Journal of Behavioral Decision Making
J. Behav. Dec. Making, **13**: 233–250 (2000)

The Economics of Immediate Gratification

TED O'DONOGHUE[1]* and MATTHEW RABIN[2]
[1]*Cornell University, USA*
[2]*University of California, Berkeley, USA*

ABSTRACT

People have self-control problems: We pursue immediate gratification in a way that we ourselves do not appreciate in the long run. Only recently have economists considered the behavioral and welfare implications of such time-inconsistent preferences. This paper outlines a simple formal model of self-control problems, applies this model to some specific economic applications, and discusses some general lessons and open questions in the economic analysis of immediate gratification. We emphasize the importance of the timing of the rewards and costs of an activity, as well as a person's awareness of future self-control problems. We identify situations where knowing about self-control problems helps a person *and* situations where it hurts her, and also identify situations where even mild self-control problems can severely damage a person. In the process, we describe specific implications of self-control problems for addiction, incentive theory, and consumer choice and marketing. Copyright © 2000 John Wiley & Sons, Ltd.

KEY WORDS hyperbolic discounting; immediate gratification; procrastination; self control; time inconsistency

Casual observation, introspection, millennia of folk wisdom, and a mass of psychological research all suggest that people have self-control problems: We would 'like' to behave in one manner, but instead 'choose' to behave in another. In particular, we tend to pursue immediate gratification in a way that we ourselves do not appreciate in the long run.

This simple fact of human nature is ignored in the traditional economic model of intertemporal choice. Economists almost always assume that intertemporal preferences are *time-consistent*: A person's preferences over how to trade off well-being between two moments is the same no matter when she is asked. In contrast, a preference for immediate gratification implies time-*in*-consistent preferences. From a long-run perspective a person has one set of preferences (e.g., she prefers not to smoke in the

* Correspondence to: Ted O'Donoghue, Department of Economics, Cornell University, 414 Uris Hall, Ithaca, NY 14853-7601, USA.

Contract grant sponsor: National Science Foundation.
Contract grant number: 9709485.
Contract grant sponsor: Alfred P. Sloan Foundation.
Contract grant sponsor: Russell Sage Foundation.
Contract grant sponsor: MacArthur Foundation.

future), but when the future arrives she has a different set of preferences (she wants to smoke now).[1] To take another example, suppose a person must choose between doing 5 hours of an unpleasant task on 1 April versus $5\frac{1}{2}$ hours on 2 April. Time consistency implies that the person's choice will be the same whether she is asked on 1 February or on 1 April. Most people, however, would choose 1 April if asked on 1 February, but 2 April if asked on 1 April. When looking two months into the future, people don't intrinsically value single-day delays in performing unpleasant tasks, so we prefer the day that involves less work. But we do have a tendency to put off work until tomorrow rather than do it today.

Self-control problems are central to many economic and social phenomena — how much we (under)save, how we (over)eat, whether we get addicted to alcohol, cigarettes, or other drugs, and whether we finish tasks punctually or inefficiently procrastinate. As such, self-control problems merit extended analysis within the framework of formal economics. Following a trickle of research over the years, a recent spurt of interest indicates that economists have begun the task of incorporating self-control problems into economics.[2] In this essay, we outline a formal model of such time-inconsistent preferences and describe its implications in some specific economic applications, then conclude by discussing some general lessons and open questions in such a model.

We begin by describing the economic model of a time-inconsistent preference for immediate gratification and discussing the related issue of whether people are aware of self-control problems. We then apply this model to a very simple (but abstract) environment where a person must complete some task exactly once in a finite set of periods. Our goal is to highlight some general lessons and intuitions that extend to richer environments. We then apply the immediate-gratification model to three economic applications: engaging in addictive activities, designing incentive schemes for people who procrastinate, and consumer choice and marketing decisions.

Two distinctions arise repeatedly in our applications, both of which have important implications. The first concerns the type of decision a person faces, and in particular the timing of the benefits and costs of activities. A preference for immediate gratification leads one to under-indulge in activities which involve immediate costs and delayed rewards (e.g., putting off an unpleasant task), but to over-indulge in activities with immediate rewards and delayed costs (e.g., overeating). Though phenomena such as procrastination and overeating have often been discussed separately, they can usefully be analyzed from the same underlying propensity for immediate gratification.

The second distinction is a psychological one: whether people are *sophisticated*, and foresee that they will have future self-control problems, or are *naive*, and don't foresee these self-control problems. As befits the profession's maintained assumption of rational expectations, most economists assume sophistication when modeling time-inconsistent preferences. Yet because behavioral evidence falls far short of a general endorsement of this assumption, and because it turns out that whether people are aware of their self-control problems has important implications, we spend considerable time discussing how sophistication affects the economic implications of immediate gratification. We emphasize an interesting phenomenon: Although, as intuition might suggest, awareness of one's self-control problems more often than not helps a person to mitigate them, in many situations sophistication actually *exacerbates* self-control problems. Sophistication makes a person aware that she is going to misbehave in the future, and in many situations, if a person knows she is going to misbehave in the future, she may reason that she might as well misbehave now.

---

[1] Several researchers have specifically posited hyperbolic discounting to account for such preferences. See Ainslie (1991, 1992), Ainslie and Herrnstein (1981), Ainslie and Haslam (1992b), Chung and Herrnstein (1967), and Loewenstein and Prelec (1992).

[2] Over the years, a small set of economists has proposed formal, general models of time-inconsistent preferences. See, for instance, Strotz (1956), Phelps and Pollak (1968), Pollak (1968), and Goldman (1979, 1980). More recently, economists have considered simple formal models of such preferences to examine the implications in various environments. See Akerlof (1991), Laibson (1994, 1995, 1997), Fischer (1997), O'Donoghue and Rabin (forthcoming a, b, 1997).

In addition to behavior, it is often of interest to consider the 'welfare costs' of self-control problems — how much worse off a person with self-control problems is than she would have been had she behaved according to her long-run preferences. We are particularly interested in the possibility that even a mild preference for immediate gratification can harm a person severely. Along the lines of the intuition discussed above, we show that in certain situations even mild self-control problems can hurt a person severely if she is naive about them, while in other situations sophistication about self-control problems can hurt a person severely.

We conclude with a discussion of some general themes and open questions.

THE TIME-INCONSISTENT TASTE FOR IMMEDIATE GRATIFICATION

Economists generally assume that at any point in time a person has well-defined preferences representable by a utility function, and that she attempts to maximize her happiness given these preferences. More precisely, suppose u_t represents a person's *instantaneous utility* in period t. Since a person in period t may care about both her present instantaneous utility and her future instantaneous utilities, economists represent her *intertemporal preferences* from the perspective of period t by a function $U^t(u_t, u_{t+1}, \ldots, u_T)$. For instance, each period might represent one day, so that each u_t represents how happy a person is on a given day, and $U^t(\,.\,)$ describes the person's preferences at day t over how happy she is on that and succeeding days.[3] Economists typically assume that people discount exponentially, yielding the intertemporal utility function

$$U^t(u_t, u_{t+1}, \ldots, u_T) = \sum_{\tau=t}^{T} \delta^{\tau} u_{\tau}.$$

This simplification has the crucial implication that intertemporal preferences are *time-consistent*: A person's relative preference for well-being at an earlier date over a later date is the same no matter when she is asked (as long as neither date has passed).

A preference for immediate gratification, however, implies time-*in*-consistent preferences. From a long-run perspective, a person has one set of preferences (she prefers not to indulge in period t), whereas come period t, she has a different set of preferences (she now wants to indulge). Phelps and Pollak (1968) developed an elegant model of time-inconsistent preferences in the context of intergenerational altruism, which Laibson (1994) later used to capture a time-inconsistent taste for immediate gratification within an individual. Recently Laibson (1994, 1995, 1997), Fischer (1997), and O'Donoghue and Rabin (forthcoming a, b, 1997) have employed this model in economic applications. This model slightly modifies exponential discounting:

$$\text{For every } t,\ U^t(u_t, u_{t+1}, \ldots, u_T) = \delta^t u_t + \beta \sum_{\tau=t+1}^{T} \delta^{\tau} u_{\tau}.$$

In this formulation, both δ and β are parameters lying between 0 and 1. The discount factor δ models long-run, time-consistent impatience, whereas β models a person's preference for immediate gratification — how much she favors now over later. If $\beta = 1$, these preferences are simply exponential discounting. But $\beta < 1$ implies that a person has self-control problems. If $\delta = 1$ and $\beta = \frac{1}{2}$, for

[3] This formulation does not contradict emotions such as wishing to get unpleasant activities over with or to delay pleasant events. If people dislike negative anticipation and like positive anticipation, then these sensations can be correctly and usefully conceptualized as part of their instantaneous utilities. In other words, u_t may depend on anticipation of an event at date $t' > t$.

example, then on Monday the person gives equal weight to her well-being on Friday and Saturday, while on Friday she cares twice as much about her well-being on Friday — which is now 'today' — than about her well-being on Saturday — which is now 'tomorrow'.

When a person wishes to behave a certain way in the future, but then behaves some other way when the future arrives, we say that she has a self-control problem. To model the behavior of a person with self-control problems, therefore, we must ask what a person believes about her future behavior. Strotz (1956) and Pollak (1968) carefully lay out two extreme assumptions. On one extreme, a person could be *sophisticated* and correctly predict how she will behave in the future. At the other extreme, a person could be *naive* and believe she will behave in the future exactly as she would currently like herself to behave.

Are people sophisticated or naive? Most economists modeling time-inconsistent preferences assume sophistication.[4] Indeed, sophistication implies that people have 'rational expectations' about future behavior, so it is a natural assumption for economists. The partial validity of the assumption is indicated by the use of such self-commitment devices as alcohol clinics, Christmas clubs, and fat farms. A naive person does not perceive herself to have a self-control problem, and thus would never limit her future options.[5] Despite the existence of some sophistication, however, it does appear that people underestimate the severity of their future self-control problems. For example, people who repeatedly do not have the 'will power' to forgo tempting foods or quit smoking predict that tomorrow they *will* have this will power. Although there is relatively little evidence about people's degree of sophistication, our impression is that there are elements of both sophistication and naivete in the way people anticipate their future behavior, and our approach is to understand the implications of both assumptions.

To examine the implications of a preference for immediate gratification, and the implications of sophistication about this preference, it is useful to compare three idealized types of people: time-consistent people (whom we will call *TC*s), sophisticated time-inconsistent people (*sophisticates*), and naive time-inconsistent people (*naifs*). These three types can be described by distinct decision-making processes:

(1) TCs have time-consistent preferences with exponential discount rate δ. At any point in time, TCs choose today's behavior by determining the optimal lifetime plan given today's preferences. (Since preferences are time-consistent, at any later time they find it optimal to stick to this plan.)
(2) Naifs have time-inconsistent preferences, with $\beta < 1$ and δ the same as that for TCs. At any point in time, naifs choose today's behavior according to today's preferences under the incorrect belief that they will behave in the future according to their current preferences (and, hence, mimic TCs).
(3) Sophisticates have the same time-inconsistent preferences as naifs. In contrast to naifs, however, at any point in time, sophisticates choose today's behavior while correctly predicting how they will behave in the future.

DOING IT NOW OR LATER

Before examining the preference for immediate gratification in concrete economic contexts, we highlight some general lessons by considering a simple (though fairly abstract) environment where a

[4] To our knowledge, only Akerlof (1991) and O'Donoghue and Rabin (forthcoming a, b, 1997) assume naivete.
[5] In addition, the very term 'self-control' implies that people are aware that it may be prudent to control their future selves. For analyses of self control in people, see Ainslie (1974, 1975, 1987, 1992), Ainslie and Haslam (1992a), Funder and Block (1989), Glazer and Weiss (1992), Hoch and Loewenstein (1991), Laibson (1994, 1995, 1997), Schelling (1978, 1984, 1992), Shefrin and Thaler (1992), Thaler (1980), Thaler and Shefrin (1981), and Wertenbroch (1998). Ainslie (1974) explores similar issues with pigeons.

person must do some activity exactly once in some finite set of periods. In any given period, the person chooses only whether to do the activity then, and there are no external commitment devices available to commit future behavior.

O'Donoghue and Rabin (forthcoming a) show that in this environment the implications of self-control problems and the implications of sophistication versus naivete both depend crucially on the type of activity being considered. The activity may be immediately unpleasant to perform, with any benefits delayed to the future, or it may be immediately pleasurable to perform, with any costs delayed to the future. To illustrate the differences between *immediate costs* and *immediate rewards*, we consider two examples, the first of which involves immediate costs.

Example 1: A person usually goes to the movies on Saturdays. The schedule at the local cinema consists of a mediocre movie this week, a good movie next week, a great movie in two weeks, and (best of all) a Johnny Depp movie in three weeks. Unfortunately, the person must complete a report for work within four weeks, and to do so she must skip exactly one of the four Saturday movies. When does she complete the report?

We suppose that the penalty for not getting the report done is so severe that the person feels she 'must' complete it. Even so, the benefits of completing the report are in the future, and do not depend at all on which of the four Saturdays she completes the report. While the reward for doing the report is delayed, the cost of doing the report on a given Saturday is immediate — not seeing the movie shown that day. Suppose the person's valuations of the four movies are the following:

Movie	*Value*
Mediocre	3
Good	5
Great	8
Johnny Depp	13

To compare the three types of people, we will use the parameters $\delta = 1$ and $\beta = \frac{1}{2}$. Since $\delta = 1$, and the benefits of the report do not depend on when it's written, TCs complete the report in the period with the minimum cost. Hence, TCs will do the report on the first Saturday, skipping the mediocre movie.

Naifs, sadly, will procrastinate until the last Saturday, forcing themselves to miss the Depp movie. Why? On the first Saturday, naifs *believe* that if they don't write the report this week, they will be sensible and write it next week. Given $\beta = \frac{1}{2}$, naifs in week 1 prefer incurring a cost of 5 *next week* to a cost of 3 *now*, so they choose to wait. When the second Saturday arrives, however, naifs decide to go see the good movie, *now* believing they will write the report in week 3 and still get to see the Depp movie; they'd rather pay the cost of 8 next week than 5 now. Finally, when the third Saturday arrives, naifs have self-control problems for a third time and see the great movie, forcing themselves to miss the Depp movie.

Sophisticates also procrastinate, but for only one week — they complete the report on the second Saturday. Sophisticates correctly predict that they would have self-control problems on the third Saturday and see the great movie rather than write the report. Anticipating this, in week 2 they *will* write the report, since they correctly view their choice as between skipping the good movie at a cost of 5 and skipping the Depp movie at a cost of 13. Because in week 1 they correctly predict they will not procrastinate in week 2, sophisticates can safely procrastinate in week 1 and see the mediocre movie.

This example illustrates an intuition for how sophistication can mitigate self-control problems. Naifs repeatedly put off the task because they incorrectly predict they will not procrastinate in the future. Sophisticates, in contrast, foresee exactly how much they will procrastinate in the future, and will not procrastinate now when doing so is too costly. They procrastinate some in this example, but not as

much as naifs. Sophistication need not help, however, when rewards are immediate. Consider a similar scenario:

> **Example 2**: Suppose a person has a coupon to see one movie over the next four Saturdays, and her allowance is such that she cannot afford to pay for a movie. The schedule at the local cinema is the same as for the above example — a mediocre movie this week, a good movie next week, a great movie in two weeks, and (best of all) a Johnny Depp movie in three weeks. Which movie does she see?

We assume the same movie valuations as before, and again consider the parameters $\delta = 1$ and $\beta = \frac{1}{2}$. Now, however, there are no costs associated with seeing a movie, while seeing a movie creates an immediate reward.

TCs will wait and see the Depp movie, since it yields the highest reward. Naifs will see merely the great movie. On the first two Saturdays, naifs skip the mediocre and good movies, believing (incorrectly) that they are waiting to see the Depp movie; with $\beta = \frac{1}{2}$ they prefer a reward of 13 in the future to an immediate reward of either 3 or 5. Unfortunately, on the third Saturday, they cannot control themselves, and see the great movie. Because rewards are immediate, the self-control problem leads naifs to do the activity too soon.

But sophisticates have even worse self-control problems in this situation. They see merely the mediocre movie due to an unraveling similar to that of the finitely repeated prisoners' dilemma. Because sophisticates correctly predict week-3 self-control problems, they will realize in week 2 that waiting means seeing merely the great movie the following week, and therefore they will see the good movie in week 2. Similarly, they will see the mediocre movie in week 1 because they correctly predict that waiting means seeing merely the good movie.

This example illustrates an intuition for how sophistication can *exacerbate* self-control problems. Naifs incorrectly predict they will wait for the best movie, but this naive overoptimism helps motivate them to wait. In contrast, sophisticates recognize that they will misbehave in the future, which unfortunately makes waiting less attractive.

The behavior in Examples 1 and 2 is illustrative of more general patterns in this one-activity environment. First, the timing of rewards and costs is very important for the implications of a preference for immediate gratification. When costs are immediate, people with a preference for immediate gratification tend to *procrastinate* — wait when they should do it — while when rewards are immediate they tend to *preproperate* — do it when they should wait.[6] This result conveys the more general intuition that people tend to under-indulge in activities with immediate costs and delayed benefits, and over-indulge in activities with immediate rewards and delayed costs.

Examples 1 and 2 also illustrate that whether sophistication helps overcome self-control problems depends on the environment. In this one-activity environment, for immediate costs, sophistication *mitigates* the tendency to procrastinate, whereas for immediate rewards, sophistication *exacerbates* the tendency to preproperate. Indeed, a more general lesson is implicit in the discussion above. If sophisticated, a person is more 'pessimistic' about her own future conduct than if naive, since she foresees her future self-control problems. Hence, when pessimism about future behavior increases a person's incentive to behave herself now, sophistication mitigates self-control problems; when it decreases a person's incentive to behave herself now, sophistication exacerbates self-control problems.

Our emphasis on the distinction between sophistication's helping to 'overcome' procrastination but exacerbating preproperation disguises a simpler description of the effects of sophistication in a one-activity setting: It always leads a person to do the activity sooner than if she were naive. This follows

[6] The word 'preproperate' derives from the Latin root 'praeproperum', which means 'to do before the proper time'. While we thought we invented the term in our own research, it can in fact be found in sufficiently unabridged dictionaries.

from the fact that if a person is pessimistic about whether she will do the activity at the right time in the future if she doesn't do it today, it becomes more attractive for her to do it today. This is true irrespective of the timing of the rewards and costs of the activity. In this environment, comparing the behavior of naifs and sophisticates to that of TCs — i.e., comparing how people actually behave to how they would like to behave from a long-run perspective — is also straightforward. Naifs exhibit solely the effects of a preference for immediate gratification: For any activity with immediate costs, naifs will procrastinate; and for any activity with immediate rewards, naifs will preproperate. Sophisticates, on the other hand, may complete an onerous task *before* the best time as a form of 'self-control' — doing it now can preempt costly procrastination in the future. Such 'preemptive overcontrol' by sophisticates crops up repeatedly in specific economic applications, often leading sophisticates to behave in ways that seemingly contradict their preference for immediate gratification.

In addition to examining behavior, it is natural to ask about the 'welfare costs' of self-control problems. How much is a person hurt by her propensity to pursue immediate gratification? The question of welfare can be problematic when a person has time-inconsistent preferences: Since the person has different preferences at different times, it is not obvious 'whose' preferences to use for welfare comparisons. If, however, one takes literally that a preference for immediate gratification represents a self-control problem, then a natural perspective is the long-run one, which ignores a preference for immediate gratification when assessing welfare. We adopt this philosophy.[7]

Because TCs, naifs and sophisticates all have the same time-consistent discount factor δ, we are comparing three types who have the same long-run preferences. (Naifs and sophisticates also have the short-run preference for immediate gratification captured by the parameter β.) We can formally represent these long-run preferences by a long-run utility function:

$$U^{\mathrm{LR}}(u_t, u_{t+1}, \ldots, u_T) = \sum_{\tau=t}^{T} \delta^{\tau} u_{\tau}.$$

We are particularly interested in when mild self-control problems can cause severe welfare losses. If a person has mild self-control problems (i.e., β close to one), the welfare loss from any single decision to indulge now is small. Even so, a person can suffer severe welfare losses when self-control problems are compounded. In our simple one-activity environment, when activities have immediate costs, mild self-control problems can cause severe welfare losses if and only if a person is naive. As we saw in Example 1, naifs can repeatedly procrastinate, each time believing they will do the activity next period. With each decision to procrastinate, they incur a small welfare loss, but their total welfare loss is the sum of these increments. Sophistication prevents costly procrastination and severe welfare loss. In Example 1, the long-run utility for TCs is 26 (=5 + 8 + 13), for naifs is 16 (=3 + 5 + 8), and for sophisticates is 24 (=3 + 8 + 13). The welfare cost of the self-control problem is therefore 10 if the person is naive and 2 if the person is sophisticated. In Example 1, we have (for illustrative purposes) assumed that the self-control problem is unreasonably severe. Given the values for movies we have assumed, if β were close to 1, both naifs and sophisticates would behave like TCs. But when costs are immediate, it can be shown more generally that even if β is very close to 1, naifs — and only naifs — can be severely hurt by their self-control problems by making repeated mistakes. This can only happen in situations where (unlike Example 1) each individual mistake costs very little, but where naifs make so many little mistakes that they add up to a big error.

In contrast, for activities with immediate rewards, mild self-control problems can cause severe welfare losses if and only if a person is sophisticated. As we saw in Example 2, sophisticates' pessimism

[7] An alternative philosophy, a natural by-product of economic tradition, is to use Pareto comparisons, where one behavior is better than another if every period-self prefers it. Goldman (1979, 1980) and Laibson (1994, 1995, 1997) use this approach.

can compound their self-control problems: In the end, sophisticates will give in; because they realize this, near the end they will give in; realizing this, they give in a little sooner, etc. For each step of this unraveling, the welfare loss may be small, but the total welfare loss is the sum of multiple steps. Naivete helps motivate a person not to indulge (she always believes she will wait for the best period), and thus prevents severe welfare losses. In Example 2, the person gets to experience the instantaneous utility of seeing a movie on only one Saturday. Hence, the long-run utility for TCs is 13, for naifs is 8, and for sophisticates is 3, and the welfare cost of the self-control problem is 5 if the person is naive and 10 if the person is sophisticated.

The more general intuition is that the welfare cost of self-control problems depends on an interaction between the situation a person is in and whether the person is sophisticated or naive. In situations where believing she will behave herself in the future makes a person perceive it as less costly to indulge today (e.g., activities with immediate costs), being naive about self-control problems can hurt the person severely. On the other hand, in situations where believing she will behave herself in the future makes a person perceive it as more costly to indulge today (e.g., activities with immediate rewards), being sophisticated can hurt the person severely.

ECONOMIC APPLICATIONS

We now explore the implications of a preference for immediate gratification in three specific economic applications: engaging in addictive activities, designing incentive schemes for people who procrastinate, and consumer behavior and marketing strategy. (These examples reflect our own research interests; conspicuously absent from our discussion is what is probably the most important economic application — saving behavior.)

Addiction

A preference for immediate gratification has important implications for the consumption of harmful addictive products, since such consumption involves yielding today to a salient desire that has costs in the future. Recently, economists have proposed models of 'rational addiction' that capture the essence of harmful addictive products by identifying two crucial properties of addictive products (Becker and Murphy, 1988; Becker, Grossman and Murphy, 1991, 1994). First, they involve *negative internalities*: Current consumption negatively affects future well-being.[8] Negative internalities include future health, career, and personal problems caused by consumption, as well as 'tolerance' — future hits yield less and less pleasure. Second, harmful addictive activities involve *habit formation*: Current consumption increases the marginal utility from future consumption (i.e., the future desire for the product). Negative internalities and habit formation combine to create the well-known trap of addiction: As a person consumes more and more of an addictive product, she gets less and less pleasure from this consumption, yet she may continue to consume the product because refraining becomes more and more painful.

These models of rational addiction, however, *a priori* rule out self-control problems, which many observers associate with addictive behavior. O'Donoghue and Rabin (1997) propose a model of self-control problems in addiction, a simplified version of which we present here. Suppose that in any given period a person must choose whether to consume an addictive product, where her choices are to either

[8] We borrow the term 'internalities' from Herrnstein *et al.* (1993), who define an internality to be a 'within-person externality'. The temporal internality we consider is merely one possible type of internality. Since our example assumes a person fully understands how current consumption affects future well-being, we are in fact assuming that the person 'internalizes the internality'; more generally, this need not be the case.

'hit' or 'refrain'. Consider an example where a person's life has three periods: 'youth', 'middle age', and 'old age'. Suppose further that in any given period a person is either currently 'hooked' or 'unhooked', depending on past consumption. Assume the person enters her youth unhooked, but hitting in one period of her life causes her to be hooked in the next period. Hence, she enters her middle age hooked if and only if she hit in her youth, and enters her old age hooked if and only if she hit in her middle age. Finally, suppose that the person's preferences in each of the three periods can be represented by the following instantaneous utilities:

	Utility from hitting	*Utility from refraining*
When unhooked	10	0
When hooked	−8	−25

There are negative internalities here, since taking a hit now causes the person to be hooked in the future, and her instantaneous utility is always lower if she is hooked than if she is unhooked. There is also habit formation here, since the *net* benefit of hitting now is greater if the person is hooked ($(-8) - (-25) = 17$) than if she is unhooked ($10 - 0 = 10$).

As in the previous section, we compare the behavior of the three types of people using the parameters $\delta = 1$ and $\beta = \frac{1}{2}$. TCs will refrain in their youth and in their middle age, but will hit in their old age (when they no longer worry about being hooked in the future); naifs will refrain in their youth but hit in their middle and old age; and sophisticates will hit throughout their lives. (See the Appendix for calculations.)

In this example, naifs indulge in the addictive activity more than TCs. This result turns out to be quite general. A preference for immediate gratification combined with a belief that in the future she will not have such a preference will always lead a person to overconsume addictive products.

The more surprising observation is that sophisticates in this example indulge in the addictive activity more than naifs. This result reflects the role of sophisticates' correct pessimism about future behavior. While naifs in their youth optimistically *and incorrectly* believe they will surely refrain during middle age, sophisticates know they will hit no matter what during middle age. If a good is habit-forming, however, getting hooked now is more costly when a person thinks she will refrain in the future than when she thinks she will indulge in the future. Hence, the very nature of addictive activities implies that the more a person thinks she will misbehave herself in the future, the less costly she perceives current indulgence to be. As discussed in the previous section, this is precisely the type of environment where sophistication can exacerbate self-control problems. Knowing she will eventually get addicted makes a person reason that she might as well get addicted now.

This example of addiction, however, has the crucial feature that the instantaneous utilities are stationary — the utilities from hitting and refraining depend only on whether the person is currently hooked, and not directly on her age. It is presumably more realistic to assume that a person's temptation to hit is higher in youth than old age, reflecting both the young body's physical resilience and forces such as peer pressure. Indeed, the prediction that all three types consume the addictive product in their old age hints at the unrealism of the above example. In fact, modifying the example to capture the idea that consumption may be more attractive in youth has some important implications. Consider the following instantaneous utilities:

	Utility from hitting	*Utility from refraining*
In youth when unhooked	14	0
In middle age when unhooked	10	0
In middle age when hooked	−8	−25
In old age when unhooked	−5	0
In old age when hooked	−23	−25

The utilities during middle age are as before, while we have changed the utilities for both youth and old age. The person has a decreased desire to hit in her old age; in fact, each type will hit in her old age if and only if she is hooked. In addition, we have modified the example so that the person has an increased desire to hit in her youth, which is more than offset by the decreased desire to hit in old age (+4 versus −15), so that a lifetime of addiction is less desirable than in the previous example.

This modification leads to a very different pattern of behavior: TCs and sophisticates now refrain throughout their lives, and naifs now hit throughout their lives. (The calculations are left to the reader.) That TCs now refrain throughout is not surprising, since consuming in old age is less desirable. Of more interest is how modifying the example in this way affects sophisticates and naifs in opposite directions. Sophisticates indulge less than before (they now never hit whereas before they hit always), and naifs indulge more than before (they now hit always whereas before they refrained in their youth).

The behavior in these examples reflects a couple of important points. First, the pessimism intuition described above and in the previous section ignores the possibility that future behavior may depend on today's behavior. In the modified example, sophisticates would like to hit in their youth and then refrain in both middle and old age. However, they correctly predict that if they enter middle age hooked, then they will hit in both middle age and old age, whereas if they enter middle age unhooked, then they will refrain in both middle age and old age. Hence, sophisticates refrain in their youth because it is required to induce good behavior in the future. This result reflects the sophisticates' more general tendency to look for ways to influence future behavior, or to make attempts at 'self-control' — forgoing an otherwise enticing option *solely* to modify future conduct. Indeed, though the numbers in this example don't generate the phenomenon, it is easy to construct addiction examples in which attempts at 'self-control' can in fact lead to preemptive overcontrol of the sort discussed in the previous section: In situations where TCs consume while young and then later quit, sophisticates may never consume because they correctly predict they will not be able to quit. A sufficiently foresighted young person aware of her self-control problem may consume *less* of an addictive good than would be ideal from a long-run perspective.

Our modified addiction example also illustrates an important phenomenon missing from our first addiction example: Even when it is *well worth their while*, naifs never quit once they get hooked. More than failing to quit because the immediate cost is too high, naifs may not quit because they incorrectly believe they will quit next period. This failure to quit addictive activities is strikingly similar, psychologically, behaviorally, and in the magnitude of harm done, to the procrastination example of the previous section, and to procrastination more generally. 'Quitting' is a task with immediate costs and long-run benefits, exactly the type of thing people constantly put off doing under the belief that we will do it tomorrow. Because naifs predict that they will quit even very harmful addictions, they may be quite susceptible to large temptations during youth (or after traumatic experiences such as divorce). In the face of such large temptations, naifs decide it is optimal to indulge now and then quit in the future. But when the future arrives, they decide not to quit. This is precisely what happens to naifs in our modified example — naifs hit in their youth believing they will refrain in their middle age.

When are people susceptible to addictions that cause severe welfare losses due to the compounding of self-control problems? Sophisticates suffer severe addictions when addiction is nearly inevitable, as in our first stationary addiction example. In such cases, even a relatively mild self-control problem can lead to severely harmful addiction because of the type of unraveling discussed in the previous section. But in more realistic environments where eventually a person will no longer be tempted as long as she controls herself now, as in our second example, sophisticates do not suffer harmful addictions. Because we suspect that most addictive products eventually lose their appeal, these results may suggest that 'sophisticated self-control problems' are not a major source of harmful addictions. Even so, we are cautious in interpreting these results too strongly, because sophistication could still lead to costly addictions in one's youth. If a person sees it as inevitable that she will consume in her

20's (and then quit in her 30's), sophistication may lead her to decide she might as well get addicted in her teens.

Naifs, in contrast, can suffer severe welfare losses in situations where they indulge while naively believing they will later quit. Hence, naifs are particularly susceptible to severe addictions in environments such as the youthful examples above where there is a temporarily large temptation to indulge. Naifs indulge every day during college, planning to quit as soon as they graduate — and then indulge every day after graduation for the rest of their lives, always planning to quit the next day.

Our discussion here has, of course, ignored a number of other factors that probably play a role in addiction. For instance, we have assumed throughout that a person has perfect awareness of how current consumption affects future instantaneous utilities. However, it seems likely that people often may not fully understand the habit-forming and negative-internality properties of harmful addictive products. A person may smoke when she is young merely because she doesn't how bad for her, or addictive, smoking is. We have also assumed that the form of a person's myopia is the same at all ages and in all contingencies. Instead, a young person might have an extreme time-consistent myopia, and/or the severity of the preference for immediate gratification might depend on how hooked a person currently is. These and other factors missing from models of immediate gratification along the lines presented here are likely to be major explanations for harmful addictions.

Incentive schemes for agents who procrastinate

The model in the third section illustrates how self-control problems can lead a person to procrastinate in completing an unpleasant task. The fact that people tend to inefficiently procrastinate gives rise to the question of how to provide incentives to combat it. O'Donoghue and Rabin (forthcoming b) examine the design of incentive schemes for procrastinators that provide rewards based on when the person completes some task. Indeed, such *temporal incentive schemes* are a central aspect to organizational design and many types of contracts. People face punishments for delay, both explicit in the form of decreased compensation, and implicit in admonitions from supervisors and decay in reputation.

We analyze a specific case. A single person, the 'principal', hires another person, the 'agent', to perform a single task, where 'wages' provide the sole incentive with which to induce the agent to complete the task. The lessons from this model, however, carry over to more general settings if we interpret the 'principal' as an organization or supervisor, the 'wages' as non-wage compensation and tacit rewards and punishments, and the 'incentive scheme' as a general organizational understanding about the effects of a delay.

Suppose a principal hires an agent to complete some task. The principal faces a *delay cost*: The task is more valuable to her when done sooner rather than later. But the agent faces a *task cost*, which incorporates the difficulty (and unpleasantness) of completing the task, as well as the opportunity costs of forgone outside options. This task cost will typically be stochastic, varying from day to day, and only the agent can observe it. Efficient behavior minimizes the sum of expected delay costs and expected task costs. If the variation in task costs is sufficiently high, efficient behavior calls for the agent to wait when the task-cost realization is high, and to complete the task when the task-cost realization is low. For example, suppose the principal faces a cost of \$1 for each day the task is delayed (with no deadline, but the task must eventually be done). Suppose further that the task requires a single day's work by the agent, and that on any given day the task cost for the agent translated into dollar terms is either \$4 or \$8, with an equal probability of each. In this example, it is efficient for the agent to be selective — wait for the first time he gets the low-cost realization of \$4.[9]

[9] If the agent is selective, then the expected task cost is \$4 and the expected delay cost is \$1 (i.e., $\frac{1}{2} \cdot \$0 + \frac{1}{4} \cdot \$1 + \frac{1}{8} \cdot \$2 + \ldots$), for a total of \$5. The relevant alternative is to complete the task immediately, which involves no delay cost, but an expected task cost of \$6.

Our focus is whether efficient behavior can be induced by temporal incentive schemes, where a temporal incentive scheme is a wage contract in which the wage depends only on *when* the agent completes the task. Importantly, the wage cannot be contingent on the task-cost realization, possibly because the principal cannot observe the task cost. In contrast, on any given day the agent observes the task cost before choosing whether to perform the task. Task costs are incurred immediately, whereas all wages are delayed, which implies that agents with a preference for immediate gratification will tend to procrastinate.

To simplify our discussion, we restrict attention to a comparison of TCs and naifs, and assume $\delta = 1$. TCs will behave efficiently if the punishment for delay is exactly identical to the delay cost incurred by the principal. Intuitively, since TCs do not procrastinate, it is optimal for the incentive scheme to internalize the principal's delay costs, in which case TCs will weigh these delay costs against their task cost exactly as efficiency requires. In our simple example above, since there is a constant delay cost of \$1, TCs will behave efficiently for any temporal incentive scheme with a constant punishment for delay of \$1.

For naifs, in contrast, merely internalizing the principal's delay costs into the incentive scheme is often not sufficient to induce efficient behavior because they have an inefficient tendency to procrastinate. In our example, for instance, a constant punishment for delay of \$1 is not enough to overcome the procrastinatory tendencies of naifs with $\beta = \frac{1}{2}$. The principal can nonetheless induce efficient behavior by increasing the punishment for delay so as to counteract the propensity to procrastinate. In our example, any punishment for delay between \$2 and \$10 will induce efficient behavior by naifs (see the Appendix for calculations). More generally, if people have a preference for immediate gratification, then incentive schemes designed to combat procrastination must punish delay by more than its true cost.

Exactly *how much* extra punishment is required depends critically on the agent's propensity to procrastinate. For instance, suppose we were to modify our example so that the task cost was either \$14 or \$18, again with equal probabilities. In other words, there is a larger mean task cost, but the same distribution around the mean. This change might reflect that the agent finds the task more difficult, or has better outside options. With this modification, it will still be efficient for the agent to be selective, since being selective reduces expected task cost by the same amount as before. In addition, a constant punishment for delay of \$1 will still induce efficiency for TCs. But now naifs are more prone to procrastinate, and inducing efficient behavior now requires a punishment for delay between \$12 and \$20. Thus, when task costs are on average higher, and the agent is therefore more prone to procrastinate, a larger punishment for delay is required.

Up to this point we have assumed that the principal knows the agent's propensity to procrastinate (i.e., she knows β and the distribution of task costs). If instead we suppose she is uncertain about the agent's propensity to procrastinate, we get a qualitatively different result. Suppose some people have a small expected task cost (either \$4 or \$8), while other people have a large expected task cost (either \$14 or \$18). If people do not procrastinate (i.e., everyone is a TC), this additional uncertainty would be irrelevant, since a constant punishment for delay of \$1 would induce efficient behavior for both types of people.

But if people procrastinate, the additional uncertainty changes everything. Because the small-cost agent requires a punishment between \$2 and \$10 to induce efficient behavior, while the large-cost agent requires a punishment between \$12 and \$20, a constant punishment for delay cannot induce efficient behavior for both types. In fact, no temporal incentive scheme of any kind can induce fully efficient behavior for both types. While full efficiency is not possible, the *most* efficient scheme in this example will be a sort of 'deadline' scheme: At first, delay will be punished mildly or not at all; but after some date it will be punished severely. Such incentives encourage those with little propensity to procrastinate to wait until it is efficient to do the task (i.e., not do the task right away if it turns out to be too costly), while assuring that those with greater propensity to procrastinate do not delay too long.

To summarize, then, a preference for immediate gratification implies that people have a tendency to procrastinate on unpleasant tasks. To combat procrastination and induce efficient behavior, incentive schemes must punish delay by more than its true cost. Furthermore, if the principal is uncertain about the agent's propensity to procrastinate, the most efficient incentive schemes will often involve 'deadlines' as a way of combatting the differing degrees of procrastination.

Before leaving this example, there are two technical points that we should mention. First, the discrete 'deadline' result relies on there being discrete types, where a deadline represents incentives adjusting to address the next type. If there is a continuum of unobserved types, the more general result is that the punishment for delay should become more and more severe as time passes.

Second, a principal hiring a naive agent may choose not to induce efficient behavior, but rather to take advantage of the agent's incorrect predictions about his own behavior and thereby 'bilk' him of money. While such 'bilking contracts' are technically possible, we feel there are a number of reasons why such contracts are likely not to be used. Perhaps most importantly, if a principal enters into numerous contracts, it is likely the principal would want to develop a reputation for offering contracts which agents find on average ex post acceptable. For a more complete discussion of these issues, see O'Donoghue and Rabin (forthcoming b).

Consumer choice and marketing

The role of self-control in purchasing decisions is well known among marketing experts. For instance, Wertenbroch (1998) describes how naughty goods are often sold in small packages, since people who know they have a regrettable tendency to overconsume a product tend not to buy it in large packages.[10] More generally, a preference for immediate gratification has important implications for a broad array of day-to-day consumer choices, and therefore also for how firms package and price consumer goods. We know of no extended formal economic analysis in this realm, but we present a couple of examples to illustrate the potential for the general model of immediate gratification to elucidate marketing and consumer-economic issues.

We begin with the example of people buying tempting goods in small packages even when those goods are also available in larger packages at smaller per-unit prices. For example, suppose a person can buy ice cream either in pints or quarts at virtually the same price, so that pints have nearly double the per-unit price. (Indeed, the point of this example will have little to do with price — assume that the price is sufficiently low that it is satiation or fat that determines consumption, not price.) Suppose further that, from a long-run perspective, the person would like to consume one pint each night; however, her preference for immediate gratification means she will eat a quart if she has it at home. Finally, suppose that the person can only go shopping in the morning, and so is out of luck (or in luck, as the case may be) if she gets home and wants to eat more ice cream than she has available.

TCs would optimally consume a pint a day, but to save money, they would buy a quart every other day. Naifs would also buy quarts, but do so every day. Each time they buy a quart, they plan to consume only a pint the first night, but instead they eat the whole quart when they get home. Sophisticates, however, correctly predict that they will eat a quart if they buy it. Therefore, they might buy one pint each day as a means of 'self-control'. In this example, naifs overconsume but purchase at the lower per-unit price, while sophisticates consume optimally but purchase at the high per-unit cost. As a result, firms pricing in this way can take advantage of both sophisticates and naifs. Discounts on large packages stimulate demand from naifs; while small packages provide a commitment device for sophisticates.

[10] For related discussions of self-control in the context of consumer choice, see Thaler (1980), Rook (1987), and Hoch and Loewenstein (1991).

A second example where immediate gratification would seem to clearly play a role is impulse purchases, such as purchasing a candy bar or trashy magazine at the checkout counter in grocery stores. In particular, the potential for impulse purchases driven by a preference for immediate gratification may affect purchase behavior for other goods where immediate gratification would not seem to play a role. For instance, suppose a preference for immediate gratification implies that a person will always make an impulse purchase whenever she goes to the store for other items. Naifs would not foresee this problem, and hence would buy more than they would like. Sophisticates, in contrast, would foresee this problem, and therefore may sometimes, as a means of 'self-control', choose not to go to the store, even when they would benefit from doing so *if* they could control themselves once there. Once again, our point is that a general model of immediate gratification could help flesh out these and other consumer-economic issues.[11]

DISCUSSION AND CONCLUSION

There is a mass of evidence that people are characterized by a preference for immediate gratification and self-control problems. Our goal in this essay has been to outline some of the insights from economic analysis of self-control problems. In this section, we discuss some general themes and open questions.

Perhaps our most explicit theme is that an underlying preference for immediate gratification can explain a variety of behaviors. We have outlined a simple model of a preference for immediate gratification, and have shown how such preferences give rise to procrastination, overindulgence in addictive activities, seemingly excessive punishments for delay in completing a task, and over-consumption of basic consumer goods. There would seem to be a number of other behaviors that such preferences could explain as well.

In addition to the basic preference for immediate gratification, we have also emphasized the importance of whether a person is aware of her self-control problem. For people who are naive about future self-control problems, behavior is simple and intuitive. Their behavior reflects the preference for immediate gratification in the straightforward way: Naifs always procrastinate on unpleasant tasks, always overindulge in addictive activities, etc. For people who are sophisticated about future self-control problems, in contrast, behavior is more complicated. Sophisticates are influenced by the preference for immediate gratification, but they also make attempts at 'self-control' — avoiding otherwise enticing behavior now in an attempt to induce good behavior in the future. In extreme cases, such attempts at self-control can lead to behavior seemingly contrary to that implied by a taste for immediate gratification.

A related theme is whether knowing about future self-control problems makes a person better off. Common intuition suggests that sophistication should help a person overcome self-control problems, and in a number of situations we have shown this to be the case. Nonetheless, we have also described how in a variety of situations sophistication can exacerbate self-control problems. More than just pointing out that sophistication can both help and hurt a person, however, formal analysis enables us to say something about *when* it helps and when it hurts. The answer hinges on whether knowing about future misbehavior increases or decreases a person's perceived cost of current indulgence.

---

[11] While the above examples are usefully analyzed using the simple model of immediate gratification, consumer marketing illustrates some limitations of the simple model. We have focused throughout on self-control problems generated by *temporal* salience — the timing of the rewards and costs associated with a good. Impulsive choice can also be influenced by other forms of salience as well (see Loewenstein, 1996). Different types of goods might induce different levels of desire for immediate gratification. Indeed, firms seem to actively attempt to spark a desire for immediate gratification. Stores tend to have certain goods enticingly packaged and prominently displayed at checkout counters.

An implicit theme throughout our discussion is a classification of situations where a preference for immediate gratification is likely to have significant implications. These situations have something in common: They all involve incremental day-to-day (or moment-to-moment) decisions of how to behave now. A person decides whether to have another cigarette right now, to do a minor task today, or to buy a quart of ice cream to eat soon. Such small-scale day-to-day decisions are where self-control problems are most likely to influence behavior. For long-run decisions, such as how to divide one's retirement savings among investment accounts, a preference for immediate gratification just isn't likely to play a big role (though a person may hurt herself severely by procrastinating for years in making or implementing such decisions). This theme leads us, in turn, to a final observation. As we have shown with examples where mild self-control problems cause severe welfare losses, when making a long sequence of day-to-day decisions, none of which seem important in isolation, even a small bit of a self-control problem can lead a person to behave in ways different from how she would like to behave from a long-run perspective. Even a person who would strike most of us as unusually self-controlled, but who is — after all — merely human, will behave significantly differently than the time-consistent super-human who inhabits the traditional economics model. Hence, if millennia of folk wisdom and a mass of psychological research are even *a little* right, the virtues of formal economic models can be fully exploited only if the taste for immediate gratification is integrated into them.

APPENDIX: CALCULATIONS

Behavior of naifs and sophisticates in addiction example

Consider first TCs with $\delta = 1$. Since the utility from hitting is larger than the utility from refraining — whether or not they are hooked — TCs will hit no matter what in their old age. In their middle age, TCs decide whether to hit knowing that they will hit no matter what in their old age. It is straightforward to show that TCs will refrain no matter what in their middle age (e.g., when hooked in middle age, refraining yields intertemporal utility $(-25) + 10 = -15$, while hitting yields utility $(-8) + (-8) = -16$). In their youth, TCs know they will refrain in their middle age no matter what they do now, and they prefer to refrain (because refraining yields $0 + 0 + 10 = 10$ while hitting yields $10 + (-25) + 10 = -5$). Hence, TCs with $\delta = 1$ will refrain in their youth and middle age, but then hit in their old age.

Consider next naifs with $\delta = 1$ and $\beta = \frac{1}{2}$. In their youth, naifs believe they will behave like TCs in the future, and therefore refrain in middle age and hit in old age, no matter what they do now. The time-inconsistent preference for immediate gratification implies that naifs have an increased desire to hit while young relative to TCs. Even so, naifs with $\beta = \frac{1}{2}$ will refrain while young, because refraining yields $0 + \frac{1}{2}0 + \frac{1}{2}10 = 5$, while they perceive that hitting yields $10 + \frac{1}{2}(-25) + \frac{1}{2}10 = 2.5$. In their middle age, naifs are aware that they will hit no matter what in their old age. Now the self-control problem leads naifs to hit no matter what: Even if they are unhooked, hitting yields $10 + \frac{1}{2}(-8) = 6$, while refraining yields $0 + \frac{1}{2}10 = 5$. Finally, in their old age, naifs, like TCs, will hit no matter what. Hence, naifs with $\delta = 1$ and $\beta = \frac{1}{2}$ will refrain in their youth, but hit in both their middle age and old age.

Finally, consider sophisticates with $\delta = 1$ and $\beta = \frac{1}{2}$. In their middle age, sophisticates correctly perceive that they, like TCs, will hit no matter what in their old age. Given this belief, it is in fact optimal for them to hit no matter what in their middle age (the comparison is identical to that for naifs). In their youth, sophisticates realize that they will hit for the rest of their life no matter what they do now. As a result, it is optimal to hit during their youth as well (because hitting yields $10 + \frac{1}{2}(-8) +$

$\frac{1}{2}(-8) = 2$ while refraining yields $0 + \frac{1}{2}(10) + \frac{1}{2}(-8) = 1$). Hence, sophisticates with $\delta = 1$ and $\beta = \frac{1}{2}$ will hit throughout their lives.

Behavior of naifs in incentive-scheme example

Let X be the constant punishment for delay — that is, if the wage for completing the task today is W, then the wage for completing the task tomorrow is $W - X$. On any given day, naifs believe they will behave like TCs beginning tomorrow. Given the task cost can be either \$4 or \$8 (with equal probability), TCs will be selective (i.e., behave efficiently) for any $X <$ \$2, and TCs will complete the task immediately for any $X \geqslant$ \$2.

Suppose $X <$ \$2, in which case on any given day naifs perceive that they will be selective in all future periods. Naifs therefore perceive the total costs incurred if they delay now to be $\$2X + \$4$, which is the expected lost wages from delaying until the first future low task-cost realization plus the low task-cost realization that is eventually incurred. To decide whether to complete the task now, naifs compare the cost of completing the task now (which is incurred immediately) to the total costs of delay (which are incurred in the future). Naifs will be selective today only if $\$4 \leqslant \beta[\$2X + \$4]$, but for $\beta = \frac{1}{2}$ any $X <$ \$2 yields $\$4 > \beta[\$2X + \$4]$. We can therefore conclude that for any $X <$ \$2 naifs with $\beta = \frac{1}{2}$ will procrastinate forever.

Suppose $X \geqslant$ \$2, in which case on any given day naifs perceive that they will complete the task for sure on the next day. In this case, naifs perceive the total costs incurred if they delay now to be $\$X + \$6$, which is the lost wages from a one-day delay plus the expected task cost incurred tomorrow. To decide whether to complete the task today, naifs again compare the cost of completing the task today (which is incurred immediately) to the total costs of delay (which are incurred in the future). Naifs will be selective today if and only if $\$4 \leqslant \beta[\$X + \$6] \leqslant \$8$, and then $\beta = \frac{1}{2}$ yields $\$2 \leqslant \$X \leqslant \$10$. We can therefore conclude that any punishment for delay between \$2 and \$10 will induce naifs with $\beta = \frac{1}{2}$ to behave efficiently.

For the case where the task cost can be either \$14 or \$18, analogous reasoning yields that the punishment for delay must be between \$12 and \$20. The calculations are omitted.

ACKNOWLEDGEMENTS

For financial support, we thank the National Science Foundation (Award 9709485), O'Donoghue thanks the Alfred P. Sloan Foundation, and Rabin thanks the Russell Sage, MacArthur, and Alfred P. Sloan Foundations. Some of the research for this project was completed when both authors were visiting the Math Center at Northwestern University. We are extremely grateful for their hospitality, intellectual feedback, and financial support. We thank Erik Eyster for valuable research assistance, and Daniel Read and three anonymous referees for helpful comments.

REFERENCES

Ainslie, G. 'Impulse control in pigeons', *Journal of the Experimental Analysis of Behavior*, **21** (1974), 485–489.

Ainslie, G. 'A behavioral theory of impulsiveness and impulse control', *Psychological Bulletin*, **82** (1975), 463–496.

Ainslie, G. 'Self-reported tactics of impulse control', *International Journal of the Addictions*, **22** (1987), 167–179.

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

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

Ainslie, G. and Haslam, N. 'Self-control', in Elster, J. and Loewenstein, G. (eds), *Choice Over Time* (177–209), New York: Russell Sage Foundation, 1992a.

Ainslie, G. and Haslam, N. 'Hyperbolic discounting', in Elster, J. and Loewenstein, G. (eds), *Choice Over Time* (57–92), New York: Russell Sage Foundation, 1992b.

Ainslie, G. and Herrnstein, R. J. 'Preference reversal and delayed reinforcement', *Animal Learning & Behavior*, **9** (1981), 476–482.

Akerlof, G. A. 'Procrastination and obedience', *American Economic Review*, **81** (1991), 1–19.

Becker, G. S., Grossman, M. and Murphy, K. M. 'Rational addiction and the effect of price on consumption', *American Economic Review*, **81** (1991), 237–41.

Becker, G. S., Grossman, M. and Murphy, K. M. 'An empirical analysis of cigarette addiction', *American Economic Review*, **84** (1994), 396–418.

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

Chung, S.-H and Herrnstein, R. J. 'Choice and delay of reinforcement', *Journal of the Experimental Analysis of Behavior*, **10** (1967), 67–74.

Fischer, C. 'Read this paper even later: procrastination with time-inconsistent preferences', mimeo, University of Michigan (1997).

Funder, D. C. and Block, J. 'The role of ego-control, ego-resiliency, and IQ in delay of gratification in adolescence', *Journal of Personality & Social Psychology*, **57** (1989), 1041–1050.

Glazer, J. and Weiss, A. 'Conflicting preferences and voluntary restrictions on choices', mimeo, Tel Aviv University and Boston University (1992).

Goldman, S. M. 'Intertemporally inconsistent preferences and the rate of consumption', *Econometrica*, **47** (1979), 621–626.

Goldman, S. M. 'Consistent plans', *Review of Economic Studies*, **47** (1980), 533–537.

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

Hoch, S. and Loewenstein, G. 'Time-inconsistent preferences and consumer self-control', *Journal of Consumer Research*, **17** (1991), 492–507.

Laibson, D. *Essays in hyperbolic discounting*, PhD dissertation, MIT (1994).

Laibson, D. 'Hyperbolic discount functions, undersaving, and savings policy', mimeo, Harvard University (1995).

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

Loewenstein, G. 'Out of control: visceral influences on behavior', *Organizational Behavior and Human Decision Processes*, **65** (1996), 272–292.

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

O'Donoghue, T. and Rabin, M. 'Doing it now or later', *American Economic Review*, forthcoming a.

O'Donoghue, T. and Rabin, M. 'Incentives for procrastinators', *Quarterly Journal of Economics*, forthcoming b.

O'Donoghue, T. and Rabin, M. 'Addiction and self control', mimeo, Cornell University and UC Berkeley (1997).

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

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

Rook, D. W. 'The buying impulse', *Journal of Consumer Research*, **14** (1987), 189–199.

Schelling, T. C. 'Egonomics, or the art of self-management', *American Economic Review*, **68** (1978), 290–294.

Schelling, T. C. 'Self-command in practice, in policy, and in a theory of rational choice', *American Economic Review*, **74** (1984), 1–11.

Schelling, T. C. 'Self-command: a new discipline', in Elster, J. and Loewenstein, G. (eds), *Choice Over Time* (167–176), New York: Russell Sage Foundation, 1992.

Shefrin, H. M. and Thaler, R. H. 'Mental accounting, saving, and self-control', in Elster, J. and Loewenstein, G. (eds), *Choice Over Time* (287–329), New York: Russell Sage Foundation, 1992.

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

Thaler, R. H. 'Toward a positive theory of consumer choice', *Journal of Economic Behavior and Organization*, **1** (1980), 39–60.

Thaler, R. H. and Shefrin, H. M. 'An economic theory of self-control', *Journal of Political Economy*, **89** (1981), 392–406.
Wertenbroch, K. 'Consumption self-control via purchase quantity rationing', mimeo, Yale University (1998).

Authors' biographies:
Ted O'Donoghue is Assistant Professor of Economics at Cornell University. His research interests include time-inconsistent preferences and self-control problems and the design of patent policy for cumulative innovation.

Matthew Rabin is Associate Professor of Economics at University of California at Berkeley. His personality traits include time-inconsistent preferences and self-control problems. His research interests include peak-load pricing during hyper-inflation.

Authors' addresses:
Ted O'Donoghue, Department of Economics, Cornell University, 414 Uris Hall, Ithaca, NY 14853-7601, USA.

Matthew Rabin, Department of Economics, 549 Evans Hall #3880, University of California, Berkeley, CA 94720-3880, USA.

