INCENTIVES AND SELF CONTROL[1]

Ted O'Donoghue
and
Matthew Rabin

We investigate the design of incentives for people subject to self-control problems in the form of a time-inconsistent taste for immediate gratification. Because such present-biased people may not behave in their own long-run best interests, there is scope for firms, policymakers, friends and family, and the people themselves to create incentives for "better" behavior. Moreover, whereas for standard agents only ultimate well-being from different courses of action matter, for present-biased people the detailed structure of incentives becomes of critical importance. Optimal incentive design, therefore, will attend to details that the conventional model would say are essentially irrelevant. We describe some general principles that have emerged in recent and ongoing research on incentives, highlighting the importance of heterogeneity among agents and providing for flexibility, and illustrate these principles with some simple examples.

Keywords: Hyperbolic Discounting, Present Bias, Time-Inconsistent Preferences, Sophistication, Naivete, Incentive Design

---

[1] For helpful comments, we thank Torsten Persson, Paige Skiba, and Ariel Rubinstein. For financial support, we thank the National Science Foundation (grants SES-0214043, and SES-0214147).

"The problem sets should have been graded. I had no incentive to do them, and as a result did poorly on the exams."

— comment from anonymous teacher evaluation, undergraduate game theory course.

1. Introduction

Ever since receiving the above comment, one of us (Rabin) has presented it to students on the first day of class as the first exercise for his microeconomics courses: In the above situation, find the incentive! This exercise is a good — albeit somewhat easy — introductory economics question. If the student wanted to do well on the exams, and if doing the problem sets would have helped him to do well on the exams, then he should have had incentive enough to do the problem sets. For this student, however, wanting to do well on the exams was apparently not sufficient motivation to prepare for the exams by doing the problem sets. Indeed, this student believed that, had an additional constraint been added in the form of graded problem sets, he would have been better off.

While it is possible that the student just didn't comprehend the basics of incentives, we believe that this comment reflects two related features of human nature that are missing from the conventional economic approach to incentives. First, people have self-control problems in the form of a *present bias*: our short-term inclinations of what to do — watching TV rather than studying — often don't accord with our own assessment of what is in our long-term best interests — studying rather than watching TV. The point is *not* that watching TV is necessarily a worse way to spend one's time than studying game theory.[2] Rather, the point is that there is a meaningful sense in which people may wish they would study rather than watch TV, and yet they watch TV. Although pursuit of immediate gratification over future gratification is often the sensible thing to do, there are situations in which people by their own assessment tend to over-pursue immediate gratification.

Second, the comment also reflects that incentives cannot be understood solely by the total payoffs associated with different actions. The details matter. Two incentive schemes A and B may each yield the same aggregate payoffs associated with two courses of action x and y, yet, depending on

[2] That assessment would depend on the TV show being watched and the game-theoretic solution concept being studied.

the details, a person may choose x under scheme A but y under scheme B. Students may study for 100 hours and get a good grade rather than study for 50 hours and get a bad grade if induced by problem sets, but study the 50 hours for the bad grade rather than the 100 hours for a good grade if no problem sets are graded — even if all they care about is their total hours of effort and their grades.

In this paper, we investigate the design of incentives when people have present bias. Our main themes correspond to the two features above. First, because present-biased people may not behave in their own long-run best interests, there is scope for firms, policymakers, friends and family, and perhaps even the people themselves to create incentives for "better" behavior.[3] Second, whereas the standard economic model says that only ultimate payoffs matter, for present-biased people, the detailed structure of incentives becomes of critical importance. Hence, optimal incentive design must pay attention to — and can sometimes efficiently use — details that the standard model would say are essentially irrelevant.

Casual evidence suggests that people vary a great deal in their degree of self-control problems, their degree of awareness of self-control problems, and (of course) in their intrinsic tastes for different activities. This heterogeneity of course complicates incentive design, because incentives that help some types — e.g., people with significant present bias — often harm other types — e.g., people with little or no present bias. Moreover, even individual agents themselves will typically face uncertainty about their own future tastes, needs, outside options, and constraints. If so, then optimal incentives may need to permit some flexibility to agents, and may even need to deal with heterogeneity in the need for flexibility. Hence, a third major theme will be the importance of heterogeneity and flexibility in designing incentives.

In Section 2, we describe the simple model of (what we call) present bias that was adopted in the seminal work of Laibson (1997) and has been widely used by many researchers ever since. We then discuss a variety of reasons why we believe the study of present bias is important for economics, and also describe recent empirical papers that conclude that present bias may provide better explanations of many observed economic behaviors than does the standard model.

---

[3] There is also scope for firms or opportunistic individuals to design incentives to exploit people's present bias. While our main focus will be beneficial incentives, we occasionally discuss exploitative incentives as well.

In Sections 3 and 4, we move to our main topic: the design of incentives when (some) people have present bias. The optimal design of incentives will, of course, depend on the details of the particular environment under consideration. Even so, existing research on how present-biased people react to natural incentives, as well as emerging research on designing incentives in specific contexts, suggests some general principles. In Section 3, we describe these general principles, and then analyze a series of examples that illustrate these principles. In Section 4, we address more directly the roles of heterogeneity and flexibility. We describe how recent research has started to deal with heterogeneity and the need for flexibility, suggest some principles that have been missing from this literature, and illustrate these principles by discussing some examples of how to combat procrastination.

In Section 5, we discuss some broader issues and conclude. In particular, while our analysis focuses on the effects of incorporating present bias into the standard economic model, we discuss how similar conclusions might arise from other recent attempts to improve the standard economic model. We also discuss how market forces might influence incentive design when some consumers have present bias.

2. Present-Biased Preferences and Economics

Most economic models, intuitions, and policy prescriptions assume that people's short-run preferences match our long-run preferences. More precisely, most economics assumes that preferences are time-consistent: As long as no new information is revealed, people have the same preferences now over future behavior as we will have when the future arrives. Hence, economists have traditionally assumed that if today we prefer to start spending less and saving more tomorrow, tomorrow we'll want to start spending less and saving more immediately; if today we prefer to quit smoking tomorrow, then tomorrow we'll want to quit smoking immediately; if today we prefer to diet tomorrow, tomorrow we'll want to diet immediately; and if today we prefer to work on our dissertation, prepare our taxes, and clean house tomorrow, tomorrow we'll want to do these things immediately.

Evidence suggests, however, that such assumptions are wrong, and more than just being wrong, they are wrong in a systematic direction: People have a bias towards immediate gratification. When thinking about two future dates, people care roughly equally about well-being on those two dates, but when the future arrives and the first of these dates becomes "today", we care more about that

first date than about the second date. We buy more, smoke more, eat more, and work less tomorrow than we wish our tomorrow selves to do.[4]

Although the stuff of literature, philosophy, and everyday folk wisdom, economists have mostly ignored this general human tendency. There have over the years been a few exceptions. Strotz (1956) provided the first formal analysis of time-inconsistent preferences, and suggested that the relevant form of time inconsistency was a preference for immediate gratification. But the literature that built upon Strotz chose to focus on general theoretical issues that arise with time-inconsistent preferences — most notably, the existence of a "sophisticated path". With the exception of a few researchers (e.g., Thomas Schelling, Richard Thaler, and George Akerlof), Strotz's suggestion to focus on a preference for immediate gratification was lost.[5] Fortunately, it was given a rebirth due to the neo-seminal work of Laibson (1997).

In order to formally analyze the implications of present bias, Laibson adopted a simple model that Phelps and Pollak (1968) used to study intergenerational altruism to instead study intrapersonal intertemporal choice. Specifically, a person's intertemporal preferences at time t are given by

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

where u_τ is her instantaneous utility in period τ. This two-parameter model is a simple modification of the standard exponential-discounting model. The parameter δ represents standard time-consistent impatience; for $\beta = 1$ these preferences reduce to exponential discounting. The parameter β represents a time-inconsistent preference for immediate gratification, where $\beta < 1$ implies

---

[4] For evidence that most humans do, indeed, have such a tendency, see Ainslie (1991, 1992), Ainslie and Haslam (1992*a*, 1992*b*), Loewenstein and Prelec (1992), Thaler (1981), and Thaler and Loewenstein (1989). For a recent overview, see Frederick, Loewenstein, and O'Donoghue (2002). This tendency is often referred to as "hyperbolic discounting" or "quasi-hyperbolic discounting".

[5] See, for instance, Schelling (1978, 1984), Thaler and Shefrin (1981), and Akerlof (1991).

an extra bias for now over the future.[6]

The general feature of present bias is of course more general than the simple β, δ model. Indeed, psychologists who estimate discount functions often focus on a one-parameter hyperbolic functional form wherein a person puts weight $a/(a+d)$ on payoffs that are received with delay d — hence the label "hyperbolic discounting". We suspect that most qualitative and quantitative predictions of the β, δ model will hold for more nuanced forms of present bias, but this remains to be seen.

As discussed by Strotz (1956) and Pollak (1968), the behavior of people with time-inconsistent preferences depends on their beliefs about their own future behavior. Two extreme assumptions have appeared in the literature: *Sophisticated* people are fully aware of their future self-control problems and therefore correctly predict how their future selves will behave, and *naive* people are fully *un*aware of their future self-control problems and therefore believe their future selves will behave exactly as they currently would like them to behave.

Much of the research on present bias has followed Laibson in assuming sophistication — perhaps reflecting economists' natural inclination to assume people have "rational expectations". In our own work (beginning with O'Donoghue and Rabin (1999*a*)), we have emphasized both sophistication and naivete, in part to emphasize that many observed behaviors may be better understood in terms of naivete. Indeed, to demonstrate that many behaviors we believe reflect naivete need not reflect 100% naivete, we have developed and analyzed a model of *partial naivete* wherein a person is aware that she has future self-control problems but she underestimates their magnitude

---

[6] This model is a generalization of standard exponential discounting — that is, exponential discounting is a special case that imposes the restriction that $\beta = 1$. Moreover, virtually every paper that applies this model, whether theoretical or empirical, makes explicit how the implications for $\beta < 1$ differ from the implications for $\beta = 1$ in the particular context being studied. Hence, unlike the vast majority of papers in economics that *a priori* impose exponential discounting, this literature permits a careful comparison of the relative merits of exponential discounting vs. incorporating present bias. Of course, permitting $\beta < 1$ is not the only direction for generalizing economic models, and we, like many researchers studying this model, are quite inclined to explore other generalizations. Nonetheless, the model as it stands is clearly less narrow and specialized than the model of time preference it is replacing.

(see O'Donoghue and Rabin 2001).[7] In this paper, we will focus for simplicity on the two extreme types. But many intuitions that arise under both complete naivete and complete sophistication also apply under partial naivete.

While there is little formal evidence on the distribution of types, casual evidence certainly suggests that not everyone has present bias, and that people with present bias differ in their degree of sophistication. We suspect, for instance, that there are many activities — e.g., drinking alcohol — which many people engage in optimally, but which non-trivial numbers of people over-do because of present bias. We emphasize, however, that the heterogeneity in observed behavior, and the harm from this behavior, may be due even more to heterogeneity in underlying tastes for different activities than to differences in present bias or awareness of present bias, and that the implications of present bias depend on these underlying tastes. If you don't like smoking or drinking or eating potato chips, present bias does not induce you to over-indulge in consuming these things. Although many papers that explore the implications of present bias assume for simplicity that there is no heterogeneity (or more to the point, explore behavior as a function of type), a more complete study of incentive design must account for heterogeneity. We address this issue in Section 4.

A major theme in the literature on present bias is that present bias can generate inefficient behavior — that is, can lead people not to behave in their own best interests. This theme will play a particularly prominent role in our discussion on designing incentives. Of course, such a claim requires a criterion for how to measure welfare. Economists have sometimes been perplexed about how to conduct welfare analysis when people have time-inconsistent preferences, since the person has different preferences at different points in time. One response has been to apply a Pareto criterion, where one outcome is deemed better than another if and only if the person views it as better at all points in time. But we view the Pareto criterion as too agnostic, and have instead encouraged using people's "long-run preferences" to measure welfare, where long-run preferences are given

---

[7] Formally, we let $\hat{\beta}$ be a person's beliefs about her future self-control problems, by which we mean that the person believes that in the future she will behave like a sophisticated person with self-control problem $\hat{\beta}$. With this formulation, a fully sophisticated person has perceptions $\hat{\beta} = \beta$ because she knows exactly her future self-control problems; a fully naive person has perceptions $\hat{\beta} = 1$ because she believes she will not have future self-control problems; and a partially naive person has perceptions $\hat{\beta} \in (\beta, 1)$.

by

$$U^0(u_t, ..., u_T) \equiv \sum_{\tau=t}^{T} \delta^{\tau-t} u_\tau.$$

A person's long-run preferences reflect her preferences when asked from a prior perspective when she has no option to indulge immediate gratification — hence why U^0 does not depend on β. We prefer this welfare criterion for several reasons. First, the Pareto criterion is clearly too agnostic about when one outcome is better than another. Indeed, it is not even close to the welfare criterion economists apply when studying standard exponential discounters. It does not, for instance, allow us to say that providing someone (or someone providing himself) with a large immediate benefit at a very small future cost is necessarily a good thing; we have, after all, made the current self better off at the cost of his future self. Second, we view the present bias as an "error". Consider, for instance, the person's present bias that applies on March 1, 2008. On every other day of her life, the person is upset with this present bias, and it is only at that moment that she is at peace with it. Hence, we believe that her present bias on March 1, 2008, or on any other day, should be given no normative weight. Third, if we were to consider a person who is choosing future incentives for herself, she would make this choice using exactly her long-run preferences. In fact, we will discuss many such examples below.

Hence, whenever we discuss welfare below, we will be using long-run preferences as our welfare measure, and this perspective is starting to take hold in the literature. Even so, it is worth making one final comment: most instances of "bad behavior" that we discuss below and in the broader literature are bad by essentially any definition. As a practical matter, especially in applied work, conceptual difficulties over how to measure welfare within the present-bias framework often simply do not matter.

A great deal of research has built off of Laibson's initial model of (what we call) present-biased preferences. Much of the initial work has been theoretical in nature, investigating the implications of present bias both for behavior and for welfare in a variety of economic contexts.[8] More recently, a nascent empirical research program has begun to test present-biased preferences vs. exponential discounting in economic field data (as we discuss below).

There are several reasons why present bias is important for economics. Perhaps the simplest

---

[8] See, for instance, Laibson (1997, 1998), O'Donoghue and Rabin (1999*a*, 1999*b*, 2001), Carrillo and Mariotti (2000), Benabou and Tirole (2002), and Gruber and Koszegi (2001, 2004).

and most straightforward has been, we feel, insufficiently emphasized: without it, economists have no coherent model of short-term impatience. Most economists wouldn't bat an eye if someone suggested that people care 1% more about today than tomorrow — it's just discounting. But under the standard model of exponential discounting, such seemingly reasonable short-term impatience implies manifestly unrealistic long-term impatience. Specifically, because $.99^{365} \approx .0255$, such discounting implies that people care roughly 40 times more about now than a year from now — which none of us do. Without assuming transparently counterfactual and extreme long-term discounting, the exponential model is inconsistent with any noticeable taste for immediate gratification. Even an unrealistically low yearly discount factor of 0.7 (which says, counterfactually, that people care 6 times more about their well-being 10 years from now than 15 years from now) implies that people care only 0.1% more about today vs. tomorrow. Hence, not only does exponential discounting imply that people's current preferences over future behavior accord with their future preferences, but for any plausible discounting it further implies our desired behavior tomorrow is virtually identical to our desired behavior today. In practical terms, then, exponential discounting implies that, if today we prefer to quit smoking, diet, and study tomorrow, then we prefer to quit smoking, diet, and study *today*.

There are two related implications of this logic. First, the strong emphasis in the literature on proving present bias by finding time inconsistency via preference reversals has been misplaced. Finding short-term impatience is sufficient. Second, present bias is *not* a theory of long-term impatience. Economists sometimes seem to associate present bias with extreme myopia. This view is incorrect. Indeed, an important feature of the present-bias model is that it permits the study of immediate gratification without assuming insane myopia. It allows, for instance, that people might care 5% more about today than tomorrow, 10% more about today than a year from now, and 5% more about a year from now than two years from now — which holds in the β, δ model above if $\beta = .95$ and $\delta^{365} = .95$. This distinction is, in fact, a major theme in recent empirical research. We often observe the same people exhibiting long-run planning (committing to retirement plans or to attend college) *and* short-term indulgence (procrastination, credit-card borrowing). Whereas exponential discounting has a hard time accounting for this combination, present bias provides a natural explanation.

A second — and more important — reason why present bias is so important for economics is that many real-world economic behaviors seem to clearly indicate short-term impatience. Topics now

or soon to be studied with the perspective that present-biased preferences matter include: general savings rates, credit-card borrowing, cigarette consumption, welfare enrollment, procrastination in personal investment, unemployment and procrastination in employment search, purchase quantities under non-linear pricing for "virtue" vs. "vice" products, the effects of payday timing on monthly consumption patterns, food stamps, the effects of coupons and rebates on demand for products, organizational incentives, retirement timing, finishing school, returning to school for G.E.Ds, unsafe sex and AIDS and pregnancy, procrastination in seeking medical attention, compliance with medical prescriptions, alcohol consumption, unhealthy eating, exercise, obesity, procrastination in research, and seeing too few Johnny Depp movies.

Is all pursuit of immediate gratification in these domains "irrational"? Of course not. As all economists understand, rational behavior often involves pursuing immediate gratification at future expense. If chocolate or a cigarette brings you immediate gratification, it may be optimal to consume these items despite possible future health costs; if watching TV brings you immediate gratification, it may be optimal to watch despite possible future costs of poor grades; and if having unprotected sex brings you immediate gratification, it may be optimal to do so despite possible future costs to health or life or family. None of these activities should be deemed mistakes *per se*. The point of this paper and of the broader literature is that it is worth investigating whether, in some cases, people engage in these activities against their *own* long-run preferences due to an *over*-pursuit of immediate gratification.[9]

Indeed, in many situations, we feel the evidence seems to clearly implicate self-harm. Consider, for instance, a major facet of U.S. economic activity: credit-card borrowing.[10] There are roughly 1.2

---

[9] Of course, many of these activities may also generate negative externalities on neighbors and taxpayers. While such externalities are and should be the topic of economic research, this research is to some extent orthogonal to research on the self-harm caused by over-pursuit of immediate gratification.

[10] We discuss credit-card borrowing and "fringe borrowing" in the context of U.S. consumers because we are more familiar with the data for the U.S. and because the amount and potential maladaptation of consumer debt seems to be greatest in the U.S. But problems with consumer debt have recently been growing in other countries as well — indeed, this international variation is itself a worthy topic for study.

billion credit cards in use in the United States, with an average annual interest rate of 18.9%. About 60% of active credit-card accounts are not paid off monthly. The average credit-card debt among all American households is $8,400, with an average card debt of $9,205 among people who have at least one card. A typical American family today pays about $1200 annually in credit-card interest.[11] Although further research is needed, it is hard to see how the lower "steady-state" consumption *and* lower financial security that result from common credit-card behavior can be optimal. Hence, while economists often debate whether Americans save enough for retirement, we feel an equally important question is whether Americans' steady-state consumption throughout their lives is too low due to the credit-card debts that they carry.

In fact, a smaller — but still worryingly large — proportion of Americans borrow with instruments that are far more costly than credit cards. This "fringe" set of financial institutions is even less studied by economists than the credit-card industry. There is a massive industry of pawn shops — putting household items into store as collateral for extremely high-interest loans — in the U.S. and in many other countries (this institution dates back millennia, and is described by some as the second oldest profession in the world). And many Americans buy household items such as furniture on a "rent-to-own" plan: instead of paying money to buy a couch, people rent it on a per-month basis, and after a while are given ownership, typically paying three times the price they could have paid initially.

And most dramatically in terms of recent growth, massive numbers of employed (and pensioned) people in U.S. have taken out one or more payday loans — borrowing against their next pay check. A typical transaction is that a person borrows $300 using his or her paycheck or social-security check as collateral, and pays back $354 when this check arrives. This is 18% interest for a month or (more commonly) *two weeks*. The size of this rapidly growing industry is hard to firmly establish. Stephens Incorporated Investment Bankers (2004) estimate that 9 to 14 million U.S. households — 8.5% of households — took out payday loans in 2002, from 22,000 payday-loan outlets. In a *New York Times* article, Henriques (2004) reports that 26% of military households took out such loans in the same year. The Stephens Incorporated report estimates that the total volume of payday loans in 2003 in the U.S. was $40 billion. It appears that most who take out such payday loans do so

[11] These particular estimates are from Bannister (2004), but the general scale of the estimates broadly match other sources.

repeatedly, so that it doesn't appear to be solely as a source for one-time emergencies.

Most of this behavior seems to us hard to reconcile with full rationality as economists model it. People should buffer themselves with some savings to avoid all this borrowing. No matter how bad their situation, to repeatedly borrow at exorbitant rates puts people in worse situations. And the logic of present bias can readily predict that those facing both higher immediate temptation to consume and higher cost of doing so are particularly prone to overconsume. Because the very poor are more often in this situation than the not very poor, they are more prone to damaging overconsumption even if (as we believe probable) they do not have an intrinsically higher tendency to overconsume.

Beyond such examples that compel speculation about present bias, over the past decade a number of researchers have indeed demonstrated empirically ways in which present bias can explain economic field data better than exponential discounting. Laibson (1997, 1998) demonstrates theoretically that hyperbolic discounting could account for many well-known empirical anomalies in the saving-consumption literature — that is, it could account for many empirical findings that were difficult to understand when viewed with the maintained hypothesis of exponential discounting. For instance, whereas evidence suggests that there is too much income-consumption comovement in household consumption data to be consistent with the standard life-cycle/permanent-income model, Laibson demonstrates that hyperbolic discounting can generate such comovement because hyperbolic discounters often find themselves liquidity constrained. Similarly, whereas evidence suggests a large drop in consumption at retirement that is inconsistent with exponential discounting, hyperbolic discounting can generate such a drop.

In follow-up work, Laibson and his collaborators have conducted more explicit empirical tests of hyperbolic discounting vs. exponential discounting in saving-consumption data. Laibson, Repetto, and Tobacman (1998) and Angeletos, Laibson, Repetto, Tobacman, and Weinberg (2001) simulate a buffer-stock saving-consumption model for both an exponential economy and a hyperbolic economy. Importantly, they incorporate in the model the existence of an illiquid asset and the existence of credit-card borrowing, two features that are important for hyperbolic discounters. They calibrate the model to match the distribution of pre-retirement wealth across households in their observed population, and then compare their simulated data to actual data on other dimensions. They find that the hyperbolic economy performs significantly better at matching the data, particularly on the dimensions of amount of wealth held in liquid assets, amount of credit-card borrowing, and income-

consumption comovement. More recently, Laibson, Repetto, and Tobacman (2005) have used the same data to estimate the discounting parameters. In their benchmark specification, they find a yearly $\beta \approx 0.7$ and a yearly $\delta \approx 0.96$, with both tightly estimated.

The key pattern in the data that drives these conclusions is that the same households that accumulate significant wealth by their pre-retirement years also accumulate a lot of credit-card debt when they are young. The fact that households accumulate significant pre-retirement wealth suggests a relatively low long-term discount rate, while the fact that they accumulate significant credit-card debts suggests a relatively high short-term discount rate. Because exponential discounting imposes the same discount rate for both the long term and the short term, it has a hard time generating the observed pattern. In contrast, present bias involves exactly this feature of high short-term and low long-term discount rates.

Recently, researchers have tested for hyperbolic discounting in other environments. DellaVigna and Paserman (2005) test for hyperbolic discounting in job-search decisions. They first develop a model in which unemployed people face two types of decisions: (i) they must decide each day how intensely to search for a job, and (ii) when they search and receive a wage offer, they must decide whether to accept that wage. In this framework, impatience (of whatever form) makes people less prone to search and more prone to accept lower wage offers. In principle, then, the impact of impatience on exit rates from unemployment is ambiguous — being less prone to search decreases exit rates, while being willing to accept lower wages increases exit rates. The authors demonstrate, however, that under exponential discounting with plausible discount rates, the wage effect dominates and hence increased impatience should lead to increased exit rates from unemployment. In contrast, under β, δ preferences, if the increased impatience comes in the form of short-term impatience (smaller β), then the search effect dominates and hence increased impatience should lead to decreased exit rates from unemployment. DellaVigna and Paserman use data from the National Longitudinal Survey of Youth (NLSY) and from the Panel Study of Income Dynamics (PSID) to test the relationships between proxies for impatience and exit rates from unemployment, and conclude that the evidence supports the hyperbolic model. In follow-up work, Paserman (2004) estimates a structural model of job search of this form (using males in the NLSY), and finds support for $\beta < 1$.

Shui and Ausubel (2004) find evidence of hyperbolic discounting in credit-card borrowing. They have access to data in which consumers were randomly assigned to receive credit-card offers with different features. They find that consumers are more prone to accept offers with a low introductory

interest rate for a short duration, even when, given actual borrowing behavior, consumers would have been better off accepting offers with a larger introductory interest rate that lasts for a longer duration. Moreover, consumers seem reluctant to switch balances even after the introductory rate expires. Shui and Ausubel demonstrate that these findings are inconsistent with exponential discounting but consistent with hyperbolic discounting.

Fang and Silverman (2004) use the NLSY to test for hyperbolic discounting in welfare-vs.-work decisions among unmarried women with children. They develop a dynamic structural model in which women choose each period whether to work, stay home and take welfare, or stay home and not take welfare, where they incorporate the welfare rules that the women face (which differ across states). Their estimates for β are significantly different from 1; however, a puzzling aspect of their analysis is that it yields a surprisingly low estimate of short-term impatience — a yearly β on the order of .3.

DellaVigna and Malmendier (forthcoming) find evidence suggestive of hyperbolic discounting in health-club data that tracks members' usage over time. They find that, among those who sign up for a monthly contract, people pay an average price per visit of about $17 over the first six months, despite that fact that they could have chosen instead to pay $10 per visit by purchasing daily passes rather than the monthly contract. Moreover, the monthly contract has an automatic-renewal feature — the person's credit card is automatically charged every month unless the person cancels either by mail or by visiting the club — and they find evidence of procrastination in cancellation in the form of a significant duration (2.29 months) between last use and cancellation. Finally, they find that the average price paid per visit in initial months is positively correlated with the cancellation lag. They argue that these findings are suggestive of naive present bias: people sign up for the monthly contract naively expecting to use the health club frequently, and then later naively procrastinate cancellation.

Ariely and Wertenbroch (2002) provide experimental evidence on present bias and procrastination. They studied executive-education students who had to write three short papers for a class. For one group of students, each student chose a deadline for each of the three papers, while for a second group, evenly spaced deadlines were exogenously imposed. Students in the first group chose to impose deadlines on themselves, instead of making all papers due at the end of the term. But the deadlines that they chose allowed for more delay than the evenly spaced deadlines, and by various performance measures — e.g., their grades for the class — they fared worse than the

students with exogenously imposed, evenly spaced deadlines. They interpret these results as being driven by partially naive present bias. The fact that people choose to make commitments implies that they are worried about future misbehavior; however, the fact that these commitments seem to allow too much delay implies that people underestimate their need for commitment.

Gruber and Mullainathan (2005) find a different form of evidence of present bias. Their starting point is to recognize that models of hyperbolic discounting predict that people might smoke despite preferring not to smoke, and hence might be made better off by cigarette taxes. They attempt to directly test this proposition by investigating how survey measures of subjective well-being depend on local cigarette taxes. Using both U.S. and Canadian data, they indeed find that cigarette taxes seem to increase happiness among those prone to smoke. (As a reliability check, they also test whether other types of excise taxes have similar effects, and confirm that they do not.)

Finally, Shapiro (2005) examines consumption patterns among food-stamp recipients. He finds that, over the month between food-stamp deliveries, caloric intake declines by about 10-15 percent. He demonstrates that this behavior is inconsistent with exponential discounting and any plausible discount factor. After ruling out several other explanations, he argues that present bias seems to be the best explanation, and demonstrates that the evidence is roughly consistent with people having a daily $\beta \approx 0.9$ and a daily $\delta \approx 1$.

3. Incentives and Present Bias

We now move to the main topic of this article: the design of incentives when (some) people have present bias. In particular, our interest is how principals faced with present-biased people are likely to set up incentives, and how these incentives are likely to differ from the incentives that would be used when facing exponential discounters.

There are at least four categories of designed incentives, differing by who is the principal and what is the principal's motive. Perhaps the most obvious category is when benevolent third-parties — parents, teachers, friends, or the government — seek to achieve better outcomes for present-biased people. Because this category revolves around attempting to help the present-biased, it is

often referred to as designing "paternalistic" incentives.[12] The second category is closely related: present-biased people sometimes design their own future incentives in an attempt to achieve better outcomes for themselves. This category revolves around sophisticated attempts at self control through commitment devices.[13]

The third category of designed incentives is when profit-seeking firms or individuals attempt to exploit present-biased people by designing incentives that lead to outcomes that the individuals don't really want or (in the case of naifs) expect. This might take the form of pricing schemes designed to sell more output or to sell output at a higher per-unit price. Alternatively, this might take the form of hiring workers with incentive schemes that lead to low per-unit wages. Such exploitation is of interest, and may underlie many incentive structures that we see in the world. Even so, in many cases even profit-seeking firms will be motivated to seek efficient exchange for the usual mutual-gains-of-trade reasons. For instance, if a firm employs a procrastinator, it might attempt to bilk that person of money or extract extra work out of that person, but it is perhaps more likely to want to induce the worker to efficiently perform the tasks she was hired to do. Hence, the fourth category of designed incentives is when profit-seeking firms attempt to induce efficient exchange.

Whichever the category, the optimal design of incentives will depend on the nature of the particular environment under consideration. Even so, existing research on how present-biased people react to the natural incentives that they face in the environment around them, as well as emerging research on designing incentives in specific contexts, suggests some general principles for incentive design. In this section, we describe some general principles we have identified, and then illustrate

---

[12] The exploration of paternalism often sparks strong resistance among economists, on various grounds. Although we certainly concur that much further research is needed before advocating paternalism in practice, we feel it is a useful task to explore the nature of paternalistic incentives. Indeed, one interesting use of such analyses is that they might enable us to identify the motives behind real-world incentives — e.g., to assess whether they reflect a desire to help or a desire to exploit.

[13] The study of commitment devices has a special place in the history of the study of self-control problems (as the term "self-control" suggests), since the most blatant evidence that the conventional economic model isn't quite right is the existence of efforts to commit oneself.

them with a series of examples.

Principle #1: For exponential discounters, only gestalt incentives matter. For standard exponential discounters, all that matters are the overall costs and benefits of the available options. Hence, when designing incentives, altering details in a way that doesn't change gestalt incentives will have little effect on exponential discounters.

Principle #2: For present-biased people, the details of incentives matter. Because of their focus on immediate payoffs, many intricacies of incentives that would be irrelevant for exponential discounters can become quite important for present-biased people. In particular, they are sensitive to the exact timing of when costs and benefits are experienced. In addition, they are sensitive to exactly how decisions are made — e.g., choosing in advance vs. in the moment.

Principle #3: There often exist changes in incentives that influence present-biased people without much affecting exponential discounters. A direct implication of the first two principles is that by focusing on altering details without changing gestalt incentives, we can alter the behavior of present-biased people without changing the behavior of exponential discounters. We emphasize this principle because it will be important for dealing with heterogeneity in the world. In particular, even if we cannot identify who is and is not present-biased, it is often possible to modify incentives so as to help (or exploit) the present-biased without much affecting the efficiency of behavior by exponential discounters.

Principle #4: For present-biased people, it can be useful to "magnify" future costs or rewards. Because present-biased people put too little weight on future payoffs, behavior can often be improved by magnifying future costs or rewards.

Principle #5: For present-biased people, it can be useful to encourage prospective choice. When all consequences of a decision are sufficiently far in the future, present bias is not a problem. Hence, if it is possible to induce people to make decisions now about future behavior, we may be able to induce better behavior. For instance, by putting restrictions on when people have access to a good (e.g., limited buying hours), people

are forced to decide in advance whether they want to consume that good. Even more attractive, we can sometimes use prospective voluntary restrictions.

To illustrate these principles, we begin with a simple example that is a variant of the one-activity model that we studied in O'Donoghue and Rabin (1999*a*). Suppose there is an activity that a person must complete exactly once during the next 7 days. Completion of this activity generates both a reward and a cost as described in Table I.

TABLE I

REWARD AND COST AS A FUNCTION OF COMPLETION DAY

Completion Day	Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7
Reward	5	9	17	33	33	33	33
Cost	5	5	5	5	9	17	33
Net benefit	0	4	12	28	24	16	0

Table I reflects the gestalt reward and cost as a function of when the activity is completed; in this example, the reward and the cost both increase over time, with the net benefit peaking at Day 4. But the exact timing of when the person *experiences* the reward and the cost might vary depending on the type of activity under consideration. If this pattern of rewards and costs is generated by an onerous activity, such as completing a school assignment, the cost is incurred immediately while the reward is received in the future. For instance, if she completes the problem set on Day 3, then she incurs a cost of 5 on Day 3, and receives a reward of 17 sometime after Day 3. If instead this pattern of rewards and costs is generated by a pleasurable activity, such as eating a tempting food, the reward is received immediately while the cost (e.g., health cost) is incurred in the future. For instance, if she eats the food on Day 3, then she receives a reward of 17 on Day 3, and incurs a cost of 5 sometime after Day 3.

In this example, an exponential discounter with $\delta = 1$ will complete the activity on Day 4 regardless of the type of activity. This example highlights how exponential discounters focus solely on the gestalt rewards and costs of their available actions; the timing of rewards and costs is (virtually)

irrelevant to their decision.[14]

Now consider a present-biased person with $\delta = 1$ and $\beta = 1/2$. From a prior perspective (or even from a Day-1 perspective), the person would most like to complete the activity on Day 4, regardless of the timing of rewards and costs. The person's actual behavior, however, depends crucially on this timing. In particular, if it is an onerous activity, the person is prone to procrastinate incurring the immediate cost, and hence delays beyond Day 4. In contrast, if it is a pleasurable activity, the person is prone to preproperate — grab the reward too soon — and hence does the activity before Day 4. Exactly how much the person procrastinates or preproperates depends on whether she is sophisticated or naive. If it is an onerous activity, naifs procrastinate until Day 7 while sophisticates complete the activity on Day 5. If it is a pleasurable activity, naifs do it on Day 3, while sophisticates do it on Day 1.

In addition to demonstrating how the timing matters for present-biased people in a way that is irrelevant for exponential discounters, this example also illustrates the potential value of prospective choice. It is only because present-biased people are making day-to-day, in-the-moment decisions that they end up not doing the activity in Day 4. If instead the person were forced in advance to commit to a day, people would choose Day 4 irrespective of sophistication and irrespective of the timing of rewards and costs. Indeed, even if the person were merely forced each day to commit to tomorrow's behavior (wait vs. do it), the activity would always be performed on Day 4.

Finally, this example illustrates some simple principles with regard to voluntary prospective commitments. As discussed above, if the person commits in advance, she'll commit to Day 4. But would the person want to make such a commitment? The answer depends on whether the person is sophisticated. Sophisticates correctly predict their behavior, and so they would see value to making a commitment. Hence, sophisticates would even be willing to pay for the ability to commit. More generally, sophisticates are often on the lookout for ways to commit their future behavior, and are sometimes even willing to pay for such opportunities. In contrast, naifs incorrectly think that they will behave themselves in the future, and so they see no value in making a commitment.

---

[14] If exponential discounters are impatient, the timing of rewards and costs can matter somewhat. The more general point, however, is that exponential discounters are influenced by only the overall present discounted value of costs and benefits — that is, holding their present discounted value constant, altering the timing cannot alter their behavior.

But note that this does not mean that naifs won't make commitments. If making a commitment is costly, either directly or indirectly via reduced flexibility, they will turn it down, as would an exponential discounter. If, however, there were some incentive to commit — perhaps we offer to pay the person $1 if they commit in advance to a day — then they would take it. More generally, naifs, like sophisticates, can benefit from commitment, but since they themselves don't see the need, principals need to create incentives for them to make commitments.

We move next to a somewhat richer context that captures a common situation that we call "cumulative procrastination". Consider a student who must read P pages in T days — that is, she must choose a path $(p_1, ..., p_T)$, where p_t is the number of pages read on day t, such that $\sum_{t=1}^{T} p_t = P$. On any given day, her disutility is equal to the number of hours spent reading — that is, $u_t = -h_t$. The number of hours required to read p pages is p^α for some $\alpha > 1$ — that is, $h_t = (p_t)^\alpha$. The key feature of this situation is that, because there are increasing marginal costs to reading on any given day, efficiency calls for the person to spread out her work regularly rather than doing it all in a few days.

First consider *June Mae*, an exponential discounter with $\delta = 1$. How will she react to these circumstances? The answer is simple by the construction: June Mae will read the same amount on every day, which means reading P/T pages per day and therefore $(P/T)^\alpha$ hours per day.

Now consider *April Mae*, a naive person with $\delta = 1$ and $\beta < 1$. What will she do? On day 1, she expects to read the same amount on all remaining days beginning tomorrow. Specifically, she expects that if she reads p_1 pages today, then she'll read $(P - p_1)/(T - 1)$ pages on each of the $T - 1$ remaining days. Hence, on day 1 she'll choose p_1 to minimize

$$(p_1)^\alpha + \beta \left[(T-1) \left(\frac{P - p_1}{T-1} \right)^\alpha \right].$$

It follows that on day 1 April Mae will choose

$$p_1 = \frac{\beta^{1/(\alpha-1)}}{T - 1 + \beta^{1/(\alpha-1)}} P.$$

Iterating this logic, one can show that on day $t > 1$ April Mae will choose

$$p_t = \left(\prod_{n=1}^{t-1} \frac{T-n}{T-n+\beta^{1/(\alpha-1)}} \right) \frac{\beta^{1/(\alpha-1)}}{T-t+\beta^{1/(\alpha-1)}} P.$$

Table II uses this equation to present how the two types behave for some specific examples.

TABLE II

BEHAVIOR IN CUMULATIVE-PROCRASTINATION EXAMPLE

	$T = 30$, $P = 30$, $\alpha = 2$		$T = 30$, $P = 60$, $\alpha = 1.7$	
	June Mae	April Mae with $\beta = 1/2$	June Mae	April Mae with $\beta = 3/4$
h_1	1 hour	$15\frac{1}{2}$ minutes	$3\frac{1}{4}$ hours	99 minutes
h_2	1 hour	16 minutes	$3\frac{1}{4}$ hours	101 minutes
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
h_{T-1}	1 hour	5.9 hours	$3\frac{1}{4}$ hours	9.6 hours
h_T	1 hour	23.76 hours	$3\frac{1}{4}$ hours	19.25 hours
Total Hours	30 hours	51 hours	$97\frac{1}{2}$ hours	108 hours

Consider the first example, where students must read 30 pages in 30 days. June Mae spreads out this work evenly, and hence reads 1 hour per day and a total of 30 hours. On Day 1, April Mae plans to spread out her work evenly, but her preference for immediate gratification leads her to choose to read less today and more on all future days. Specifically, on day 1 she reads for $15\ 1/2$ minutes, planning to read for 62 minutes on each of the remaining 29 days. Hence, when she makes this decision to indulge immediate gratification, she has decided to increase total hours by $13\ 1/2$ minutes.

Unfortunately for April Mae, on Day 2 she changes her mind and decides to indulge immediate gratification once again, this time reading for 16 minutes and planning to read for 64 minutes on each of the remaining 28 days. Indeed, she repeatedly revises her plans to indulge immediate gratification, and as a result she ends up reading for many hours just before the deadline — nearly 6 hours on day 29, and $23\ 3/4$ hours on day 30 — an "all-nighter". Overall, whereas June Mae reads for a total of 30 hours, April Mae reads for a total of 51 hours.[15]

Notice in this example that April Mae did not plan to pull an all-nighter on Day 30. Rather, due to a series of decisions to indulge immediate gratification, she kept backloading more and more work, and ended up forced to do an all-nighter. This example illustrates the more general point that people with present bias do not choose optimal life courses in the way that exponential discounters do. For people who are (to some extent) naive, eventual life courses can be the result of a series of revisions of plans, and look very different from the person's initial plan. Even for fully sophisticated people,

[15] If we were to consider *September "Sally" Mae* who is sophisticated with $\delta = 1$ and $\beta = 1/2$, we'd find (as the solution to a fairly tedious problem) that Sally Mae reads for a total of 39 hours.

who never deviate from their plans, to the extent that they cannot control their future behavior they also do not choose optimal life courses.

We next explore the role of incentives in this example, and in particular how some micro-details that would have small effects for exponential discounters can have significant impact on a naive person with present bias, like April Mae. Consider first a minor modification of this environment: People experience a small bit of enjoyment from studying with their friends — specifically, whereas $u_t = -h_t$ if a person studies alone, $u_t = -.99h_t$ if the person studies with a friend. Suppose further, however, that studying together requires scheduling a day *in advance*, and that the cost of backing out from planned studying is significant. For exponential discounters, this modification would have no effect on hours worked, although students would now work together. For naifs like April Mae, however, the enjoyment of studying together would serve as a serendipitous commitment. In particular, they would always plan in advance to work the optimal one hour per day, and hence end up working 30 hours instead of 51 hours.

If, like April Mae, the person were completely naive, then it is important that she enjoy studying with friends. If instead she had a slight distaste for studying with friends, April Mae would never do so, and end up again with the 51 hours total. For people who are (to some degree) sophisticated, they might make use of a friend as a commitment device even if they had a slight (or even significant) distaste for studying with friends.

If a teacher wanted to help April Mae study more efficiently, what might he do? Taking the model literally, there is an obvious answer: just require a page every day. Such a policy will induce April Mae to choose the efficient one hour per day. Moreover, this policy would have no effect on June Mae. Hence, here we have an example of a policy that can have big benefits for people with present bias that does not affect standard exponential discounters.

A problem with this simple solution is that we probably shouldn't take the model literally. It is more realistic to assume day-to-day variation in the costs of reading, in which case the optimal strategy would involve more reading on low-cost days, less reading on high-cost days, and some "precautionary reading" early on to protect against high-cost days. If so, then requiring a page a day would likely make June Mae worse off. Even so, there is a natural compromise: Impose interim but not too micro-managed incentives. For instance, requiring 5 pages every 5 days in the certainty case would reduce April Mae's reading from 51 hours to about 39 hours, and even with uncertainty leaves June Mae some flexibility.

In fact, an even better way to help April Mae with limited harm to June Mae is to use voluntary restrictions. Instead of requiring 30 pages in 30 days, as in our benchmark example, suppose the teacher offers students *at the start of the month* a choice between two sets of requirements. If they like, they can read 31 pages in 30 days, without any interim deadlines. Alternatively, they can be required to read at least 10 pages after 10 days, at least 20 pages after 20 days, and 30 pages after 30 days. Why the extra page in the former case? Because adding this extra work creates a small incentive to choose the interim deadlines. As a result, if the interim deadlines are too costly in terms of reducing flexibility, all types will choose the former option. But if the interim deadlines are not a big deal, everybody — naifs, sophisticates, and exponential discounters — will all choose the interim incentives.[16]

Finally, there is a way to employ magnification here with a mechanism that we see quite a lot (largely for other reasons, but partly for present bias): Give people false deadlines. This raises perceived benefits of reading more now. A teacher can induce more efficient reading early on by people with present bias by saying that the 30 pages must be read in 15 days rather than 30 days — indeed, in our benchmark example, this would induce an hour of reading on day 1. Then, after a few days, the teacher can tell students that they really have 20 days to do the reading, and so forth. If students believed the teacher (which they shouldn't, and for credibility, moral, and other reasons, this is not a wise place to use false deadlines) the teacher could, even with uncertainty, greatly improve efficiency. (But unlike our other schemes, this *would* harm exponential discounters significantly.)

For our third example, we leave to realm of procrastination and consider instead overconsumption of harmful goods — specifically, whether people develop and maintain "bad habits". For both classically addictive substances and for all sorts of activities like diet, exercise, and lifestyle, researchers have started to pay more attention to the role of habit formation. We consider here in highly stylized form an example of a bad habit. Suppose your happiness on day t is given by:

$$u_t \equiv \begin{cases} \rho_L \ln(x_t + 1) - px_t & \text{if } x_{t-1} = 0 \\ \\ \rho_H \ln(x_t + 1) - px_t - (\Gamma + \gamma \ln(x_{t-1} + 1)) & \text{if } x_{t-1} > 0, \end{cases}$$

where p is the price of the addictive good, $x_t \geq 0$ is consumption of the addictive good on day t,

[16] As we discussed in Section 2, Ariely and Wertenbroch (2002) provide an experimental study of voluntary interim incentives.

and $\rho_L, \rho_H, \Gamma, \gamma > 0$ are taste parameters, with $\rho_H > \rho_L \geq \gamma$. This formulation reflects that the person allocates current income between the addictive good and a composite good — specifically, $-px_t$ reflects that expenditures on the addictive good reduce funds available for consumption of the composite good.

This formulation assumes a highly stylized version of addiction: if a person didn't consume at all last period, then she is unaddicted, whereas if she consumed any positive amount, she is addicted. But the model incorporates the two essential components of bad habits that have been much discussed in the economics literature. First, current consumption imposes a negative cost in the future — as reflected by the fact that if the person consumes x now, then she will incur a utility cost of $\Gamma + \gamma \ln(x+1)$ next period.[17] Second, current consumption creates habit formation in the sense that it increases the benefits from future consumption — as reflected by the fact that consuming now raises the benefits from next period's consumption from $\rho_L \ln(x+1)$ to $\rho_H \ln(x+1)$.

It could be optimal to develop a bad habit, and whenever it is, people with self-control problems do so, although their steady-state consumption is suboptimally large. Here, we focus on an example where it is not optimal to develop a bad habit. Specifically, we suppose that

$$u_t \equiv \begin{cases} 15\ln(x_t+1) - x_t & \text{if } x_{t-1} = 0 \\ 20\ln(x_t+1) - x_t - (20 + 10\ln(x_{t-1}+1)) & \text{if } x_{t-1} > 0. \end{cases}$$

Consider how an exponential discounter with $\delta \approx 1$ would behave in this example. Because the incentive to consume is stronger when addicted than when unaddicted, this person would choose either an unaddicted steady state or an addicted steady state. In the unaddicted steady state, the person consumes $x_t = 0$ every period, and hence her instantaneous utility in all periods is 0. In the addicted steady state, she consumes $x_t = 9$ every period, and hence her instantaneous utility in all periods is $20\ln 10 - 9 - (20 + 10\ln 10) = -5.97$.[18] Because the unaddicted steady state is clearly better, an exponential discounter with $\delta \approx 1$ will choose the unaddicted steady state. Notice that this conclusion would hold even if the person started out addicted. For example, even if she had formed the habit in her youth, she would immediately stop consuming.

Consider next a naive hyperbolic discounter with $\beta = .8$ and $\delta \approx 1$. When thinking about the

---

[17] Herrnstein, Loewenstein, Prelec, & Vaughan (1993) labeled such effects "negative internalities".

[18] When addicted and consuming, optimal consumption x maximizes $20\ln(x+1) - x - \delta 10\ln(x+1) \approx 10\ln(x+1) - x$, which yields $x = 9$.

future, the naif, like the exponential discounter, views the unaddicted steady state as clearly better than the addicted steady state, and hence plans and expects never to consume in the future. Even so, her preference for immediate gratification might lead her to choose short-term consumption. In this example, if she starts out unaddicted, then she will not consume. Given the relatively small temptation to consume when unaddicted, she is not willing to incur the future costs even given her preference for immediate gratification.[19]

But suppose instead that she starts out with the habit — again, perhaps because she had consumed in her youth. If she thought her choice was to quit now or never, then she would quit, because she views the unaddicted steady state as clearly (essentially infinitely) better than the addicted steady state, and so no short-term temptation would lead her to choose a long-term addiction. But her choice is not to quit now or never. Rather, she also has the option to consume now and then quit tomorrow, and her preference for immediate gratification combined with the higher temptation when addicted in fact makes this her preferred option.[20] Moreover, because she feels this way on every day of her life, she would naively spend her entire life planning to quit in the near future.

Note several features of this example. First, once again we see that details that are irrelevant for standard exponential discounters matter for people with present bias. Here, initial conditions — whether a person starts out unaddicted or addicted — is irrelevant to exponential discounters. While this conclusion is not true for any example, it reflects that, for a reasonably patient exponential discounter, unless short-term transition costs are very large, the person will choose the best steady state. In contrast, because a person with present bias might naively procrastinate initiating the transition to a different steady state, initial conditions can determine lifetime outcomes.

Second, notice the distinction between people's choices and their life courses. Initially addicted naifs end up with a lifetime addiction, but they never chose this life course. Rather, on each day they are choosing merely one more day of addiction. But the end result is that they maintain a

---

[19] When unaddicted, the best x for short-term consumption maximizes $15\ln(x+1) - x - \beta\delta 10\ln(x+1) \approx 7\ln(x+1) - x$, which yields $x = 6$. But since $15\ln(7) - 6 - \beta\delta[20 + 10\ln(7)] < 0$, the person prefers $x = 0$ to $x = 6$.

[20] When addicted, the best x for short-term consumption maximizes $20\ln(x+1) - x - \beta\delta 10\ln(x+1) \approx 12\ln(x+1) - x$, which yields $x = 11$. And since $20\ln(12) - 11 - \beta\delta[20 + 10\ln(12)] > 0$, the person indeed prefers $x = 11$ to $x = 0$.

lifelong addiction despite the fact that on every day of their life they'd prefer quitting immediately.

Third, notice the power of prospective choice: If by market technology or by government regulation, she had to decide ahead of time — just a day ahead of time would be enough — whether to consume, then no matter the initial conditions she would end up choosing the unaddicted steady state.[21]

Our example suggests that people may be over-prone to maintain (and more generally to develop) bad habits. A natural policy response would be to tax such goods, because by raising the market price we can induce less consumption. Such a policy is a nice example of magnification: because agents are prone to underweight the future costs of consumption, we might be able to help them to behave better by raising the monetary cost of consumption. Although such "sin" taxes might hurt those who consume optimally, Gruber and Koszegi (2004) and O'Donoghue and Rabin (2003, 2005) investigate this trade-off, and demonstrate that sin taxes can often be welfare-improving; we'll return to this point in the next section.

It is also worth considering richer forms of incentives. For instance, sin licenses (or tax vouchers) may be a useful technique to induce people to make prospective decisions. Specifically, suppose we set a high presumptive tax but then let people buy in advance a license that exempts them from the tax this year. By setting the license fee equal to, say, the tax that would be paid on 100 packs, we can induce people to assess whether they plan to purchase more or less than 100 packs this year. Moreover, for those who choose not to buy the license, the high presumptive tax provides a motivation not to buy in the moment.

Consider a somewhat more speculative proposal. Suppose we impose a "clean-living" surcharge: anybody unhooked on this habit at any point pays $2 a day, and anybody hooked pays nothing. In addition, we offer a "quitting bonus" of $100 to anybody hooked who quits. How would people react? Unhooked people, whether an exponential discounter or a naif, will choose to become an addict only if being an addict is not very bad — less bad than paying $2 a day. And they

---

[21] Extrapolating from our example, this type of framework generates implications for how people might react to quantity discounts, which are quite prevalent in the market (cartons of cigarettes vs. packs, big bags of chips vs. small). Such reactions can provide a hint as to whether people want or expect to consume the good in the future, and indeed Wertenbroch (1998) has documented that people tend to buy "vice" goods in smaller quantities than "virtue" goods.

might choose constantly starting and stopping, but only if doing so is worth $49 a day.

What about those who are addicted? First note that an exponential discounter, no matter her utility function, would not plausibly quit in response to this policy. For any plausible impatience, it would never be worth receiving a one-time $100 payment in return for having to pay $2 for the rest of her life. Hence, any rational addict will not be influenced by this policy. In contrast, naifs might quit in response to this policy. The reason is *not* because they are tempted by the $100 now in exchange for $2 every day for life — like exponential discounters, for plausible impatience they view that as a bad deal. The reason is rather that naive addicts might already be *planning* to quit soon, and so they expect to be paying the $2 per day for life in any event. If so, their choice boils down to choosing between $100 and pain of quitting now vs. $100 and pain of quitting tomorrow, and the former might be better.

4. Heterogeneity and Flexibility

A central theme in the emerging literature on incentives for present-biased people is the importance of accommodating heterogeneity among agents, and facilitating flexibility in behavior. Casual evidence suggests that people vary a great deal in their degree of self-control problems, their degree of awareness of self-control problems, and (of course) in their intrinsic tastes for different activities. This heterogeneity complicates incentive design because, as we highlighted in Section 3, optimal incentives differ for different types. Moreover, even individual agents will typically face uncertainty about their own future tastes, needs, outside options, and constraints. As a result, optimal incentives may need to permit some flexibility to agents.

Researchers have begun to address both issues. Heterogeneity has been most discussed in the recent literature on designing paternalistic incentives, where there is concern that interventions designed to help people who make errors might cause significant harm to people who do not make errors. One strand of literature has dealt with this issue by identifying non-intrusive interventions that would combat errors while attempting to impose little or no harm on fully rational people. This basic approach has been put forth by various authors under various labels: "cautious paternalism" (O'Donoghue and Rabin 1999*c*), "asymmetric paternalism" (Camerer, Issacharoff, Loewenstein, O'Donoghue, and Rabin 2003), "libertarian paternalism" (Sunstein and Thaler 2003*a*,2003*b*), and "benign paternalism" (Choi, Laibson, Madrian, and Metrick 2003). Although this approach has

been applied to a broad set of errors, present bias is a common focus, and principles #1, #2, and #3 from Section 3 are particularly pertinent for this goal.

In O'Donoghue and Rabin (2003, 2005), we directly address heterogeneity by deriving optimal policy as a function of the distribution of types in the population. Specifically, we investigate consumption of goods that generate immediate consumption utility but future health costs, where people differ both in terms of their degree of present bias and in terms of their tastes for the good. We show that taxing such goods can yield significant benefits for present-biased people by inducing them to reduce over-consumption, while at the same time such taxes may have very little effect on the utility of fully rational people. Moreover, if tax proceeds are returned to consumers as a uniform lump sum, then there is a net redistribution from present-biased people to fully rational people, and as a result such taxes can sometimes end up making everybody better off — that is, they can create a Pareto improvement.

Heterogeneity is also relevant for firms concerned with either efficiency or exploitation — e.g., incentives designed to exploit people with present bias might end up losing revenue on those without it. Indeed, DellaVigna and Malmendier (2004) derive optimal two-part tariffs for a monopolist facing present-biased consumers, and they show that the optimal contract may look very different for people with sophisticated present bias, people with naive present bias, and those without present bias. They do not, however, consider the screening problem. More recently, Eliaz and Spiegler (forthcoming) investigate a contracting model in which a firm attempts to screen sophisticates from naifs so as to promote efficient exchange (efficient commitment) with sophisticates while at the

same time exploiting naifs.[22]

Researchers have also started to address the need for flexibility. In many analyses, it has been implicit. For instance, O'Donoghue and Rabin (1999*b*) investigate incentives to combat procrastination when agents face day-to-day uncertainty in the opportunity cost of engaging in an activity. In such an environment, it is typically not optimal to impose an immediate deadline, because the agent should be permitted some flexibility to find a good time to complete the task. Recently, Amador, Werning, and Angeletos (2005) explicitly investigate the optimal trade-off between commitment and flexibility in a savings-consumption model. (They find, for instance, that imposing a minimal level of savings is always a part of the optimal solution.)

An important aspect, in our view, that is missing from existing research is the use of more subtle, "voluntary" screening devices. If our goal is to implement a policy that combats present bias, but we are worried that this policy might hurt people who don't have present bias, why not let people voluntarily select in advance whether to be subject to the policy. If everyone were fully sophisticated, such a scheme can be very effective, because we can count on all agents to choose whatever incentives are best for them. A problem arises, however, when there are naive agents, because they may not see the value of the policy. Even so, we can deal with this problem by providing an ex ante incentive to volunteer — perhaps framed in reverse where the default is being subject to the policy but for a small price people can buy out the restrictions. We suspect this technique will play a prominent role as the literature evolves.

---

[22] Formal analyses of heterogeneous populations are important to advance research on present bias. Demonstrating, for instance, the extent to which proposed incentives to combat present bias are robust to the existence of vast numbers of fully self-controlled agents will highlight that the merits of such incentive schemes do not rely on an assumption that virtually everyone is present-biased. It is also our hope that research acknowledging and investigating the prevalence of different degrees of self-control problems will help to rebut the over-frequent supposition among doctors, preachers, and laypeople that certain activities reflect purely unwarranted pursuit of immediate gratification. Research that incorporates self-control problems and carefully identifies those instances when consumption is actually optimal presumably has a better chance of persuading skeptics that such instances are indeed optimal than does research that a priori assumes that whatever a person does is optimal.

Voluntary incentives may be particularly important as a means to deal with the need for flexibility. Perhaps the most important aspect of the need for flexibility is heterogeneity in this need. For some agents, flexibility is quite important, while for others it is not. Much as for heterogeneity more generally, we suspect that the best way to deal with heterogeneity in the need for flexibility is via voluntary commitments. In our cumulative-procrastination example, for instance, we discussed how we could let people choose between having slightly more total work and lots of flexibility vs. slightly less total work but reduced flexibility. If uncertainty is high and thus flexibility is valuable, people will choose the former; if they have relatively little value for flexibility, they will choose the latter.

Indeed, procrastination is a useful context in which to make these ideas even more concrete. Our examples in Section 3 illustrate the usefulness of deadlines for combatting procrastination. Deadlines are perhaps most important for "optional" tasks. People face many tasks in their lifetime that are on net beneficial but need not be done. These might be mundane tasks — such as building shelves in the garage or investing in a better technology for mowing the lawn — or they might be more economically relevant tasks — such as investing in human capital or improving one's portfolio of financial investments. Because such tasks have no inherent deadline, procrastination can lead to infinite delay — that is, present-biased people may never complete the task even when they would prefer to complete the task early on rather than never. In such cases, imposing *any* not-too-long-away deadline can yield significant benefits.

A major worry with any final deadline, however, is the need for flexibility — with a final deadline, if due to bad luck a person misses the deadline, she has lost the opportunity to ever do the task. To address this worry, there is another type of deadline that makes use of the difference between the nature of procrastinatory delay and that of rational delay. To the extent that procrastination is driven by a naive belief that a task will be completed in the near future, putting limits on the frequency with which one can carry out tasks can help to combat procrastination. If a person can take a job-certification course on any day for the rest of her life, then she might repeatedly delay based on a daily belief that she'll do it tomorrow or within a few days. If instead the job-certification course is offered, say, only once per month, then the person can no longer justify delay based on a naive belief that she'll do it soon. Rather, she'll be forced to recognize that delay now means at least a one-month delay, and if the consequences of a one-month delay — e.g., lower wages for the

next 30 days — are sufficiently large, she won't delay.[23]

While this technique solves the final-deadline problem, it can still be critiqued for lack of flexibility — if it turns out that a person cannot make the course on the night it is offered, she must wait a month before her next opportunity. Here, voluntary incentives can be used as a means to address this concern. For instance, we might offer people two choices for a job-certification course. For an upfront fee of $100, they can take the course on any day they like; but for an upfront fee of $90, they can take the course only on the first of any given month (and we do not permit people to switch from their initial choice). People who face significant day-to-day uncertainty and thus put a high value on flexibility will choose the former option; people who face less day-to-day uncertainty will choose the latter option.

In addition to flexibility, we should also be worried about heterogeneity in the propensity to procrastinate, due either to heterogeneity in present bias or heterogeneity in how unpleasant a person finds the task. There are in fact other ways to combat procrastination that deal with such heterogeneity. One technique is to create direct incentives (rewards) for people to accelerate their plans. Suppose, for instance, that $20 vouchers are given to heroin addicts who maintain clean urine tests for a month. Heroin addicts who rationally prefer to remain heroin addicts presumably will not be influenced much by such vouchers — if they are influenced, it is either because they barely preferred to remain heroin addicts or because temporary restraint is not very costly. But among heroin addicts who had been planning to commence withdrawal, such small financial rewards can lead them to immediately initiate such plans when they otherwise would have procrastinated. Similarly, small vouchers, especially dated ones, by commercial enterprises that encourage people to sign up for a service are much more likely to have an effect on people planning to join such a service than those that never plan to.

A second technique is to focus less on combatting procrastination and more on limiting the harm from procrastination. In particular, in many cases, policymakers or firms get to specify default outcomes that are implemented unless people opt out. By choosing these defaults wisely, it can

---

[23] Even better — to the extent it is possible — the course could be offered every day, but certifications for the month are processed only on, say, the last day of the month. Then the person could still take the course on any day, but the last day of the month would loom large as a deadline to receive the benefits for the upcoming month.

be possible to implement better outcomes. A nice example of such defaults are in the Save More Tomorrow (SMART) Plan that has recently been marketed by Thaler and Benartzi (2004). They operate under the presumption that people want to initiate savings plans (for retirement), but have difficulty initiating those plans. To overcome this tendency, they offer people the option to set their default outcome to increased saving beginning in the future, with regular increases in the amount saved up to some limit, all implemented via automatic withdrawals from one's paycheck. They document that many people indeed choose to sign up for such plans, and end up saving more than those who do not sign up for the plan.

A third technique is to promote "active choice" — to the extent possible, force people to actually make a decision, or at least make them perceive that they are forced to make a decision. This technique was first proposed by Choi, Laibson, Madrian, and Metrick (2004), who studied a firm whose 401(k) plan unintentionally set up such a situation. Specifically, new employees were given the enrollment form for the company's 401(k) plan as part of a package of forms that needed to be filled out within 30 days. While some of these forms really were required, the 401(k) enrollment form was not — if the employee failed to turn it in, the employee just wasn't enrolled in the plan. Even so, employees seem to have viewed it as a requirement, and as a result more employees signed up for the 401(k) plan relative to what happened when the company later moved to an online system that permitted enrollment on any day.

5. Discussion and Conclusion

In this section, we discuss some broader issues and conclude. We have described how people seem to have a time-inconsistent taste for immediate gratification, and we have explored how this taste can lead to different reactions to incentives than does the standard economic model of choice. Present bias is, however, only one improvement in how to conceptualize people's behavior. Researchers have started to investigate other improvements, and these also can lead to different reactions to incentives.

In the realm of intertemporal choice, a related improvement is "temptation utility," as formulated by Gul and Pesendorfer (2001), wherein people find it unpleasant to have unchosen tempting options. In our example of a bad habit, for instance, if consuming the good is more tempting than not consuming, then a person may prefer to consume the good so as to avoid the temptation disutil-

ity that she would experience while not consuming. Moreover, even if she would be able to resist the temptation, she might in advance attempt to eliminate her access to the good so as to eliminate the temptation disutility that she would feel while not consuming. Because the details of incentives can influence the extent to which people experience temptation disutility, then, much as for present bias, more than just gestalt incentives matter.[24] Another related improvement in the realm of intertemporal choice is "projection bias," as formulated by Loewenstein, O'Donoghue, and Rabin (2003), wherein people underappreciate taste changes — specifically, perceived future tastes are in between actual future tastes and current tastes. In our example of a bad habit, if an unaddicted person underappreciates the negative cost associated with being addicted, she might choose to become an addict even when she shouldn't. Again, because the details of incentives can influence both current and future tastes, more than just gestalt incentives matter.

It is instructive to compare the implications for incentive design of these different ways to improve the standard economic model. The models of temptation utility and of fully sophisticated present bias share two features: (i) people correctly predict their future behavior, and (ii) people's ex ante preferences correspond to their normative preferences. Two implications follow. First, from the perspective of designing incentives to promote better outcomes, we can never do better than eliciting from individuals what they view as their optimal commitments. Second, from the perspective of designing exploitative incentives, the best we can do is to create temptations and to undermine commitments. Contrast these conclusions with the models of projection bias and of naive present bias, where people might incorrectly predict future behavior. For such problems, designing incentives becomes more difficult. For incentives designed to produce better outcomes, it may be necessary, as we have discussed for naive present bias, to create incentives to encourage reluctant agents to take on commitments. And for exploitative incentives, it may be possible to induce people to pay for things that they will end up not valuing.

---

[24] Indeed, a few recent papers have studied the design of incentives when people experience temptation utility — see for instance Amador, Werning, and Angeletos (2005) and Esteban, Miyagawa, and Shum (2005). Temptation utility and (fully sophisticated) present bias are sometimes discussed as mutually exclusive models, in part because they often yield similar predictions. In our view, however, these models reflect two distinct psychological phenomena that might both influence behavior.

Indeed, there exist real-world examples that seem naturally interpreted as exploitation of naivete, where firms offer long-term contracts that take advantage of naive mispredictions. For instance, DellaVigna and Malmendier (forthcoming) argue that health clubs may encourage long-term contracts to take advantage of people's naive beliefs that they will frequently use the health club. Shui and Ausubel (2004) describe how consumers seem to underpredict future credit-card borrowing, and hence are too sensitive to the introductory interest rate and pay too little attention to the long-term (post-introductory) interest rate. If so, then we might expect credit-card companies to take advantage of this propensity by pushing cards with low introductory rates and high long-term rates (although it is unclear how competition will affect their ability to make money in this way). More generally, there seem to be numerous examples of firms attempting to make default outcomes continued subscription to some service. The health club studied by DellaVigna and Malmendier, for instance, used the common approach of having a monthly membership that is automatically renewed by credit card, and they found evidence of significant procrastination in cancelling the membership — specifically, people continued to pay for the membership long after their last use of the health club.

We have emphasized throughout how the existence of present bias creates a need for incentives to promote better outcomes. An important issue is to what extent will the market provide such incentives. Researchers have long discussed how sophisticates would demand commitment devices, and have pointed to the existence of Christmas clubs, fat farms, and other commitment devices provided by the market as evidence that people have self-control problems. One might further infer from these examples that the market is an effective provider of commitment devices, and hence there is limited need for any "paternalistic" interventions.

Recently, however, economists have demonstrated that markets might, in fact, undermine the scope for commitment devices. In particular, a commitment device will work only to the extent that another firm doesn't come along offering ways to undo that commitment. For instance, Laibson (1997) demonstrates how illiquid savings instruments that tie up people's wealth in the short term can be a valuable commitment device to combat under-saving due to present bias. But he then discusses how the rapid expansion of the credit-card industry in the early 1990's might have undermined this commitment device by giving people the ability to borrow against that illiquid wealth. DellaVigna and Malmendier (2004) and Koszegi (forthcoming) explore more systematically the types of situations in which the market is likely to be able or unable to provide commitment

devices. They show, for instance, that for investment goods for which the optimal commitment involves low per-unit prices (below marginal cost), there is little scope for other firms to undermine this commitment, whereas for indulgent goods for which the optimal commitment involves high per-unit prices (above marginal cost), there is significant scope for other firms to undermine this commitment.

Researchers are only beginning to study incentive design when people have present bias (and other behavioral phenomena). For such people, the usual intuitions still apply — e.g., in a standard moral-hazard model, the principal still faces the usual trade-off between insurance and incentives. But new intuitions also arise, and in particular details that would be mostly irrelevant to standard agents can become quite important. We hope that by describing some general principles, and by highlighting the importance of heterogeneity and flexibility, this paper will help lay the groundwork for this emerging literature.

Department of Economics, Cornell University, 414 Uris Hall, Ithaca, NY 14853-7601, U.S.A.; edo1@cornell.edu; http://www.people.cornell.edu/pages/edo1/,

and

Department of Economics, 549 Evans Hall #3880, University of California at Berkeley, Berkeley, CA 94720-3880, U.S.A.; rabin@econ.berkeley.edu; http://emlab.berkeley.edu/users/rabin/.

References

AINSLIE, G. (1991): "Derivation of 'Rational' Economic Behavior from Hyperbolic Discount Curves," *American Economic Review*, 81, 334-340.

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

AINSLIE, G. AND N. HASLAM (1992*a*): "Self-Control," in *Choice Over Time*, ed. by G. Loewenstein and J. Elster. New York: Russell Sage Foundation, 177-209.

— (1992*b*): "Hyperbolic Discounting," in *Choice Over Time*, ed. by G. Loewenstein and J. Elster. New York: Russell Sage Foundation, 57-92.

AKERLOF, G. (1991): "Procrastination and Obedience," *American Economic Review*, 81, 1-19.

AMADOR, M., I. WERNING, and G.-M. ANGELETOS (2005): "Commitment vs. Flexibility," Unpublished Manuscript, MIT.

ANGELETOS, G.-M., D. LAIBSON, A. REPETTO, J. TOBACMAN, AND S. WEINBERG (2001): "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15, 47-68.

ARIELY, D. AND K. WERTENBROCH (2002): "Procrastination, Deadlines, and Performance: Self-Control by Precommitment," *Psychological Science*, 13, 219-224.

BANNISTER, P. (2004): "25 Fascinating Facts About Personal Debt," *www.bankrate.com*, September 20, 2004 (http://www.bankrate.com/brm/news/debt/debtguide2004/debt-trivia1.asp).

BENABOU, R. AND J. TIROLE (2002): "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 117, 871-915.

CAMERER, C., S. ISSACHAROFF, G. LOEWENSTEIN, T. O'DONOGHUE, AND M. RABIN (2003): "Regulation for Conservatives: Behavioral Economics and the Case for 'Asymmetric Paternalism'," *University of Pennsylvania Law Review*, 151, 1211-1254.

CARRILLO, J. AND T. MARIOTTI (2000): "Strategic Ignorance as a Self-Disciplining Device," *Review of Economic Studies*, 67, 529-544.

CHOI, J., D. LAIBSON, B. MADRIAN, AND A. METRICK (2003): "Optimal Defaults," *American Economic Review (Papers and Proceedings)*, 93, 180-185.

— (2004): "Optimal Defaults and Active Decisions," Unpublished Manuscript, Harvard University.

DELLAVIGNA, S. AND U. MALMENDIER (2004): "Contract Design and Self-Control: Theory and Evidence," *Quarterly Journal of Economics*, 119, 353-402.

— (forthcoming): "Paying Not to Go to the Gym," *American Economic Review*.

DELLAVIGNA, S. AND M. D. PASERMAN (2005): "Job Search and Impatience," *Journal of Labor Economics*, 23, 527-588.

ELIAZ, K. AND R. SPIEGLER (forthcoming): "Contracting with Diversely Naive Agents," *Review of Economic Studies*.

ESTEBAN, S., E. MIYAGAWA, AND M. SHUM (2005): "Nonlinear Pricing with Self-Control Preferences," Unpublished Manuscript, Pennsylvania State University.

FANG, H. AND D. SILVERMAN (2004): "Time-Inconsistency and Welfare Program Participation: Evidence from the NLSY," Unpublished Manuscript, Yale University and University of Michigan.

FREDERICK, S., G. LOEWENSTEIN, AND T. O'DONOGHUE (2002): "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40, 351-401.

GUL, F. AND W. PESENDORFER (2001): "Temptation and Self-Control," *Econometrica*, 69, 1403-1435.

GRUBER, J. AND B. KOSZEGI (2001): "Is Addiction 'Rational'? Theory and Evidence," *Quarterly Journal of Economics*, 116, 1261-1303.

— (2004): "Tax Incidence When Individuals are Time-Inconsistent: The Case of Cigarette Excise Taxes," *Journal of Public Economics*, 88, 1959-1987.

GRUBER, J. AND S. MULLAINATHAN (2005): "Do Cigarette Taxes Make Smokers Happier?" *B.E. Journals: Advances in Economic Analysis & Policy*, 5, Article 4.

HENRIQUES, D. B. (2004): "Seeking Quick Loans, Soldiers Race Into High-Interest Traps," *New York Times*, December 7, 2004.

HERRNSTEIN, R., G. LOEWENSTEIN, D. PRELEC, AND W. VAUGHAN, JR. (1993): "Utility Maximization and Melioration: Internalities in Individual Choice," *Journal of Behavioral Decision Making*, 6, 149-185.

KOSZEGI, B. (forthcoming): "On the Feasibility of Market Solutions to Self-Control Problems," *Swedish Economic Policy Review*.

LAIBSON, D. (1997): "Hyperbolic Discounting and Golden Eggs," *Quarterly Journal of Economics*, 112, 443-477.

— (1998): "Life-Cycle Consumption and Hyperbolic Discount Functions," *European Economic Review*, 42, 861-871.

LAIBSON, D., A. REPETTO, AND J. TOBACMAN (1998): "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity*, 1, 91-196.

— (2005): "Estimating Discount Functions from Lifecycle Consumption Choices," Unpublished Manuscript, Harvard University.

LOEWENSTEIN, G. AND D. PRELEC (1992): "Anomalies in Intertemporal Choice: Evidence and an Interpretation," *Quarterly Journal of Economics*, 107, 573-597.

LOEWENSTEIN, G., T. O'DONOGHUE, AND M. RABIN (2003): "Projection Bias in Predicting Future Utility," *Quarterly Journal of Economics*, 118, 1209-1248.

O'DONOGHUE, T. AND M. RABIN (1999*a*): "Doing It Now or Later," *American Economic Review*, 89, 103-124.

— (1999*b*): "Incentives for Procrastinators," *Quarterly Journal of Economics*, 114, 769-816.

— (1999*c*): "Procrastination in Preparing for Retirement," in *Behavioral Dimensions of Retirement Economics*, ed. by H. Aaron. Washington DC and New York: Brookings Institution Press and Russell Sage Foundation, 125-156.

— (2001): "Choice and Procrastination," *Quarterly Journal of Economics*, 116, 121-160.

— (2003): "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes," *American Economic Review (Papers and Proceedings)*, 93, 186-191.

— (2005): "Optimal Sin Taxes," Unpublished Manuscript, Cornell University.

PASERMAN, M. D. (2004): "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation," Unpublished Manuscript, Hebrew University.

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

POLLAK, R. A. (1968): "Consistent Planning," *Review of Economic Studies*, 35, 201-208.

SCHELLING, T. C. (1978): "Egonomics, or the Art of Self-Management," *American Economic Review*, 68, 290-294.

— (1984): "Self-Command in Practice, in Policy, and in a Theory of Rational Choice," *American Economic Review*, 74, 1-11.

SHAPIRO, J. M. (2005): "Is There a Daily Discount Rate? Evidence from the Food Stamp Nutrition Cycle," *Journal of Public Economics*, 89, 303-325

SHUI, H. AND L. M. AUSUBEL (2004): "Time Inconsistency in the Credit Market," Unpublished Manuscript, University of Maryland.

STEPHENS INCORPORATED INVESTMENT BANKERS (2004): "Stephens Payday Loan Industry Report, 2004."

STROTZ, R. H. (1956): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165-180.

SUNSTEIN, C. AND R. H. THALER (2003*a*): "Libertarian Paternalism," *American Economic Review (Papers and Proceedings)*, 93, 175-179.

— (2003*b*): "Libertarian Paternalism Is Not An Oxymoron," *University of Chicago Law Review*, 70, 1159-1202.

THALER, R. H. (1981): "Some Empirical Evidence on Dynamic Inconsistency," *Economic Letters*, 8, 201-207.

THALER, R. H. AND S. BENARTZI (2004): "Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving," *Journal of Political Economy*, 112, 164-187.

THALER, R. H. AND G. LOEWENSTEIN (1989): "Anomalies: Intertemporal Choice," *Journal of Economic Perspectives*, 3, 181-193.

THALER, R. H. AND H. M. SHEFRIN (1981): "An Economic Theory of Self-Control," *Journal of Political Economy*, 89, 392-406.

WERTENBROCH, K. (1998): "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice," *Marketing Science*, 17, 317-337.