2005) 65-94

On the feasibility of market solutions to self-control problems

Botond Köszegi*

Summary

■ I consider behavior and welfare in competitive markets supplying harmful or beneficial goods when consumers, at each moment in time, prefer immediate gratification more than they would themselves approve. In a spot market and without advance commitment to her consumption level, a decision maker overconsumes harmful goods and underconsumes beneficial goods from her own point of view. The optimal level of consumption can be achieved by adjusting prices in a way that forces a consumer to fully internalize the future consequences of consumption. Calibrations in the case of tobacco and exercise indicate that the optimal interventions can be very large. I examine the extent to which market forces can provide such "self-control" interventions, when firms and consumers can agree on price schedules ex ante. If consumers are bound to the firm whose offer they accept and price schedules are restricted to two-part tariffs, consumption is much closer to optimal than in a spot market. But if consumers cannot be prevented from purchasing from other firms ex post, the consumption of harmful goods is as if only a spot market was available. And if firms can engage in non-linear pricing, perfectly sophisticated consumers consume optimally in equilibrium, but individuals with an arbitrarily small amount of overoptimism regarding their self-control problem consume as if they were buying on the spot market. Hence, government intervention may typically be necessary to correct self-control problems. ■

JEL classification: D49, D69, D91.
Key words: intertemporal choice, hyperbolic discounting, self-control, market solution.

* *Botond Köszegi is Assistant Professor at the Department of Economics, University of California, Berkeley.*

SWEDISH ECONOMIC POLICY REVIEW 12 (2005) 65-94

On the feasibility of market solutions to self-control problems

Botond Köszegi*

Evidence and intuition indicate a fundamental conflict in how people approach choices over time: while most generally recognize that they should invest into the future, they act more short-sightedly when it actually comes down to sacrificing momentary pleasure for it. A (Berkeley) professor may think that he should mostly work on papers rather than play basketball, but to his general dismay, at each moment in time, he still prefers to play basketball. A dieting person may know that she should not have dessert with most meals, but still want to order one every time she is actually confronted with the immediate choice. And a smoker may believe that she should soon quit, but still prefers to smoke a few cigarettes when attending a party.

This paper considers the implications of the above conflict for the consumption of "intertemporal goods" such as tobacco, fatty or salty foods, vegetables, and exercise, that have consequences for future utility. I show that a consumer's taste for immediate gratification leads her, by her own reckoning, to systematically overconsume harmful products and underconsume beneficial ones. This "self-control problem" raises the specter for welfare-increasing interventions in her behavior. I argue that in many environments, there are fundamental limitations to the market's ability to provide consumers with self-control, especially when it comes to harmful products. The results indicate that in many cases, the government or another social planner may be in a unique position to mitigate self-control problems.

Section 1 introduces a simple three-period model with periods 0, 1, and 2, where the consumption of a standard good a and an intertemporal good b occurs only in period 1. Consumption of a unit of good b yields immediate pleasure or displeasure as well as harm h in period 2; h could be positive or (for a beneficial good) negative. To capture a

* *This paper draws heavily on joint work with Jonathan Gruber, as well as on numerous conversations over the years with many colleagues. I would especially like to thank Stefano DellaVigna and Matthew Rabin for helpful comments.*

taste for immediate gratification of which the consumer would herself disapprove, I follow the methodology of Strotz (1956), Phelps and Pollak (1968), and Laibson (1997) in assuming that she is a "hyperbolic discounter:" whereas self 1, her period-1 incarnation, discounts the future harm from consuming b using the discount factor $\beta\delta$, self 0's utility and the consumer's overall welfare are computed using $\delta > \beta\delta$ instead. To capture differences in the consumer's understanding of her future impatience, I assume that self 0 believes self 1's discount factor between periods 1 and 2 to be $\hat{\beta}\delta$ (O'Donoghue and Rabin, 2001). Then, $\hat{\beta} = \beta$ means that the consumer is perfectly "sophisticated", while $\hat{\beta} = 1$ means that she is perfectly "naïve" about her future impatience.

To illustrate the basic problem, suppose good b is cigarettes (so $b > 0$), and that it is supplied by a competitive spot market in period 1. Because the consumer makes her smoking decision in period 1, exactly when she is most impatient about it, she tends to smoke too much. Section 2 shows that a simple way of correcting this self-control problem is to introduce a lump-sum-redistributed price increase for cigarettes that forces self 1 to fully internalize the future harm from smoking. Calibrations based on Gruber and Köszegi (2001) indicate that the price adjustment on a pack of cigarettes should be upwards of 3 USD even for small self-control problems, and could be as high as 12 USD. I perform a similar calculation for exercise, and find that the optimal subsidy per thirty-minute exercise session is between 2.50 USD and 10 USD. These price adjustments could be achieved by the government through a lump-sum-redistributed tax or, in principle, through the market by a variety of ex-ante contracting or other commitment arrangements.

Section 3 considers whether the market would indeed be able to provide this self-control tool to consumers in various environments and, by implication, whether government intervention may be warranted if the market fails. I begin by assuming that in period 0, competitive firms offer two-part tariffs that specify a lump-sum payment and a per-unit price that will apply in period 1, and once a consumer accepts a firm's offer, she cannot or does not want to buy from other firms. This environment is very similar to that of DellaVigna and Malmendier (2004). Consider their example of exercise; the conclusions for harmful goods are symmetric. Sophisticated consumers un-

derstand that they would tend to consume too little of this beneficial good in period 1, and so value and demand tariffs where they pay up-front for a cheaper per-use price (which encourages them to exercise more often). Profit-maximizing firms will satisfy this demand and hence, the market is very effective in providing self-control to these consumers. Because a naïve gym-goer believes she will behave optimally in the future, she does not demand any help in self-control per se. Nevertheless, her misperception of future behavior leads firms to offer her the kind of contract they sell to sophisticates: since she thinks she will exercise more than she actually will, she can be "scooped" into paying a great deal for a membership with a lower per-use price (DellaVigna and Malmendier, 2004). While intended to take advantage of naïveté, the lower per-use price has the beneficial side-effect of providing self-control. This suggests that the market does not systematically undersupply self-control and consequently, that government intervention aimed at getting people to act more patiently is not necessary.

The same conclusion does not hold in other market environments. Next, suppose that firms and consumers can still agree on two-part tariffs in period 0 that govern their interaction in period 1, but consumers cannot be prevented from purchasing the good on a spot market in period 1. This applies to situations where exclusive contracts are not possible and switching costs between suppliers are small. Then, the per-unit price of good b cannot exceed its marginal cost.

Since the optimal contracts for beneficial goods involve a per-unit price below the marginal cost, these contracts can coexist with a spot market. For harmful goods, however, the optimal contracts above cannot survive competition from the spot market. In fact, for such goods, the outcome is as if firms and consumers could not interact in period 0 at all. Intuitively, no bilateral contract can commit a consumer to pay a high price for hamburgers, because she always has the option of buying her burgers elsewhere. That is, the attempt of one firm to provide self-control to a consumer is undermined by the incentive of other firms to profit from the very self-control problem the consumer is trying to solve. A major lesson, therefore, is that the market is significantly better at providing self-control in the case of beneficial goods than in the case of harmful goods. The government is in a unique position to solve the problem, because it can impose a price increase that applies to all transactions in the market.

Finally, once more assume that firms can offer pricing schedules in period 0 and that there is either no spot market in period 1, or that consumers do not use it due to switching costs. Suppose, however, that in period 0, firms can make more complicated offers, where they are not restricted to linear per-unit prices. For perfectly sophisticated consumers, the equilibrium contract is still one that yields the ex-ante optimal consumption level in period 1. The same, however, is not an equilibrium if consumers are even slightly naïve. A firm can deviate by offering the ex-ante optimal consumption level at a slightly lower total price, and introduce an option to pay a penalty and switch to an alternative choice that a consumer will barely prefer in period 1 (as in Eliaz and Spiegler, 2004). Since the consumer believes she will not switch, she accepts this contract, undermining her self-control. In fact, to make switching most tempting in period 1 at the lowest possible cost, the firm's alternative option gives self 1 the same consumption that she would choose in a spot market. Hence, for an arbitrarily small amount of naïveté, the market offers consumers no self-control whatsoever.

The key idea behind this result is that a profit-maximizing firm does not just want to sell self-control to consumers. While profiting from the demand for self-control is one of its motives, it also wants to take advantage of the consumer's impatience in period 1. If it can use non-linear prices, it can achieve both goals with a contract that the consumer believes provides enough self-control, but in fact affords none. Indeed, although credit card financing is, on average, very expensive, companies regularly offer grace periods, teaser rates, and other perks to induce more borrowing by consumers. And while casinos have very high margins on their gambling units, they use free alcohol and other bonuses to induce consumers to forget about such costs and gamble. Once again, the government is in a unique position to move consumption levels closer to the optimal, by taxing or subsidizing the production of the good in question.

The paper is organized as follows. Section 1 introduces and analyzes the basic model of consumer behavior. Section 2 solves for the optimal price intervention to correct the consumer's self-control problem, and calibrates it for smoking and exercise. Section 3 considers to what extent the market can provide self-control to consumers. Section 4 concludes.

1. A basic model of self-control problems

1.1. The framework

To address the necessity of government intervention in the market for intertemporal goods, I first introduce a model of individual choice in such products. The model combines a simple formalization of what it means for a consumption good to be harmful or beneficial, with the idea that consumers systematically behave more impatiently in the present than they would have preferred in the past. The consumer lives for three periods, $t = 0, 1, 2$. Consumption only occurs in period 1. In that period, the consumer has a choice between goods *a* and *b*. Good *a* is a classical consumption good that is pleasant to consume at the moment and has no implications for future utility. Good *b* may be pleasant or unpleasant to consume in period 1, and it also provides benefits or causes harm in period 2. To capture these properties, consider the following per-period utility functions:

	period $t = 0$	period $t = 1$	period $t = 2$
utility u_t	0	$a + u(b)$	$-h \cdot b$

Consumption of a unit of good *b* leads to harm *h* in period 2. If *h* is positive, the good is harmful, and the greater is *h*, the more harmful it is. If *h* is negative, the good is beneficial; the smaller (the more negative) is *h*, the more beneficial it is. I normalize the price of good *a* in period 1 to 1. Let *p* be the price of good *b* in that period, and denote income by I. I assume that *u* is strictly concave and parameters are such that the consumer's maximization problem has an interior solution, and market equilibrium always exists.

The second important aspect of consuming intertemporal goods concerns the nature in which the per-period utilities u_t are integrated into a global utility function. The crucial assumption of the model is that the consumer's preferences regarding the choice between *a* and *b* change over time. To highlight this distinction, I will call the person's period-0 incarnation "self 0," and her period-1 incarnation "self 1." Self 0's discounted utility function is

$$u_0 + \beta\delta u_1 + \beta\delta^2 u_2,$$

while self 1's discounted utility function is

$$u_1 + \beta\delta u_2.$$

Parameters β and δ are between zero and one. This "hyperbolic-discounting" formulation captures a specific type of self-control problem: since the discount factor self 0 applies between periods 1 and 2 (δ) is greater than the discount factor self 1 applies between the same two periods ($\beta\delta$), self 0 would like to behave more patiently in period 1 than self 1 actually does. This creates a conflict between selves 0 and 1, and means that self 0 benefits from self-control that brings period-1 behavior in line with her wishes. By making the implicit parameter restriction $\beta = 1$, standard models assume that no such self-control problem exists.

To analyze the nature of the above self-control problem and whether the market can help the consumer solve it, I assume that she cannot restrict her own behavior other than through market channels. That is, self 1 is free to choose any consumption level within her market opportunity set.

1.2. Overconsumption of harmful and underconsumption of beneficial goods

Consider first the behavior of self 1. Since there is no consumption in periods 0 and 2, she spends all her money: $a = I - pb$. Substituting this into self 1's discounted utility function, she is solving

$$\max_b I - pb + u(b) - \beta\delta hb.$$

This leads to the following first-order condition for the uncommitted level of consumption, b^u:

$$u'(b^u) = p + \beta\delta h. \tag{1}$$

Equation (1) is a variant of the familiar fact that a utility-maximizing consumer sets the marginal utility of consumption equal to its price in terms of the numeraire. When goods have future consequences, those consequences constitute an additional cost or benefit to consumption.

Thus, the consumer sets marginal utility equal to the "harm-inclusive" price of the good. Crucially, in this calculation, the harm is discounted according to self 1's discount factor between periods 1 and 2, $\beta\delta$.

Contrast the above with the consumption level self 0 would like self 1 to choose. Self 0's problem is

$$\max_b \beta\delta \cdot \left[I - pb + u(b)\right] - \beta\delta^2 \cdot hb.$$

yielding the first-order condition for self 0's preferred consumption level, b^*:

$$u'(b^*) = p + \delta h. \tag{2}$$

Similarly to self 1, self 0 sets the marginal utility from consumption equal to its harm-inclusive price. Because self 0 views the tradeoff between periods 1 and 2 differently from self 1, she uses a different discount factor, δ, in the evaluation of the harm.

For a harmful good ($h > 0$) and any $\beta < 1$, Equations (1) and (2) imply that $b^u > b^*$, so that self 1 overconsumes from the perspective of self 0. For a beneficial good ($h < 0$), the equations imply that $b^u < b^*$, so that self 1 underconsumes from the point of view of self 0. If $h = 0$, self 1's consumption is exactly at the level preferred by self 0. If a good does not have any future consequences, the fact that selves have different preferences for how the future should be discounted does not create a disagreement regarding the level of consumption.

While selves 0 and 1 disagree about consumption in period 1, from the perspective of this model there is no immediate reason to put priority on either self's preferences when evaluating the individual's overall welfare. In line with much of the literature (Gruber and Köszegi, 2001; DellaVigna and Malmendier, 2004; Gruber and Köszegi, 2004; O'Donoghue and Rabin, 2005), however, I will take the position that self 0's preferences are very close to the appropriate measure for welfare evaluation. First, in reality, consumption decisions of the type in this paper are repeated many times. In that case, a person underweights the future consequences of period-t consumption only in period *t*; *all* earlier selves would like her to place a higher weight on later outcomes. For example, when considering whether to consume cigarettes on the week of March 6, 2006, a person would at virtually any time prior to that date like to smoke very little. Only

when March 6, 2006, arrives does she prefer to smoke more. Hence, it seems reasonable to base welfare judgments on the preferences of earlier selves (O'Donoghue and Rabin, 2005). Second, as emphasized by Gruber and Köszegi (2001), bringing the decision maker's consumption level in each period of a longer-horizon model closer to the wishes of earlier selves in fact increases the discounted utility of all selves; hence, such interventions are welfare-improving under virtually any criterion. Third, one interpretation of Bernheim and Rangel's (2004) cue-conditioned-consumption model is that a self's excessive focus on the short run when making a consumption choice is a mistake. It is arguably appropriate not to respect such mistaken inclinations in welfare.[1]

2. Optimal interventions

The previous section has shown that self 1's unrestricted choice of consumption is suboptimal from the point of view of self 0, and has argued that this point of view is the appropriate one for the individual's broader welfare as well. I now turn to investigating possible interventions to correct self 1's behavior. Throughout the paper, I consider competitive-market environments and assume the marginal cost of producing good *b* to be *c*. I also assume that in the case of government intervention, revenues are lump-sum-redistributed to consumers. Since firms make no profits and the government returns all revenue, the social impact of any intervention is identical to its impact on consumer welfare.[2] And since consumers have quasi-linear preferences and consume all resources in the economy, their welfare depends entirely on the consumption of good *b*. Hence, all subsequent

[1] It is worth emphasizing why the model does not incorporate the considerations above. The focus of this paper is not on establishing the appropriate welfare criterion when self-control problems are present—but to take one criterion and investigate the market's and the government's ability to maximize welfare as defined by that criterion.

[2] If firms have market power and hence make positive profits, the welfare evaluation of a possible intervention would have to consider the impact on firms' profits. Since firms with market power tend to increase the price above the marginal cost, market power in a spot market mitigates self-control problems for harmful goods and exacerbates it for beneficial goods. This indicates that in the case of harmful products, making markets competitive without attending to consumers' self-control problems may actually lower social welfare.

analysis will address whether self 1 consumes the right amount of good b.

One important point regarding interventions must be emphasized. Many of the goods that potentially fit the model also have serious (negative or positive) externalities that might themselves call for an intervention in the market. I work with a model that ignores these externalities to keep the analysis simple, and to focus on the novelties hyperbolic discounting brings into the analysis. If externalities were present, classical reasons and methods of government intervention would be justified *on top of* any interventions implied by the current paper, and my sense is that there are no serious interactions between the two issues.

For studying the effectiveness of interventions, I take as a benchmark the situation where there is no market or government-imposed restriction on the consumer's behavior. That is, the good is available to her at marginal cost, as it would be on a competitive spot market. I start by considering "optimal linear price interventions" of the following form. An adjustment τ is added to the per-unit price p (= c) that the consumer faces in period 1. The proceeds (which are positive for $\tau > 0$ and negative for $\tau < 0$) are lump-sum redistributed to the consumer at the beginning of period 1. Although this formalism is equivalent to lump-sum-redistributed linear taxes, I instead use the term "intervention" to emphasize that a priori, it is not necessarily the case that only the government can achieve an optimal price adjustment.

A simple way of thinking about the optimal price intervention is the following. Self 0 would like self 1 to discount period 2 relative to period 1 by a factor of δ, and hence would prefer that self 1 counts the future implications of consuming each unit of good b at a value of δb. Self 1, in contrast, discounts between periods 1 and 2 by a factor of $\beta\delta$, and hence only counts the future implications of consuming b at a value of $\beta\delta b$. To align self 1's interests with those of self 0, one can impose an optimal price adjustment of $\tau = \delta b - \beta\delta b = (1 - \beta)\,\delta b$ per unit of good b. Due to this change in price, self 1 is forced to fully internalize the future benefit or harm of consumption.

To illustrate the implications of this simple formula, I consider two potential applications, cigarette consumption and exercise. In doing so, I must confront a major conceptual issue in taking the stylized model in this paper to an application. While the model is written in

discrete time, a person's life does not easily break into units corresponding to formal periods. Hence, in any application, one must choose the appropriate period length and an estimate of the parameter β corresponding to it. One way of thinking about this choice is the following. At the heart of my model and of hyperbolic-discounting models in general, there is a distinction between "present" and "future" and the assumption that a person discounts heavily between the present and all future periods, but does not discount so heavily between future periods. In any application, it is crucial to make an assumption about what it means for consumption to be in the present. For the applications below, the appropriate notion of the present period seems to be quite short, on the order of a week or a month. The estimates of β corresponding to such a horizon range between 0.6 and 0.9 (Thaler, 1981; Ainslie, 1992; Kirby and Herrnstein, 1995; Laibson, 1997).

The formula for the optimal price intervention combines β with an estimate of the discounted harm from consumption, δh. Gruber and Köszegi (2001) give one conservative estimate of the discounted harm from smoking. Although smoking has a number of adverse health and quality-of-life effects even when a person is alive, we focus exclusively on the fatal consequences of smoking. Based on Viscusi (1993), we estimate that, on average, Americans value their lives at 6.4 MUSD (in 2001 dollars). Assuming an annual discount rate of 4 percent (i.e. a yearly δ of 0.96) and taking the average age and lifespan for a smoker, an extra year at the end of life is worth 99,110 USD. Finally, using that smokers die on average 6.1 years early (Cutler et al., 2001), the average cost in terms of life years lost per pack of cigarettes is USD 30.45.

Given the high valuation people put on their lives and the tremendous damage smoking does to one's health, hyperbolic discounting implies phenomenally large intervention rates. Even for $\beta = 0.9$, a value that represents a modest self-control problem within the range indicated by the evidence, the optimal price intervention is $\tau =$ \$3.05. And going to the other end of the estimates for the short-run discount factor, $\beta = 0.6$, the optimal intervention could be as high as $\tau =$ USD 12.18. Once again, this adjustment is just for the fatal consequences of smoking, and is in addition to the taxes a social planner might want to impose for reasons other than self-control.

One important implication of this exercise is that in terms of its implications for government policy, the conventional assumption of $\beta = 1$ is a *knife-edge case*. It implies zero intervention—other than those justified by externalities—no matter how much damage people are doing to themselves through consuming harmful goods. In the case of smoking, where the self-inflicted damage is large, deviating just a little bit from the $\beta = 1$ model has tremendous consequences for the optimal intervention.

The lesson is very similar for exercise. I estimated the life-expectancy benefit of exercising in two ways. First, I found several life-expectancy calculators online that are purported to be based on accurate actuarial data.[3] I entered my own data as well as randomly chosen parameters for middle-aged Americans and, holding other answers fixed, varied the amount of exercise I reported in the lifestyle part of the calculator's questionnaire. Second, I searched the epidemiological literature for recent evidence on the health benefits of exercise (Lee and Skerrett, 2001; Bauman, 2004, for instance). Based on these sources, a reasonable estimate seems to be that 30 minutes of moderate intensity exercise 5 times a week adds 3 years to an average person's lifespan. Although large, there are reasons to believe that the estimate is in fact a lower bound. Both the life-expectancy calculators and the epidemiological studies control for blood pressure and weight in calculating the benefits of exercise, although the exercise can increase the lifespan through its effect on these risk factors as well. Naturally, any analysis of this sort also faces serious selection issues, as exercise is likely to be correlated with unobserved habits that improve one's health. While many of the studies attempt to control for such alternative causes, the endogeneity issues seem impossible to deal with perfectly.

Taking the 3-year figure seriously, however, and making similar assumptions as in the case of cigarette smoking means that the dollar value of the life-expectancy benefit of a single exercise session is USD 25.11.[4] This means that the optimal price intervention is $\tau = -2.51$ USD for $\beta = 0.9$ and as high as $\tau = -10.04$ USD for $\beta = 0.6$. And, as

[3] Two calculators I used could be found in May 2005 at http://moneycentral.msn.com/investor/calcs/n expect/main.asp and http://www.agingresearch.org/calculator/quiz.cfm.

[4] I am grateful to Jonathan Gruber for performing this calculation.

in the case of cigarettes, this ignores all benefits of exercise other than its direct effect on longevity.

A serious issue with interventions of this magnitude is that it is difficult to directly subsidize exercise. While many people will be willing to go to the gym for 10 USD, some will do little more than watch TV once there. That is, the intervention, which can only apply to easily observable behavior, may not be as effective in changing the critical behavior in question.[5]

Several comments about this analysis are in order. One question is why I focus on linear price interventions to achieve the optimal level of consumption. Since all preference parameters are known, this level can be fixed in advance without any choice left up to self 1. A linear intervention, however, remains optimal even if self 1's taste for the good is not known in advance, as long as the benefit or harm of consumption is known. Indeed, in many potential applications of the theory, the momentary pleasure of consuming a good is highly unpredictable, while the future damage changes much less dramatically from person to person, or for any given person over time.

In addition, the above analysis ignores at least two important issues that are likely to affect the optimal price intervention. First, the adverse consequences of a harmful product are often convex in the total consumption accumulated over one's lifetime, and the gains from a beneficial good are often concave. For example, smoking relatively heavily for a short period, and then quitting or cutting back significantly, has no noticeable impact on health and mortality. A moderate amount of exercise is very beneficial for health, but the return to further effort is much smaller and eventually turns negative. In this case, if consumers overestimate their future self-control—such as do the naïve consumers defined below—they will underestimate the benefits of current patience. If a smoker thinks that she will soon quit smoking, she believes that her current cigarette consumption does little harm. And if a person thinks she will exercise regularly in the future, she will not feel the need to work out so hard today. In this case, the optimal price intervention can be much higher than the estimates above, even up to the full future impact of consumption.

[5] To a lesser extent, the same problem applies to smoking. Just as parents in poorer countries tell kids to finish all food on their plates—while lots of food is wasted in richer nations—when cigarette prices increase, consumers smoke each cigarette more intensely (Adda, 2005). Hence, part of the effect of a price change is undone by consumers' behavioral response.

Second, many potential goods to which this framework applies (including smoking and exercise) are addictive. Incorporating addiction into a consumption model introduces a number of subtle behavioral consequences (Becker and Murphy, 1988), and has implications for the optimal intervention (Gruber and Köszegi, 2001; Gruber and Köszegi, 2004). Surprisingly, if consumers understand the addictive nature of a substance, such addictiveness actually decreases the optimal intervention, sometimes substantially so (Gruber and Köszegi, 2001). For a harmful substance, for instance, a consumer will be worried about getting addicted, so she restrains current consumption even by herself. If consumers do not appreciate the addictive nature of a harmful substance (Loewenstein, O'Donoghue, and Rabin, 2003), they underestimate the extent to which initial consumption will lead to future cravings and consequently more damage. This consideration is likely to increase the optimal intervention.

3. The market for self-control

The previous sections have demonstrated that the consumer's behavior in a period-1 spot market is different from what she herself would consider optimal, and that a simple price intervention can perfectly correct her behavior. This does not immediately imply, however, that government interventions in the market are necessary as a result: since the consumer may be able to participate in market transactions already in period 0, she might privately arrange for a price adjustment or other intervention in her future behavior. I now turn to analyzing whether the market would be able to provide this service in various environments.

Three caveats regarding this exercise are important to emphasize. First, the analysis is much more explorative than conclusive. While some general principles will emerge from the analysis, no specific market can easily and without careful individual consideration be categorized into any one of the environments analyzed below. In particular, it is essential to consider whether and how the contracts considered in the theory might be implemented in practice.

Second, investigating the ability of the market to provide self-control is just one of the many questions to ask regarding the necessity or desirability of government intervention. Interventions typically have a variety of costs and other limitations ignored by the model of this paper. On the one hand, enforcing a policy is costly, and even at a

very high cost, it may be impossible to avoid black-market activity that circumvents attempts at regulation. On the other hand, even if a policy can be enforced, it may not be desirable because of the distortions it causes in the behavior of fully rational and self-controlled individuals (Camereret al., 2003). O'Donoghue and Rabin (2005) and Sunstein and Thaler (2003) deal with some of these heterogeneity issues in detail.

Third, there are multiple reasons for the existence of many market features that emerge from the analysis, and it is not the goal of this paper to identify the extent to which each is driven by firms' reactions to self-control problems. Whether or not the original reason for these market features is consumer hyperbolic discounting, however, they do have the effects discussed below on such consumers.

3.1. Further assumptions

In the marketplace, the major way in which self-control can be achieved is through self 0 taking steps that influence her future behavior. In order to carry out an analysis of the market's ability to provide self-control, I must therefore make assumptions regarding what self 0 predicts about self 1's behavior—which is different from what she would prefer in the future. Two extreme assumptions can be made in this regard. At one end, one can assume naïveté, where the decision maker is completely unaware that she will be more impatient in the future than she would now like. A naïve self 0 makes plans to maximize her discounted utility $u_0 + \beta\delta u_1 + \beta\delta^2 u_2$, and takes the period-0 step in the optimal plan. She does not realize that in period 1, she may prefer to change her mind. At the other extreme, one can assume sophistication, where the decisionmaker fully realizes that she will behave suboptimally, and acts accordingly. A sophisticated self 0 calculates her optimal choice using backward induction: she first solves for self 1's behavior in each contingency, and then makes the best decision, taking this behavior into account.

To address questions of how much sophistication or naïveté is required to generate certain outcomes, it will be useful to consider intermediate levels of sophistication, not just the above two extremes. As a natural metric for the degree of sophistication, I use the framework of O'Donoghue and Rabin (2001). Suppose that self 0 believes with certainty that the discount factor between periods one and two will be $\hat{\beta}\delta$; that is, she believes that self 1's utility function is

$u_1 + \hat{\beta}\delta u_2$. Parameter $\hat{\beta}$ reflects self 0's beliefs about her future β, so that $\hat{\beta} = \beta$ corresponds to perfect sophistication, and $\hat{\beta} = 1$ corresponds to complete naïveté. For simplicity (and still following O'Donoghue and Rabin, 2001), I assume self 0 to be completely certain in her beliefs about self 1's preferences. In particular, this means that she does not make inferences about β from the contracts offered to her in the market.

I will consider various market environments and investigate the types of outcomes possible under different degrees of naïveté. As mentioned above, all cases will assume that there is perfect competition between firms, that firms' marginal cost of producing good b is c, and that good a is supplied on a competitive spot market at price 1.[6] The environments differ in the types of agreements consumers and firms are assumed to be able to undertake in period 0.

3.2. Two-part tariffs

The first market environment I consider is one where firms can offer two-part tariffs to consumers in period 0. More precisely, a firm can offer consumers a pair (T, p), where T is an up-front payment to the firm in period 1 (which could be negative), and p is the per-unit price of good b in period 1. Once a consumer accepts an offer, she is bound by its terms and cannot purchase from other firms. This assumption applies to situations where firms and consumers can write exclusive contracts and, more realistically, also to situations where switching costs prevent consumers from going to another firm once they have decided on a producer.[7] The qualitative results of this case are almost identical to those of DellaVigna and Malmendier (2004), although the analysis is adapted to the specific model in this paper.

Self 0 believes that when facing price p, self 1 will choose consumption according to

$$u'(b) = p + \hat{\beta}\delta b. \tag{3}$$

[6] The qualitative features of the contracts would be very similar if the market was instead dominated by a monopolist, although in that case, firm profits would, of course, be positive.

[7] I argue in Section 3.3 that in the case of harmful goods, small switching costs are often sufficient to deter consumers from switching.

As in the case of first-order conditions (1) and (2), the marginal utility of consumption is set equal to the harm-inclusive price of the good. Since self 0 believes that self 1 discounts between periods 1 and 2 using the discount factor $\hat{\beta}\delta$, she believes self 1 will count the harm at $\hat{\beta}\delta h$. Let $b^{\hat{\beta}}(p)$ be the solution to the first-order condition (3). In this notation, $b^{\beta}(p)$ is self 1's actual consumption of good h.

Based on the above actual and perceived behavior, a profit-maximizing firm's choice of p can be shown to satisfy

$$p - c = \underbrace{K \cdot [b^{\beta}(p) - b^{\hat{\beta}}(p)]}_{\text{Exploiting false beliefs}} + \underbrace{L[(1-\hat{\beta}) \cdot \delta h]}_{\text{Demand for self-control}}, \tag{4}$$

where K and L are positive real numbers. Equation (4) gives firms' equilibrium markup over the marginal cost. Since the markup is zero whenever $h = 0$ or $\beta = 1$, its sign and size measure how the market responds to the consumer's self-control problem regarding a good with future consequences.

To understand the intuition behind Equation (4), suppose that $h > 0$, so that the good is harmful. The intuition for beneficial goods is exactly symmetric. If the consumer is perfectly sophisticated—she correctly predicts her period-1 preferences and behavior, so that $\hat{\beta} = \beta$ and $b^{\hat{\beta}}(p) = b^{\beta}(p)$—the first under braced term is zero and in fact $L = 1$, so that the equation reduces to $p - c = (1 - \beta) \cdot \delta h$. Hence, in the contract accepted by consumers in equilibrium, firms impose exactly the optimal price intervention identified in the previous section. Sophisticates understand the exact extent to which they consume too much in period 1. Therefore, they accept the optimal compensated price increase for themselves to decrease their future consumption.

At the other extreme, consider a perfectly naïve consumer. Then, $\hat{\beta} = 1$, so the second under braced term in Equation (4) is zero. Since self 0 believes that she will behave optimally in the future, she believes she does not need to impose any self-control. Therefore, she does not want the price to be above the marginal cost on that account. A naïve consumer, however, underestimates consumption in period 1 $b^{\beta}(p) > b^{\hat{\beta}}(p)$—and a firm can exploit this misperception by raising p and lowering T in a way that seems like a good deal to the

consumer, but in fact is not. Hence, the equilibrium price is again above the marginal cost.

When the consumer is partially sophisticated $\left(\beta < \hat{\beta} < 1\right)$, both the above considerations enter firms' equilibrium markup. To the extent that the consumer believes she will overconsume in period 1, she demands a price increase as a way of exercising self-control. This is captured in the second term in Equation (4). The key determinant in her demand for self-control is $1-\hat{\beta}$, her perceived future self-control problem. To the extent that the consumer is naïve about her future behavior, the firm exploits her misperception by charging a higher price. This is captured in the first term in Equation (4). The key determinant of this term is $b^{\beta}\left(p\right)-b^{\hat{\beta}}\left(p\right)$, self 0's misperception of future behavior.

Note that whatever is the degree of naïveté, for $b > 0$ both terms in Equation (4) are nonnegative, and at least one of them is positive. Hence, the markup is positive for all types of consumers. But this means that the market provides at least some self-control to all consumers. In fact, depending on the exact form of the utility function, the markup might be higher or lower than the optimal price intervention identified in the previous section. For naive consumers, the self-control provided by the market is in a sense an accidental byproduct of firms' attempts to take advantage of consumer naïveté, but the self-control is nevertheless useful.

The central conclusion from this version of the model is that in a competitive market where firms can offer only two-part tariffs and consumers are bound to the firm whose offer they accept, the market does not systematically under provide self-control: it provides perfect self-control to perfect sophisticates, and may provide too little or too much of it to partially or completely naïve individuals. Hence, if government intervention is necessary, the right policy may even be to encourage consumers to be more short-sighted.

3.3. Two-part tariffs with a competitive spot market

The previous section has assumed that firms and consumers can write contracts that govern consumers' market behavior in period 1, or that consumers find it too costly to switch to another firm once they have decided on a producer. Often, however, there is nothing to prevent consumers from also dealing with other firms in the market in period

1. This section considers the equilibrium contracts in such an environment.

As in Section 3.2, suppose that firms and consumers can write two-part linear contracts in period 0 regarding consumption in period 1. Now suppose, however, that in period 1 good b is also available on a competitive spot market—which will therefore have an equilibrium price of c.

The analysis of this alternative model is straightforward. In addition to previous considerations, firms offering contracts in period 0 must take into account that consumers may choose to purchase from the spot market in period 1. This means that the per-unit price they offer must be less than the spot-market price: $p \leq c$.

Suppose first that the good is beneficial, so that $b < 0$. Since the equilibrium markup found in the previous section is negative, adding the non-binding constraint $p \leq c$ does not change the equilibrium. Intuitively, once a consumer pays for having access to a good at a cheaper price, she would be dumb to then buy on a more expensive spot market. Hence, the spot market does not threaten such agreements.

Health clubs seem to be a good example of this kind of outcome. Health clubs typically offer expensive memberships with unlimited use of the gym for a given period. At most gyms, it is also possible to buy a daily pass for single use. Once a gym-goer has a membership at one particular gym, however, she will try to attend her own gym instead of paying the one-time fee at a competitor's.

The story is entirely different for harmful goods, where $b > 0$. In that case, the equilibrium markup found in the previous section is positive, which does not survive competition from the spot market. In fact, given that markups have to be non-positive, the best firms can do is to replicate the spot market and offer a per-unit price of $p = c$. Hence, in the case of harmful goods, the market cannot restrain consumers' overconsumption at all.

This case exposes a major limitation of the market in providing self-control. While a market participant has a strong incentive to provide self-control to consumers who demand this service, other market participants may be able to handsomely profit from undermining such self-control. A meat producer would be happy to sell only expensive hamburgers to a consumer who prefers a high price for self-control reasons. But the numerous hamburger joints around town will still offer lower prices to entice consumers into buying their unhealthy

food. More generally, whatever self-control products (smoking cessation aids, drugs against alcohol abuse, fat farms, etc.) someone might invent, there will be others interested in undermining the use and effectiveness of these products and making their harmful goods ever more tempting to consumers.[8]

The major lesson from the current version of the model is that there is a fundamental asymmetry in the market's ability to provide self-control for harmful goods as opposed to beneficial goods. It is possible to decrease the marginal price of a good through market forces, but pressure from competitors limits a firm's ability to increase its price. Government intervention can, in principle, solve the problem. With the caveats previously discussed, the government has a unique ability to increase the prices applicable to all market transactions—as opposed to just single bilateral price increases that can be achieved by a private contract.

The prediction that the marginal market price of a harmful good cannot increase above the marginal cost largely seems to be correct for many consumer goods (e.g. fat- and sugar-laden foods) that are available at a very low cost. There are, however, potentially harmful products that are puzzlingly costly. For example, credit card borrowing seems very expensive relative to the apparent competitiveness of the industry (Ausubel, 1991). Self-control problems combined with even a small switching cost between products provides one natural explanation for this phenomenon. Suppose a consumer with self-control problems is currently consuming an expensive harmful product, and must pay a small switching cost to find an appropriate alternative. She might have to look and apply for another credit card, for example, or try a few brands of cigarettes to find another that she likes. For different reasons, neither a naïve nor a sophisticated consumer will pay the switching cost. A naïve person underestimates the amount she will consume in the future, and hence does not believe that it is worth making the investment to find an alternative to the current one. A sophisticated consumer, on the other hand, knows that she will consume a great deal, and realizes that she could save money by looking around. But she also knows that finding a cheap alternative will exacerbate her self-control problem by inducing her to con-

[8] It is my impression that the market's incentive to undermine self-control provided by other participants has been recognized by other researchers. I am not aware of any model that formalizes the issue or identifies the distinction between harmful and beneficial goods.

sume more. Because consumers are reluctant to switch to alternative producers, consumers and firms can agree on prices ex ante, so that the analysis of the previous section applies.[9]

3.4. Non-linear contracts

The analysis so far indicates that if competitive spot markets are readily available to consumers at the time of consumption, market forces cannot supply self-control for harmful goods with very low switching costs between producers. For beneficial goods, as well as harmful goods for which there are non-trivial switching costs, however, it seems that the market is effective at providing self-control. The current section shows that this conclusion does not hold when firms are not limited to two-part tariffs in their ex-ante offers to consumers.

Suppose that firms and consumers in period 0 can contract on consumption levels in period 1, instead of just on a transfer and linear price subject to which self 1 can choose consumption. Although it is typically impossible to contract perfectly on consumption, examining this extreme case is very useful for a number of reasons.

First, ex-ante contracting on consumption would appear to be the most favorable situation for a market solution to self-control problems to emerge. I will show that even under this scenario, the market does not do well.

Second, many more complicated contracting possibilities reduce to the analysis below. Most importantly, even if firms offer arbitrary, highly non-linear and non-monotonic price schedules, any offer ultimately reduces to just two relevant consumption-payment pairs: one that self 0 thinks self 1 will choose, and one that self 1 will actually choose. Thus, the analysis of the two-option choice sets below applies.

Third, as argued by Eliaz and Spiegler (2004) and DellaVigna and Malmendier (2004) and as argued below, the results and logic emerging from the analysis do have parallels in actual market institutions.

To develop the results, I first assume that the only contracts allowed in the market are ones that completely specify the consumption

[9] This also means that established producers of harmful products can have substantial market power over consumers. Such market power might itself call for government intervention into the consumption or production side of harmful goods. It is, for instance, likely that market power decreases a firm's incentive to make its product safer.

of good b in period 1. Such contracts take the form (b, T), where b is the agreed-to consumption level and T is a transfer to the firm.

Since there is no choice for self 1 to make, the consumer's perception of period-1 preferences is irrelevant for her evaluation of different contracts. In addition, when choosing the period-1 consumption level to commit to in period 0, a person chooses according to her current preferences. Thus, equilibrium consumption is the optimal level b^* and the accepted contract is (b^*, b^*i). Therefore, it appears that if contracting directly on consumption is possible, the optimal level of consumption will be achieved.

I now show that this conclusion fundamentally depends on the assumption that firms can only offer a committed level of consumption. Suppose instead that firms and consumers can agree to leave the final level of consumption partially up to self 1. In particular, suppose firms can offer decision sets of the form $\{(b_1, T_1), (b_2, T_2)\}$, where self 1 will be able to choose the consumption-transfer pair from this set. Firms can, however, still offer perfectly committed contracts (where $(b_1, T_1) = (b_2, T_2)$). These flexible contracts are similar to those considered by Eliaz and Spiegler (2004), though in their framework, consumers have a different form of naïveté about their preferences, and their focus is not on whether the market can provide self-control, but on how a monopolist should optimally screen individuals with different degrees of naïveté.[10]

First, I consider the (easier) case of a perfectly sophisticated consumer. Whatever choice set such a consumer is offered, she knows the option she will choose in the end. Hence, a contract that offers her a choice is equivalent to the committed contract that only offers the option she would eventually choose. This means that without loss of generality, I can restrict the attention to committed contracts. And I have shown above that when limited to committed contracts, the equilibrium level of consumption is optimal.

The situation is entirely different if individuals are not perfectly sophisticated. First, I argue that offering the committed contract $(b_1, T_1) = (b_2, T_2) = (b^*, b^*i)$ is not an equilibrium in that case. I do so by identifying a profitable deviation from it. Suppose that the good is harmful ($b > 0$), although the argument is essentially the same for

[10] Screening is a non-issue in the current model, since even if $\hat{\beta}$ is not known, consumers self-select into the contracts they would be offered with perfect observability of $\hat{\beta}$.

beneficial goods. Consider a firm that offers the contract $\{(b^*, b^*c - \varepsilon), (b^u, T - \varepsilon)\}$, where b^u is self 1's preferred consumption at price c, and T is chosen so that self 1 is indifferent between $(b^*, b^*c - \varepsilon)$ and (b^u, T). Since $\hat{\beta} > \beta$, this means that self 0 thinks that self 1 strictly prefers $(b^*, b^*c - \varepsilon)$ over (b^u, T), so that for a sufficiently small $\varepsilon > 0$, she also believes self 1 will choose $(b^*, b^*c - \varepsilon)$ over $(b^u, T - \varepsilon)$. Because the option $(b^*, b^*c - \varepsilon)$ entails the same consumption and a lower payment than the committed contract offered in the purported equilibrium, she strictly prefers to accept the flexible contract over the committed one. But by construction, self 1 actually strictly prefers option $(b^u, T - \varepsilon)$ in the contract. Furthermore, since b^u is self 1's preferred consumption level at price c, she is willing to pay more than $(b^u - b^*) \cdot c$ to increase consumption from b^* to b^u: $T - (b^*c - \varepsilon) = (b^u - b^*) \cdot c + k$ for some $k > 0$. This implies that $T = b^u c + k - \varepsilon$, so that for a sufficiently small $\varepsilon > 0$, the firm earns positive profits from the flexible contract.

Intuitively, a firm that offers a flexible contract can "lure" the consumer with an option that dominates the competitive committed contract, and have her pay to change her mind. Loosely speaking, it tells the consumer to take a gift in exchange for accepting an option to later change her mind regarding consumption. Thinking that she will not change her mind, the consumer takes the deal, and knowing that she will change her mind, the firm makes positive profits from the deviation. For example, casinos often offer very cheap hotel rooms to attract consumers who think they are going to gamble little and thus believe this is a great deal. Once vacationers are at the hotel, they gamble more and leave a great deal of money there.

The above considerations imply that for any $\hat{\beta} > \beta$, the equilibrium contract must induce the consumer to unexpectedly change her mind regarding consumption. I now investigate the properties of the two options in the equilibrium contract. Call the consumption level self 0 expects to choose option 1 or the "luring option", and the level self 1 chooses option 2 or the "switching option".

First, it is clear that self 1 must be indifferent between the two options. If she strictly preferred the switching option, a firm could deviate by slightly decreasing the payment in the luring option, and slightly increasing the payment in the switching option. Since self 0 still expects not to switch, but that self 1 will switch, consumers accept the deviating firm's contract, and the firm earns positive profits.

Since self 1 is indifferent between the two options and $\hat{\beta} > \beta$, self 0 expects to strictly prefer to stick with the luring option. As a result, self 0's preferences do not impose a binding constraint on the firm's choice of the switching option. Hence, the only concern in choosing the switching option is to get self 1 to pay most for changing her mind. This consideration has a striking implication: the firm offers self 1 her preferred consumption level. As a result, it ends up providing the consumer zero self-control.

It is worth emphasizing that the above holds for *any* $\hat{\beta} > \beta$—even for an arbitrarily small amount of naïveté, the market provides zero self-control. Consumers with $\hat{\beta} \approx \beta$ know that they will be much more impatient in the future than they would like to be, and, in principle, have a great demand for self-control. Nevertheless, they end up buying none in the market.

Broadly speaking, the profit-maximizing motive that leads a firm to sell self-control to a consumer who demands it also undermines the self-control itself. More precisely, a firm would like to simultaneously sell self-control to consumers, and then offer them an option to buy out of it. They can achieve both these goals by offering a package that appears to provide enough self-control to consumers, but in fact affords none.

Summarizing the results of the current section and those of Section 3.2, for perfect sophisticates, the market offers perfect self-control independently of whether firms are restricted to two-part tariffs or can engage in non-linear pricing. When consumers are not perfectly sophisticated, however, they receive much more self-control when they are facing a two-part tariff. To understand the difference, suppose again that $h > 0$ (i.e. the good is harmful). The key difference is whether a firm can exploit both consumer misperception and consumer impatience. With two-part tariffs, a firm cannot do both: in order to take advantage of naïveté, it needs to raise the per-unit price above the marginal cost; but this limits self 1's consumption to inefficiently low levels (from the firm's point of view). With non-linear prices, the firm can exploit both consumer misperception and consumer impatience. To return to the casino example, once vacationers are at the hotel, the casino offers them expensive gambling, but also free alcohol and other perks that induce people to play, even when they are not so much into it. If gamblers sometimes gamble for the perks but at other times play without fully taking advantage of such

perks, their consumption level may be largely determined by the perk-inclusive price, while their overall spending reflects much higher average prices. Casinos therefore both extract a great deal of money from guests' unexpected taste for gambling and, at the same time, do not limit their gambling very much.

Another illustration of the same idea is the institution of grace periods for credit card purchases (a similar logic applies to "teaser" interest rates). When a consumer makes a credit card purchase, her debt does not immediately begin to accrue interest. If she pays off the balance by the due date on her next bill, she does not pay any finance charges. This means that for loans, a consumer pays back relatively quickly, credit card borrowing is much less expensive than the 16-percent average annual interest rate on interest-paying credit card balances (Laibson et al., 2003). Hence, if consumers perceive that they will quickly pay back the additional debt incurred by a marginal purchase, they choose their debt level according to a marginal price significantly below the average price on inframarginal purchases. Indeed, a partially naïve borrower may believe that she can splurge this month, but cut back and pay back all extra debt next month, so that she does not have to pay interest on the extra purchase. Once again, therefore, she pays heavily overall for borrowing, but this does not deter her to any considerable extent from accumulating a great deal of debt.[11]

More generally, for both harmful and beneficial goods, quantity discounts have the feature of the contracts derived above. For harmful goods, consumers do not mind the high prices they would have to pay for much of their consumption, because they believe they will not consume a great deal. This allows a firm to extract great deal of money from consumers. At the same time, the lower marginal prices for higher levels of consumption mean that consumers get little self-control in the market. For beneficial goods, consumers like high levels of consumption to be relatively cheap on a per-use basis. But since they end up consuming less, they both pay a great deal on average and receive less self-control than would be desirable.

But the model and the results in this section are quite extreme in many ways. In most situations, the firm would not know the value of

[11] Similarly, if a cell phone user thinking about making a less important call falsely believes she is unlikely to use all her allotted minutes, she perceives the price of talking to be very low. Hence, she may, on average, pay a great deal without being restrained from talking too much.

β, so it would be unable to offer the precisely tailored switching contracts I have derived above.[12] In addition, there may be limitations in setting up contracts of this form, especially when a consumer persistently consumes something different from what she had originally planned. Consider, for instance, health clubs, where a practical implementation of a switching contract could include a membership under which the consumer can regularly use the gym, and an option to pay a penalty and switch to a per-use contract. But since the consumer at any time believes that she will exercise a great deal in the future (if she is naïve) or that she needs self-control (if she is sophisticated), she will not be willing to switch to such a contract. On the other hand, the model is also extreme in its assumptions about contracting possibilities. As mentioned above, perfect ex-ante contracting would seem to be exactly the situation where commitment is most likely to be observed. The fact that even in this case, the consumer purchases no self-control indicates a fundamental weakness in the market's ability to provide it.

Furthermore, the logic of the results in this section extends to marketing strategies beyond non-linear pricing. In general, if a firm has multiple tools—not just the linear price it charges—to profit on the consumer's demand for self-control as well as her misperception of her self-control problem, it may in the end do very little to combat that self-control problem. In the case of health clubs, for example, the fancy features many gyms use to attract consumers who think they will exercise a great deal may be effective in getting people to join, but not in motivating them to actually use the gym intensely in ways that improve health.

4. Conclusion

This paper studies the behavior of consumers who systematically more heavily discount the future relative to the present than they would themselves like, and examines the response of profit maximizing firms to such consumers. The results of the paper derive from the interaction of two straightforward and intuitive consequences of an

[12] Presumably, firms would also not know the value of $\hat{\beta}$. This, however, is irrelevant. It is easy to see that if firms offer both the equilibrium contracts above (that for sophisticates and that for all others), consumers self-select into the contract intended for them.

exceedingly strong taste for immediate gratification. On the one hand, if consumers value and demand self-control, a firm can profit from providing it and will want to provide it. For example, weight-loss companies are happy to charge a dieting person thousands of dollars to take her to a moderately nice place, feed her only 900 calories a day, and force her to do cardio exercises all day. On the other hand, other firms, or even the same firm, can make money by taking advantage of the person's short-sightedness, and so will want to undermine any self-control. No matter how hard weight-loss companies and the dieting person try to get her to be healthy while on a weight-loss vacation, when she returns home there will be plenty of others trying to sell her candy, burgers, and other nasty goods, and make these products highly conspicuous and tempting to her. As I have shown, this second consideration severely limits the market ability to provide self-control, especially for harmful goods. As a result, a government or social planner, which is in a unique position to impose rules that apply to all market transactions, may typically be necessary to bring the consumption of harmful and beneficial goods closer to optimal.

This paper does not deal with one aspect of the market response to self-control problems. Namely, firms can, in principle, develop products that make consumption of a particular harmful product impossible or very unpleasant, and the consumption of a beneficial product easier. For example, drugs exist that make drinking alcohol nearly impossible, and that prevent the "high" from taking opium and morphine. Clearly, if consumers demand self-control, pharmaceutical companies have a large incentive to develop such products. If a particular drug enters into wide use, however, other companies have a strong incentive to develop products that temporarily relieve the effects of that drug. Whether or not either of these products can be developed depends to a large part on the available technology. Hence, a purely economic analysis has little to say about which force ultimately wins out.

References

Adda, J. (2005), Prices, cigarette consumption and smoking intensity, Manuscript, University College of London.

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

Ausubel, L.M. (1991), The failure of competition in the credit card market, American Economic Review 81, 50-81.

Bauman, A.E. (2004), Updating the evidence that physical activity is good for health—an epidemiological review 2000-2003, Journal of Science and Medicine in Sport 7, 6-19.

Becker, G.S. and Murphy, K.M. (1988), A theory of rational addiction, Journal of Political Economy 96, 675-700.

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

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.

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

DellaVigna, S. and Malmendier, U. (2004), Contract design and self-control: Theory and evidence, Quarterly Journal of Economics 119, 353-402.

Eliaz, K., and Spiegler, R. (2004), Contracting with diversely naïve agents, Manuscript, Tel Aviv University.

Gruber, J. and Köszegi, B. (2001), Is addiction "rational"? Theory and evidence, Quarterly Journal of Economics 116, 1261-1305.

Gruber, J. and Köszegi, B. (2004), A theory of government regulation of addictive bads: Optimal tax levels and tax incidence for cigarette taxation, Journal of Public Economics 88, 1959-1987.

Kirby, K.N. and Herrnstein, R.J. (1995), Preference reversals due to myopic discounting of delayed reward, Psychological Science 6, 83-89.

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

Laibson, D.I., Repetto, A. and Tobacman, J. (2003), A debt puzzle, In P. Aghion, R. Frydman, J. Stiglitz and M. Woodford (eds.), Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps, Princeton University Press, NJ.

Lee, I.-M., and Skerrett, P.J. (2001), Physical activity and all-cause mortality: What is the dose-response relation?, Medicine & Science in Sports & Exercise 33, S459-S471.

Loewenstein, G., O'Donoghue, T. and Rabin, M. (2003), Projection bias in predicting future utility, Quarterly Journal of Economics 118, 81-123.

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

O'Donoghue, T. and Rabin, M. (2005), Optimal sin taxes, Manuscript, University of California, Berkeley.

Phelps, E.S. and Pollak, R.A. (1968), On second-best national saving and game-equilibrium growth, Review of Economic Studies 35, 185-199.

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

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

Thaler, R. (1981), Some empirical evidence on dynamic inconsistency, Economics Letters 8, 201-207.

Viscusi, W.K. (1993), The value of risks to life and health, Journal of Economic Literature 31, 1912-1946.