

Willpower and the Optimal Control of Visceral Urges*

Emre Ozdenoren[†] Stephen Salant[†] Dan Silverman[†]

This Version: January, 2008

Abstract

Common intuition and experimental psychology suggest that the ability to self-regulate (“willpower”) is a depletable resource. We investigate the behavior of an agent with limited willpower who optimally consumes an endowment of a tempting consumption good or “cake” over time. We assume that restraining consumption of cake below what maximizes immediate gratification requires willpower. Any willpower left over may be valuable in controlling other urges. Willpower thus links otherwise unrelated behaviors requiring self-control. Our model predicts domain-specific time preference and a linkage between expenditure on nontempting goods and the value of self control in regulating other activities. It also shows that short-sighted behavior by the poor may be attributable to their lack of wealth rather than their tastes or skills. Unlike someone with perfect self-control, an agent with limited willpower will almost never perfectly smooth his consumption, even when it is feasible to do so. Our model provides an explanation for negative time preference, a phenomenon usually attributed to direct utility from beliefs (anticipation). At the same time, it provides a novel explanation for self-commitment, intertemporal preference reversals, and procrastination.

*We thank Roy Baumeister, Gérard Gaudet and Miles Kimball for their important contributions to this paper. We also thank Roland Benabou, Doug Bernheim, Drew Fudenberg, Kai-Uwe Kuhn, Yoram Halevy, David Laibson, David Levine, Robert Mendelsohn, Muriel Niederle, Andrew Postlewaite, Matthew Rabin, Lones Smith, Joel Sobel, Kathleen Vohs, Itzhak Zilcha, and participants in seminars at Harvard, IAS/Princeton, SITE, Calgary, Central Florida, Michigan, Milano-Bicocca, Montreal, Penn and Pittsburgh for their helpful comments on a previous draft. Silverman gratefully acknowledges the support of the Institute for Advanced Study.

[†]Department of Economics, University of Michigan, 611 Tappan St., Ann Arbor, MI 48109-1220.

1 Introduction

Common patterns of intertemporal choice that are inconsistent with standard models have inspired a literature on the economics of self-control. Tendencies to seek self-commitment, to procrastinate, and to seize small immediate rewards despite their important future costs, have motivated studies of quasi-hyperbolic time discounting (e.g., Laibson 1997, O’Donoghue and Rabin, 1999), temptation costs (e.g., Gul and Pesendorfer 2001, 2004), and conflicts between selves or systems (e.g., Thaler and Shefrin, 1988; Bernheim and Rangel, 2004; Fudenberg and Levine, 2006). These models are consistent with a great deal of experimental evidence, and have been fruitfully applied to a number of economic problems ranging from portfolio choice to labor supply to health investment.

Our paper is focused on an aspect of self control that has received less attention in the literature: the *variation* in an individual’s exercise of self restraint across time and circumstances. Specifically, we are concerned with anecdotal and laboratory evidence that the exercise of self control at one time, or in one domain, leaves a person less able or willing to exert self-restraint at another time, or in another domain. This evidence suggests that the exercise of self control draws on a limited and fungible cognitive resource – a resource that is often called *willpower*.

Anecdotes consistent with limited willpower are common. Many who resist unhealthy food and fruitless websurfing all day, and who might prefer to go to bed early after a light dinner, find themselves up late, watching T.V. and gorging on junk food. Dieters often maintain their discipline for short periods but find such self-restraint unsustainable over the long term. Profligate spending or drinking to excess is frequently the “reward” for a hard week at work or home.

Experimental psychology (Baumeister *et al.*, 1994; Baumeister and Vohs, 2003) has gone beyond such anecdotes and has demonstrated that individuals depleted by prior acts of self-restraint tend to behave later as if they have less capacity for self-control. The typical experiment has two phases. Every subject participates in the second phase but only a randomly chosen subset participates in the first, with the remainder serving as a control group. In the first phase, subjects are asked to perform a task that is meant to deplete their willpower; in the second phase, their endurance in an unrelated activity requiring self control is measured.¹ Subjects who participate in the first phase display substantially less endurance in the second phase. This apparent link between the exercise of self-control in one activity and later self-discipline in another has been observed repeatedly, with many different manipulations and measures of self-regulation (Baumeister and Vohs, 2003, Vohs

¹For example, in the first phase subjects have been asked to resist eating tempting foods, to resist drinking liquids when thirsty, and to inhibit automated/habitual behaviors such as reading the subtitles of film or “reading” the color of the ink that a word is written in, rather than the word itself.

and Faber, 2004). Two experiments (Vohs and Faber, 2004 and Dewitte et al., 2005) show that willpower depletion and prior cognitive loads affect subsequent *economic* behavior.² Individual experiments have weaknesses, but we view the collection of these experimental findings as reinforcing the intuitive notion of willpower as a depletable cognitive resource.

In this paper we investigate how willpower limitations affect economic behavior. To do so, we add a willpower constraint to the simplest model of intertemporal choice: the canonical cake-eating problem. In our formulation, moderating consumption of tempting goods requires willpower; the greater restraint the consumer exercises, the faster his willpower erodes.³ We also take account of other activities besides intertemporal saving that require willpower: cramming for exams, training for performances or competitions, maintaining a diet, or preparing an important presentation.

While our model of limited willpower is simple, it generates a rich set of implications. We find that a willpower-constrained consumer would behave in ways that seem inconsistent with having a single rate of time discount. He might, for example, exhibit considerable patience when it comes to work effort and, at the same time, appear myopic when it comes to the long-term consequences of food and drink. His time preference would thus appear *domain-specific*.

In addition, a willpower-constrained consumer who is poor might appear more impatient than his richer counterpart even though they are identical in every respect besides their initial wealth. The poorer consumer would deplete his paycheck sooner and might also devote less willpower to resisting other temptations. Thus, behavioral differences between rich and poor that are often attributed to differences in rates of time discount or self-control skills may also be due simply to differences in wealth.

When a consumer is willpower-constrained, even his demand for conventional (i.e. nontempting) goods will be affected. If, for example, the value of willpower allocated to regulating alternative activities increases, the consumer will use less willpower to limit his intake of tempting goods and, as he spends more on temptations, spend less on conventional goods.

²In Vohs and Faber (2004), subjects who were willpower-depleted purchased a wider assortment of merchandise and spent a larger portion of their experimental earnings than the control group. The experiments of Dewitte et al. should be distinguished from research that showed the effects of cognitive loads on *contemporaneous* choices. In these latter studies, respondents were more likely to choose cake over fruit (Shiv and Fedorikin, 1999) or a smaller earlier reward over a larger later one (Hinson, et al. 2003) when *simultaneously* asked to perform a memory task. In contrast, Dewitte et al. (2005) found that even when such memory tasks were performed prior to the consumption decision, they affected the choice of how much candy to eat.

³Thus, we mainly treat willpower as a resource that is depleted by resisting tempting consumption. Experiments indicate that willpower may also be drained by other activities such as inhibiting automatic responses or bearing cognitive loads. Our model accommodates these causes of depletion either as anticipated uses of willpower that come after consumption choices, or as unanticipated shocks to the willpower stock.

Finally, limited willpower has important consequences for preferences over the timing of consumption. The hallmark of the canonical model, intertemporal smoothing, virtually disappears when a willpower constraint is introduced. Even when the agent has enough willpower to smooth consumption perfectly over the entire time horizon, he is almost always better off foregoing this option and using more willpower to regulate alternative activities.⁴ Indeed, as we show, under a willpower constraint it may even be optimal to consume at an *increasing* rate. Others have attributed such negative time preference to direct utility from beliefs (anticipation). Our analysis suggests it may also reflect self-control problems.

The model presented here differs from the bulk of the self-control literature by accommodating an effect of previous acts of self-restraint on current behavior that requires self-control. Models of hyperbolic discounting (Laibson 1997, O’Donoghue and Rabin, 1999) or temptation costs (Gul and Pesendorfer 2001), for example, would predict that, so long as the agent faces the same choice set, the history of his previous choices would not matter. Models of addiction, habituation, or cues, (Laibson, 2001, Gruber and Köszegi, 2001, Gul and Pesendorfer, 2007, Bernheim and Rangel, 2004) allow the history of choices to affect behavior, but in these models the consumption of a good today affects later tastes for *the same good* whereas in our model, the willpower constraint links (induced) preferences across all forms of tempting consumption. The model in Benabou and Tirole (2004) also links past behavior to future choices as agents with hyperbolic discounting attempt to develop self-reputations for being patient.⁵ However, in Benabou and Tirole (2004)—unlike our model—the successful exercise of self-control in one domain *increases* the likelihood of later self-control. We defer further discussion of the relation of our model to other particularly relevant contributions of the self-control literature (Fudenberg and Levine (2006), and Loewenstein and O’Donoghue (2005)) until the implications of our model have been presented.

The rest of the paper proceeds as follows. In section 2 we present the primary elements of our model. In the first part of that section, we focus on a particularly simple case of our model in order to highlight the novel linkages in economic behavior induced by a willpower constraint. Because the willpower constraint does not induce time preference in this case, we can show many consequences of limited willpower using elementary calculus. In the remainder of the section, we consider the

⁴Allowing willpower to have alternative uses has other implications. For example, using the typical two-stage protocol, Muraven (1998) finds that a given first-stage depletion activity has less of an effect on second stage self-control when subjects are paid more for exerting that latter control. This finding is consistent with a model where willpower has alternative uses. If we fix the marginal value of willpower remaining after the experiment, and increase the incentive for exerting self-control in the second phase activity, then the agent will optimally reallocate willpower to that second phase self-control activity and leave less in reserve for self-regulation after the experiment.

⁵Agents build self-reputations because the model allows uncertainty about preferences and imperfect recall of both preferences and choices.

more general problem and discuss the rich patterns of time preference that the willpower constraint can induce. In section 3 we discuss further how our theory relates to other models of self control. In section 4 we consider two extensions of our analysis: to the case where an individual is uncertain about the size of his willpower stock and to the case where the exercise of self control depletes willpower over the short term but builds it over the longer term. In section 5 we conclude.

2 A Model of Limited Willpower

To investigate the consequences of limited willpower in the simplest intertemporal model, we consider the canonical cake-eating problem where a consumer chooses his consumption path $c(t)$ to maximize his discounted utility $U(c(t))$ over a finite horizon. We denote the size of the cake at time t as $R(t)$ and assume that $R(0)$ is given. There is no return to saving in the model; thus at time t the cake declines at rate $c(t)$.

We depart from the canonical model by assuming that the cake is tempting and that the agent, initially endowed with willpower stock $W(0)$, depletes his willpower when he restrains his consumption. We denote the rate of willpower depletion at time t as $f(W(t), c(t))$. Thus depletion may depend both on the rate of consumption (self restraint) and on remaining willpower reserves.

We assume that instantaneous utility is maximized at a consumption rate \bar{c} and no willpower is used to consume at that rate. To restrain consumption below that rate, however, does require willpower. We assume that the more the agent restrains his consumption, the faster he depletes his willpower reserves. We also allow that the same act of self-restraint may result in (weakly) faster depletion of willpower if the agent's reserves of willpower are lower. Formally, we assume that f is strictly decreasing ($f_c < 0$) and weakly convex ($f_{cc} \geq 0$) in c and weakly decreasing in W ($f_W \leq 0$).⁶ We assume $U(0) = 0$ and, for $c \in (0, \bar{c})$, that utility is strictly concave in c , strictly increasing and twice differentiable. Our focus is on non-addictive goods like those captured in experiments by Baumeister and colleagues, so we assume that $U(\cdot)$ is a function only of contemporaneous consumption.

An important feature of willpower depletion is that as long as even a morsel of tempting cake remains, restraining consumption requires willpower; but when there is nothing left to eat, resisting eating it requires no willpower. We refer to this feature of the willpower technology as the “fundamental discontinuity of willpower depletion” and take account of it in our formulation.⁷

⁶For completeness, we assume that the cross effects are positive, $f_{cW} \geq 0$

⁷This discontinuity would persist even if $f(\cdot)$ were a function of $R(t)$ so long as $\lim_{R(t) \rightarrow 0} f(\cdot) > 0$. It would simplify our analysis if the discontinuity disappeared, but there is no reason to expect that it does.

We assume that any willpower remaining after moderating intertemporal consumption is applied to restrain urges in alternative activities. This willpower “bequest” does not generate utility directly. But it does alter the feasible set of choices in the alternative activities and therefore does *indirectly* determine the utility generated from the alternative activities. We denote this indirect utility function as $m(\cdot)$.

The agent cannot choose a consumption path which results in negative willpower. From the set of consumption paths which maintain nonnegative stocks of both willpower and cake, the consumer chooses the one which maximizes the sum of his discounted utility of consumption and the value of the willpower bequeathed to the regulation of his alternative activities.⁸

If a consumer with initial willpower \bar{W} and initial cake \bar{R} chooses his consumption path $c(t) \geq 0$ optimally, he maximizes:

$$\begin{aligned}
 V(0) &= \int_0^T e^{-\rho t} U[c(t)] dt + e^{-\rho s} m(W(s)) \\
 \text{subject to } \dot{R}(t) &= -c(t) & (P1) \\
 \dot{W}(t) &= \begin{cases} -f(W(t), c(t)) & \text{if } R(t) > 0 \\ 0, & \text{otherwise} \end{cases} \\
 R(T) &\geq 0, \quad W(T) \geq 0 \\
 R(0) &= \bar{R} \geq 0 \\
 W(0) &= \bar{W} \geq 0,
 \end{aligned}$$

where ρ is the subjective rate of time discount and $s = \sup \{t \in [0, T] : R(t) > 0\}$.⁹ The function describing changes in the stock of willpower jumps when the cake is exhausted, reflecting the fundamental discontinuity of willpower depletion discussed above.

As discussed below, the inclusion of a willpower constraint affects consumer behavior in two ways: it induces linkages between behaviors in different domains and it can induce time preference in consumption. To study these effects separately, we investigate first a particularly simple case where willpower induces linkages but not time preference and then move on to the more general

⁸As the various metrics used in experiments reflect, willpower may be measured in many ways. A natural concern is that the model’s predictions about optimal consumption depend on how willpower is measured. However, the optimal consumption path does not depend on how willpower is measured. The reason is simple. Every consumption path that is feasible when willpower is measured in one way is also feasible when it is measured in another. Since utility depends only on the consumption path chosen, the same consumption path is optimal regardless of how willpower is measured. We thank Doug Bernheim for raising this concern about measurement.

⁹We have assumed, for convenience, that the value of the willpower bequest to alternative activities is realized either when the cake is exhausted (s) or at the end of the horizon (T), whichever comes first. This assumption is not material to our results.

case where willpower concerns can also induce time preference.¹⁰

2.1 Induced Linkages

We begin our analysis by studying optimal choice in the particularly simple case where the rate of willpower depletion depends only on the level of consumption. Formally, we assume that $f(W(t), c(t)) = g(c(t))$.¹¹ Moreover, we assume that additional willpower devoted to alternative activities has a constant marginal value ($m'(W(t)) \equiv m$). To isolate the effects of limited willpower, we assume there is no time discounting ($\rho = 0$).

Under these assumptions, it is optimal to consume the tempting good at a constant rate as long as cake remains. To see why this is true, suppose the contrary. That is, suppose consumption were larger at one time than at another when cake remains. Because utility is strictly concave in tempting consumption, marginally reducing the larger consumption while marginally increasing the smaller one by an offsetting amount would strictly increase utility from consumption without violating the cake constraint. Whether this perturbation strictly increases overall utility depends on its effect on willpower. Recall, however, that willpower depletion is independent of the willpower stock and is a decreasing, convex function of consumption. It follows that the proposed perturbation will actually make more willpower available for use elsewhere: increasing the lower level of consumption releases more willpower than is depleted by reducing the larger consumption by an offsetting amount. At all other times, when consumption is left unchanged, willpower depletion is unchanged (because $f_W = 0$). Thus the perturbation releases additional willpower to be used in the alternative activities and thereby generates another gain in utility. So if the original program was feasible, the perturbed program will also be feasible but it will yield higher utility. Hence, consumption must be constant as long as cake remains.¹²

The preceding argument shows that consumption is constant when it is positive. In principle, however, cake may be exhausted prior to T , after which consumption jumps to zero. A perturbation that further smooths consumption in this situation has a decidedly different effect than the one described above. In this situation, marginally reducing consumption in the phase when it is positive

¹⁰Because the willpower depletion technology is discontinuous, problem (P1) is somewhat nonstandard. Therefore, in each investigation, we will solve problem (P1) by analyzing a simpler, conventional problem with the same solution. Details regarding the equivalent problem are provided below in Section 2.2.

¹¹Two prominent special cases of this formulation are $g(c(t)) = k \cdot (\bar{c} - c(t))$ and $g(c(t)) = k \cdot (U(\bar{c}) - U(c(t)))$, where \bar{c} is the lowest rate of consumption that maximizes instantaneous utility. In the former case, the rate of willpower depletion is proportional to the difference between the agent's ideal and his actual rate of consumption; in the latter case, it is proportional to the difference between the utilities generated by these two rates of consumption.

¹²See footnote 19 for a more formal proof of these arguments.

and marginally increasing it in the phase when it is zero depletes willpower in *both phases*. This occurs because of the fundamental discontinuity of willpower depletion: when no cake is available no willpower is depleted, but once cake becomes available in the second phase, consuming a small amount requires additional willpower. As the analysis below clarifies, it is not always optimal to stretch the phase of constant consumption to the end of the horizon. Intuitively, when the return to alternative uses of willpower is high, it will be better to complete consumption before time T and save more willpower for other purposes. We will call the situation when the agent chooses to stretch consumption to time T the “perfect-smoothing regime,” and call the alternative the “no-smoothing regime.”

2.1.1 Simple Comparative-Statics

In the rest of this section, we assume that the agent has enough cake or willpower to implement perfect smoothing if he so chooses.¹³ Then the payoff from consuming at constant rate c and exhausting the cake at time $s \leq T$ is $sU(c) + m[\bar{W} - sg(c)]$, where $s = \frac{\bar{R}}{c}$ and $s \in [0, T]$. Provisionally substituting $\frac{\bar{R}}{s}$ for c we obtain a continuous, strictly concave objective function of one variable, s . The agent’s problem is then equivalent to choosing the time to exhaust the cake, s , to solve:

$$\max_s sU\left(\frac{\bar{R}}{s}\right) + m[\bar{W} - sg\left(\frac{\bar{R}}{s}\right)]$$

It is straightforward to verify that an optimum exists and that, since the maximand is strictly concave, any solution to the first-order condition is optimal. The optimum never occurs at $s = 0$ but may occur at $s < T$ or at $s = T$. We consider each case in turn.

No-smoothing Regime When no smoothing is optimal, the first order condition implies:

$$U\left(\frac{\bar{R}}{s}\right) - U'\left(\frac{\bar{R}}{s}\right)\frac{\bar{R}}{s} - mg\left(\frac{\bar{R}}{s}\right) + \frac{\bar{R}}{s}mg'\left(\frac{\bar{R}}{s}\right) = 0. \quad (1)$$

The objective function is strictly concave, so the left-hand side of (1) is strictly decreasing in s . If it did not decline to zero for $s < T$, then perfect smoothing ($s = T$) would be optimal.

To determine the rate of consumption in the no-smoothing regime, we rewrite the first-order condition in terms of the constant consumption rate c :

$$U(c) - cU'(c) - mg(c) + cmg'(c) = 0. \quad (2)$$

Notice that the rate of consumption in the no-smoothing regime is determined entirely by the value of additional willpower in the alternative activities (m). Hence, if an agent in this regime

¹³That is, $\bar{W} - Tg\left(\frac{\bar{R}}{T}\right) > 0$. The case where he lacks enough cake or willpower to implement perfect smoothing is equally tractable but less interesting.

began with larger initial willpower, he would optimally consume at the same rate and, with a cake of unchanged size, would exhaust it at an unchanged time (s), bequeathing all of the additional willpower to the alternative activities.

Turning next to changes in initial cake sizes, suppose someone in the no-smoothing regime began with a marginally larger initial cake. Then his rate of consumption would not change (since m has not changed); but since he would take longer to exhaust his larger cake later (s increases) he would have less willpower left for the alternative activities.

Now consider changes in the value (m) of willpower applied to alternative activities. Suppose someone in the no-smoothing regime began with a marginally higher m . It is straightforward to verify that the rate of consumption is a strictly increasing function of m .¹⁴ Hence, the optimal rate of consumption would increase in response and, given that the cake size is unchanged, it would be exhausted sooner.

Perfect-smoothing Regime In contrast, when perfect-smoothing is optimal, no behavior changes in response to an exogenous increase in the value of willpower in the alternative activities (m). For, perfect smoothing requires that the cake still be exhausted at T and, given a cake of unchanged size, the consumer would devour it at the same rate and would therefore have the same amount of willpower to bequeath. If someone who would perfectly smooth received marginally more willpower, he would not alter his consumption path or date of exhaustion but would simply use the additional willpower in the alternative activities. If such a person instead began with a larger cake, then he would consume it at a faster rate, exhausting it at an unchanged date ($s = T$), and he would have more willpower left over for the alternative activities. We summarize the comparative-static results for each regime in Table 1 and then turn to the behavioral implications of these results.

2.1.2 Smoothing Consumption May Not Be Optimal

Consumption smoothing is the hallmark of the standard model with stationary utility and no discounting. However, it does not survive the addition of willpower concerns when willpower is needed not only to resist tempting consumption but also to regulate other urges. These urges may include the temptation to slack off at work or school, the urge to participate in risky financial or sexual behaviors, or the urge to express anger or jealousy to co-workers, friends or family. Suppose the return to the alternative activities is high enough, so that the agent is in the no-

¹⁴Totally differentiating the first-order condition determining c in the no-smoothing regime, we conclude that $dc/dm = (g(c) - cg'(c)) / (c[mg''(c) - U''(c)]) > 0$, where the sign follows since $g(\cdot)$ is positive, strictly decreasing, and strictly convex while U is strictly concave.

	c	s	$W(s)$
m	+	-	+
\bar{R}	0	+	-
\bar{W}	0	0	+

No Smoothing

	c	s	$W(s)$
m	0	0	0
\bar{R}	+	0	+
\bar{W}	0	0	+

Perfect Smoothing

Table 1: **Comparative-Static Results in Each Regime.** Sign of changes in optimal consumption (c), time to exhaust the cake (s), and willpower bequest $W(s)$, resulting from changes in the return to willpower applied to alternative activities (m), changes in initial cake size (\bar{R}), or changes in willpower (\bar{W}).

smoothing regime.¹⁵ If this agent had more willpower, he would use no more to restrain tempting consumption; instead he would use the additional willpower to restrain urges in other activities. (See the bottom rows in Table 1.) Therefore, if m is high enough, the agent would continue to avoid perfect smoothing of consumption no matter how large is his initial stock of willpower.

2.1.3 Domain-Specific Time Preference

This simple model also shows how, for an agent with limited willpower, time preference may differ sharply by decision domain. The individual may appear willing or able to postpone gratification in one set of activities and, at the same time, profoundly myopic about choices in another, even though in fact he discounts time at a single rate (assumed to be zero here). Consider the following concrete example. Suppose willpower is used both to regulate consumption and to exert concentrated effort on dull but professionally important tasks at work – tasks which provide only longer-term rewards. If the return (m) to more concentrated effort at work is high enough so that the agent is in the no smoothing regime, an increase in m will result in the agent consuming his tempting cake faster, exhausting it more quickly, and therefore having more willpower for work related activities. (See Figure 1 and the top rows in Table 1.) Hence, the agent would appear at the same time more myopic when it comes to consumption choices, but more forward-looking when it comes to work.

Short-sighted consumption behavior and diligence at work might also occur if the individual were truly myopic but simply enjoyed work. However, this explanation can be distinguished from the effects of willpower. In our model as the return (m) to additional willpower allocated to work

¹⁵To see which regime is optimal for given parameters m and \bar{R} , set $s = T$ in the left hand side of (1). If it is nonnegative, then perfect smoothing is optimal; otherwise no-smoothing is optimal. Notice that, for any cake size, no smoothing is optimal for any m above a critical level.

increases, the individual would appear increasingly impatient in his consumption and increasingly disciplined in his regulation of alternative activities. In contrast, an increase in returns to effort at work would not affect short-run time preference for consumption in a model that attributed such domain-specific time preference to tastes for work or to domain-specific rates of time discount. Thus, the limited-willpower model makes distinctive predictions about changes in self-restraint exercised in one activity as a function of the return to exerting self-restraint in another.

2.1.4 Poverty Affects Patience

Willpower concerns also induce a distinctive linkage between poverty and patience. In the standard view, poverty is partially caused by impatience: impatient agents are less willing to postpone gratification and invest in human and financial capital, and thus they accumulate less wealth.¹⁶

The willpower model suggests another relationship between poverty and patience. Consider two agents with the same willpower reserves and self-control technologies, but one sufficiently poor (has a small enough cake) that no smoothing is optimal and the other sufficiently rich such that perfect smoothing is optimal. The poorer one may both finish his “cake” sooner and have less willpower left over for alternative activities than his richer counterpart. (See Figure 2.) The poor man thus appears to lack self-control when compared with his wealthier counterpart; but the difference in his behavior is attributable entirely to his smaller paycheck. While impatience may lead to poverty as present-oriented consumers forego opportunities to invest in wealth or human capital, the willpower model suggests that the poverty of the poor may also make them impatient.

2.1.5 Willpower Concerns Even Affect Consumption of Nontempting Goods

So far, we have assumed that all consumption is tempting. In this subsection, we consider the case where some goods, such as deodorant, oil changes, or plain yogurt, pose no meaningful temptation. Although limiting consumption of such “nontempting goods” might require no willpower, demand for them is nonetheless influenced by willpower concerns. For example, suppose a worker is limiting his junk food and entertainment consumption over a fixed horizon while applying moderate effort to a project at work. He then learns that an influential manager has shown an interest in the project and will closely monitor various aspects of it – an unusual opportunity to impress or disappoint his superior. In response to this news, he will allocate more of his willpower budget to work and less to restricting his consumption of junk food and his drinking at happy hours. To finance the additional consumption of these tempting goods, however, he would spend less of his

¹⁶See, e.g., Lawrance (1991).

budget on nontempting consumption. Thus, news of management’s interest in the project would cause spending on nontempting goods to fall.

To see this formally, suppose that a fixed sum of money (I) can be spent either on goods requiring willpower (intertemporal consumption) or on the consumption of nontempting goods. Suppose R dollars is spent on intertemporal consumption of the tempting good (adjust units so that one unit of cake costs one dollar) and R_{CG} is spent on the set of conventional goods: $R + R_{CG} \leq I$. Assume that spending R_{CG} on the conventional goods generates utility $V(R_{CG})$, where the utility function is strictly increasing and strictly concave.

Suppose that the agent is smoothing perfectly. We denote the agent’s payoff from activities requiring willpower by $J(R, m) \equiv TU(R/T) + m(\bar{W} - Tg(R/T))$. Since U is strictly concave and g is strictly convex, J is strictly concave in R . The agent will choose R to maximize $V(I - R) + J(R, m)$. At the optimum a dollar allocated to consumption of the tempting good and the nontempting good must have the same value at the margin, that is, $-V'(I - R) + J_R(R, m) = 0$. Since both V and J are strictly concave in R , the second-order condition also holds; that is, $V''(I - R) + J_{RR}(R, m) < 0$. To see how spending on the tempting good varies with the return to willpower used in the alternative activities, we take the total derivative of the first-order condition with respect to m to obtain:

$$\frac{dR}{dm} = \frac{-J_{Rm}(R, m)}{V''(I - R) + J_{RR}(R, m)}.$$

Since $J_{Rm}(R, m) = -g'(R/T) > 0$, we conclude $dR/dm > 0$.¹⁷ Thus when the return to exercising self-control in alternative activities increases, the consumer devotes more of his budget to tempting consumption leaving less for nontempting goods.

2.2 Induced Time Preference

Besides inducing linkages between otherwise unrelated behaviors, the introduction of a willpower constraint may also induce time preference in consumption of the tempting good. To investigate this aspect of behavior, we now permit the rate of willpower depletion to depend on the stock of willpower remaining ($f_w < 0$). Analysis of this case requires a more complete examination of the consumer’s dynamic optimization problem ($P1$) since consumption need not be constant. As noted above, problem ($P1$) is non-standard because of the discontinuity in the willpower depletion process. To circumvent this difficulty, we examine a related “free endpoint” problem ($P2$), which

¹⁷Using similar arguments we can show that R and m are linked in the no smoothing regime as well.

is standard in its formulation and which has the same solution as problem (P1):

$$\begin{aligned}
V(0) &= \int_0^s e^{-\rho t} U[c(t)] dt + e^{-\rho s} m(W(s)) \\
\text{subject to } \dot{R}(t) &= -c(t) & (P2) \\
\dot{W}(t) &= -f(W(t), c(t)) \\
R(s) &\geq 0, W(s) \geq 0 \\
R(0) &= \bar{R} \geq 0 \\
W(0) &= \bar{W} \geq 0.
\end{aligned}$$

In problem (P2), the agent chooses not only an optimal consumption path but also a date $s \leq T$ beyond which depletion of each stock ceases ($\dot{W}(t) = \dot{R}(t) = 0$ for all $t \in (s, T]$). Problem (P2) therefore allows all of the paths that are feasible in (P1) and evaluates each as (P1) does. However, problem (P2) also allows additional feasible paths: those where willpower depletion ceases even though cake remains ($R(s) > 0$). To establish that the solutions to the two problems are the same, we need merely prove that these additional paths are never optimal in (P2). Intuitively, this is obvious since one can always dominate such paths by consuming marginally more cake in the neighborhood of s . This additional consumption is feasible since (by hypothesis) more cake is available; moreover, consuming it would relax the willpower constraint.¹⁸

Working with problem (P2), then, the Hamiltonian is given by

$$H(c(t), R(t), W(t), \alpha(t), \lambda(t), t) = e^{-\rho t} U(c(t)) - \alpha(t) c(t) - \lambda(t) f(W(t), c(t)) ;$$

we will refer to it as $H(t)$ when no confusion arises. The first-order conditions include:

$$c(t) \geq 0, e^{-\rho t} U'(c(t)) - \alpha(t) - \lambda(t) f_c \leq 0 \text{ and c.s.} \quad (3)$$

$$\dot{W}(t) = -f \quad (4)$$

$$\dot{R}(t) = -c \quad (5)$$

$$\dot{\alpha}(t) = 0 \quad (6)$$

$$\dot{\lambda}(t) = \lambda(t) f_W \quad (7)$$

$$T - s \geq 0, H(s) - \rho e^{-\rho s} m(W(s)) \geq 0 \text{ and c.s.} \quad (8)$$

$$R(s) \geq 0, \alpha(s) \geq 0 \text{ and c.s.} \quad (9)$$

$$W(s) \geq 0, \lambda(s) - m'(W(s)) \geq 0 \text{ and c.s.} \quad (10)$$

¹⁸For a formal proof of the equivalence of P1 and P2, see our Technical Appendix available online at www.umich.edu/~dansilv/research.

Using this dynamic system, we can study the impact of willpower concerns on the agent's time preference of consumption. To isolate the effects of the willpower constraint, we continue to assume that there is no discounting ($\rho = 0$).

2.2.1 The Time Path of Consumption

To begin, consider the intuitive tradeoffs revealed by the first-order conditions characterizing optimal consumption at t . When consumption is strictly positive, we can rewrite condition (3) as:

$$\underbrace{U'(c(t))}_{\text{direct marginal benefit}} + \underbrace{\lambda(t)(-f_c)}_{\text{indirect marginal benefit}} = \underbrace{\alpha(t)}_{\text{marginal cost}} . \quad (11)$$

Grouped in this way, condition (3) shows that consuming at a slightly faster rate at time t generates two marginal benefits and one marginal cost. The *direct* marginal benefit ($U'(c(t))$) is the usual one: an increase in utility at time t (expressed in utils at $t = 0$) that results from consuming more at time t . Similarly, the marginal cost of consuming at a faster rate at t ($\alpha(t)$) is also the standard one; it reflects the utility lost because the additional cake consumed at t can no longer be consumed some other time. What is distinctive is that increasing consumption also has an *indirect* marginal benefit ($-f_c\lambda(t)$). This indirect benefit derives from the fact that, by increasing the rate of consumption, the agent depletes willpower at a slower rate ($-f_c$) and each unit of willpower saved at time t is worth $\lambda(t)$ utils. At an interior optimum, the sum of the two marginal benefits must equal the marginal cost.

The optimality conditions also reveal an important difference between the two depletable resources, cake and willpower, when the rate of depletion depends on the stock of willpower remaining. Additional cake is equally valuable whenever it arrives since its mere availability, before it is consumed, provides no services and its future arrival can be anticipated by consuming more in advance. That explains why the shadow value of additional cake is constant over time ($\dot{\alpha}(t) = 0$). In contrast, when $f_W < 0$ and the willpower constraint is binding, additional willpower is more valuable the earlier it arrives because its mere presence provides a service: the more willpower one has available at t , the less must be depleted over a short interval to restrain consumption by a given amount. This is not a service upon which one can draw simply by recognizing that more willpower is available in the future. For this reason, the shadow value of willpower declines with time ($\dot{\lambda}(t) < 0$ from (7) whenever $f_W < 0$). This distinction between the two depletable resources disappears when $f_W = 0$. In that case, additional willpower—like additional cake—is equally valuable whenever it arrives and, as we have seen, it is optimal to consume at a constant rate until the cake is exhausted at $s \leq T$.¹⁹

¹⁹To see that $c(t)$ must be constant for $t \leq s$, note that (4) implies that $\alpha(t)$ is constant. If it is assumed that

The Possibility of Negative Time Preference As we showed in subsection 2.1, when willpower depletion depends only on the consumption level ($f_W = 0$), it is optimal to consume at a constant rate as long as cake remains. More generally, however, limited willpower may induce *time-varying* paths of consumption. Indeed, willpower concerns may result in increasing paths of consumption (negative time preference).

To see the sources of time preference in the willpower model, note that the *sum* of the direct and indirect marginal benefits of increased consumption defined in (11) cannot vary over time, since (6) insures that $\alpha(t)$ is constant. Hence, whenever the path of the indirect marginal benefit of increased consumption is increasing the direct marginal benefit path must be decreasing. Given the concavity of the utility function, this means that, in these circumstances, consumption must be increasing. By a similar argument, a decreasing indirect marginal benefit path requires a decreasing consumption path.

The following proposition characterizes the (local) rate of change in consumption and thus gives a condition for consumption to be locally increasing.

Proposition 1 *Suppose $c(t) > 0$ for $t \leq s$. Then $\dot{c} \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$ as $\lambda \frac{d}{dW} \frac{d}{dc} \ln f(W, c) \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$.*

Proof. Since $\alpha(t)$ is constant over time, differentiating equation (11) with respect to t implies $[U''(c) - \lambda f_{cc}] \dot{c} + \{\dot{\lambda}(-f_c) + \lambda f_{cW}(-\dot{W})\} = 0$. Because $[U''(c) - \lambda f_{cc}] < 0$, $\dot{c} \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$ if and only if $-f_c \lambda \left\{ \frac{f_{cW}}{-f_c}(-\dot{W}) + \frac{\dot{\lambda}}{\lambda} \right\} \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$. Then, using (7) and (4) to replace $\frac{\dot{\lambda}}{\lambda}$ and $-\dot{W}$, we obtain $\dot{c} \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$ if and only if $\lambda \{f f_{cW} - f_c f_W\} \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$. The result then follows because $\lambda \{f f_{cW} - f_c f_W\} \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$ as $\lambda \frac{d}{dW} \frac{d}{dc} \ln f(W, c) \begin{smallmatrix} \geq \\ \leq \end{smallmatrix} 0$. ■

This proposition implies that, when strictly positive, consumption is constant in only three circumstances: (1) when willpower is not scarce ($\lambda = 0$), (2) when the rate of willpower depletion is independent of the stock of willpower remaining ($f_W = 0$), or (3) when, despite neither of the preceding circumstances, the indirect marginal benefit of a marginal increase in consumption is constant. Thus, if willpower is scarce ($\lambda(t) > 0$) and the rate of willpower depletion depends on the stock ($f_W < 0$), the prediction of a constant consumption path is a knife-edge case.

Proposition 1 can also be seen to say that when willpower is scarce, consumption strictly increases whenever the logarithm of the depletion function has a strictly positive cross partial derivative.²⁰ Knowing this, it is straightforward to construct examples where the consumption path will rise or fall monotonically before jumping to zero. These examples provide useful intuition for optimal choice under a willpower constraint.

$f_W \equiv 0$ then $f_c(W(t), c(t))$ does not vary with $W(t)$ and, given (6), $\lambda(t)$ is constant. Since we are assuming that $\rho = 0$, the solution to (3) will be the same for all $t \leq s$.

²⁰The reader may recognize this condition as logsupermodularity of the depletion function.

Consider a willpower depletion function f that is the product two functions, $K(W)$ and $g(c)$. The cross partial of the logarithm of such a function is zero, so the consumption path must be flat prior to the last moment of consumption. If $f_{cW} > 0$, however, simply adding a strictly positive (respectively, strictly negative) constant A to any function of this form generates a strictly increasing (respectively, strictly decreasing) consumption path. To see this formally let

$$f(W, c) = A + K(W)g(c) \text{ with } K' < 0 \text{ and } g' < 0.$$

Here, a given act of will is more depleting when the agent is already depleted. In addition, willpower is either being drained ($A > 0$) or replenished ($A < 0$) at a constant rate.²¹ Note that,

$$\frac{d}{dW} \frac{d}{dc} \ln f(W, c) = \frac{K'(W)g'(c)A}{[A + K(W)g(c)]^2}$$

and thus, by Proposition 1

$$\text{sign}(\dot{c}) = \text{sign}(K'g'A).$$

Since g is strictly decreasing, if $K' < 0$ then $\dot{c} \begin{matrix} \geq \\ \leq \end{matrix} 0$ as $A \begin{matrix} \geq \\ \leq \end{matrix} 0$. So whenever the willpower is scarce, $K' < 0$, and $A > 0$, the path of optimal consumption must increase over time. The increasing path emerges in this example from the interaction of two features. First, because A is positive, willpower will disappear as time elapses; this is a “use it or lose it” situation. However, this feature alone would not generate an increasing path of consumption; if K' were zero, the consumption path would be flat. The consumer would anticipate the shrinking willpower stock and behave, in that circumstance, exactly as he would with a smaller but unshrinking initial willpower stock. An increasing path thus requires the second feature of this example: that willpower depletion depend on the remaining stock of willpower. Because K' is negative, a given act of will is more depleting when the willpower stock is low. So the agent optimally takes the opportunity to exert acts of will when they have a relatively low opportunity cost (when W is large) and before time erodes the willpower stock. This simple example thus shows how limited willpower may induce even *negative* time preference in the absence of time discounting.

A preference for increasing paths of consumption has previously been attributed to utility from beliefs (anticipation). Here we see how problems of self-control may also generate incentives to “save the best for last” or “get the hard part out of the way first.”

²¹This exogenous draining or replenishing of willpower is importantly different from the endogenous depletion we have been modeling all along and different from the endogenous *replenishment* modeled in section 4 where we allow willpower to behave like a muscle – drained by exertion over the short term, but built by exertion over the longer term.

While our model is flexible enough to explain both positive and negative time preference revealed in consumption choices, it can be rejected. The problem of predicting the effects of willpower on consumption profiles is analogous to predicting the effects of price on demand or labor supply; in the absence of additional information, income and substitution effects make the net effect ambiguous; but with sufficient information about preferences, unambiguous predictions can be derived. Here, as Proposition 1 shows, with sufficient information about the willpower depletion technology the model makes strong predictions about the qualitative features of the consumption path. Importantly, choice experiments can provide the requisite information. It is relatively straightforward to design, for example, choice experiments that would identify whether the rate of willpower depletion depends on the stock of remaining willpower.

2.2.2 *Almost Never Smooth Even When Feasible*

In section 2.1, we showed that perfect smoothing, the signature prediction of the standard model, may not occur when $f_W = 0$ and the marginal value of willpower allocated to alternative activities is constant and large enough. In this subsection we derive necessary and sufficient conditions for perfect smoothing to occur under more general assumptions ($f_W \leq 0$ and $m(W)$ either zero or strictly increasing and weakly concave). We find that if an agent lacks an alternative use of willpower, a necessary and sufficient condition for him to smooth perfectly is simply that he have enough willpower initially for such smoothing to be feasible. If, instead, willpower has alternative uses, in *almost* all circumstances he will refrain from perfect smoothing no matter how large his initial stock of willpower.

No Alternative Uses of Willpower We consider first the case where willpower has no alternative uses and thus willpower bequeathed to the future has no marginal value. Since our model coincides with the canonical model when the willpower constraint is slack ($\lambda(t) = 0$), perfect smoothing is optimal whenever feasible.

Proposition 2 *Let W_H be the minimum level of initial willpower such that setting $c(t) = \frac{\bar{R}}{T}$ for $t \in [0, T]$ is feasible. Denote the optimal consumption path as $c^*(t)$. Let $\rho = m(W) = 0$. Then if $\bar{W} \geq W_H$ the cake is exhausted, and $c^*(t) = \frac{\bar{R}}{T}$ for $t \in [0, T]$. If instead $\bar{W} < W_H$ then both the cake and willpower are exhausted ($R(s) = W(s) = 0$), and $\lambda(t) > 0$ for all $t \in [0, s]$.*

Proof. See Appendix. ■

Corollary 1 *If $\rho = m(W) = 0$ and $\bar{W} < W_H$ then when consumption is strictly positive it is strictly increasing (resp. constant, strictly decreasing) if and only if $\frac{d}{dW} \frac{d}{dc} \ln f(W, c)$ is strictly positive (resp. zero, strictly negative).*

Proof. This follows as a consequence of Propositions 2 and 1. ■

Remark 1 (*Procrastination*) *Corollary 1 describes the path of consumption when strictly positive. As we have seen, however, this description does not preclude intervals of zero consumption. Suppose that the agent must spend a fixed amount of time to complete a project, and the cake represents the remaining hours of leisure. Then he may work hard on the project at the beginning, slack off as time goes by, but then cram for some interval just before the deadline.*

When Willpower Has Alternative Uses Proposition 2 shows that if there is no other use of willpower, perfect smoothing is always optimal when feasible; and thus the standard model is nested in ours. We now move on to investigate when perfect smoothing is optimal if willpower has alternative uses. We assume throughout this subsection that devoting more willpower to controlling alternative urges is always useful and find that perfect smoothing is *almost* never optimal even when the cake is exhausted at T .

To see why the optimality of smoothing depends on the alternative uses of willpower, note that if it is optimal to use some willpower to regulate consumption and the remainder to regulate other urges, then it must be that willpower cannot be advantageously re-allocated between the two uses ($\lambda(s) = m'(W(s))$ by condition (10)). However, we have assumed that additional willpower devoted to alternative activities is valuable ($m'(\cdot) > 0$). Therefore $\lambda(s) > 0$. Since this shadow value weakly declines over time (condition (7)) and is strictly positive when consumption ceases, it must also be strictly positive previously. But then our previous result implies that perfect smoothing occurs only if $\frac{d}{dW} \frac{d}{dc} \ln f(W, c) = 0$.

We know from the analysis in subsection 2.1 that even if consumption is flat when positive, perfect smoothing still may not occur. It may be optimal instead to consume at a constant but rapid rate in order to exhaust the cake before T and save more willpower for the alternative activities. Intuitively, this would occur if additional willpower devoted to regulating other urges was sufficiently valuable.

How valuable must willpower be in alternative uses for perfect smoothing to fail even though optimal consumption is flat whenever positive? For perfect smoothing to be optimal it must be that, when $c(t) = \frac{\bar{R}}{T}$ and $s = T$, the marginal return from allocating a bit more willpower to alternative activities, $m'(W(T))$, is weakly less than the marginal value of additional willpower allocated to consumption under perfect smoothing. We show below that this latter marginal value can be bounded above by:

$$m_H(W, c) = (U(c) - U'(c)c) / (f(W, c) - f_c(W, c)c).^{22}$$

²²If no smoothing occurs, then $m_H(W, c) = m'(W(s))$. This generalizes equation (2).

Thus, if we denote by \hat{W} the willpower available at T if perfect smoothing has been implemented, and the rate of consumption under perfect smoothing by c_H , then if $m'(\hat{W}) > m_H(\hat{W}, c_H)$, perfect smoothing never occurs.²³

The following proposition summarizes the necessary conditions derived above and, in addition, provides sufficient conditions for perfect smoothing to be optimal.

Proposition 3 *If additional willpower devoted to alternative activities is valuable ($m'(\cdot) > 0$), and if the agent perfectly smooths consumption then $\frac{d}{dW} \frac{d}{dc} \ln f(W, c) = 0$, for $c = c_H$ and $W \in [\hat{W}, \bar{W}]$, and $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$. Moreover, if $\frac{d}{dW} \frac{d}{dc} \ln f(W, c) = 0$, for all $c \geq c_H > 0$ and $W \in [\hat{W}, \bar{W}]$ and $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$ then perfect smoothing must occur.²⁴*

Proof. We have already proved, in the text, that when $m'(\cdot) > 0$ perfect smoothing requires $\frac{d}{dW} \frac{d}{dc} \ln f(W, c) = 0$, for $c = c_H$ and $W \in [\hat{W}, \bar{W}]$. Thus, to prove the first statement, it remains only to derive that $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$. If any willpower is used to restrain consumption, optimality of this decision requires that $m'(W(s)) \leq \lambda(s)$. Next we find an upper bound on $\lambda(s)$ using the first-order conditions. Condition (3) implies that $U'(c(s)) - \lambda(s) f_c \leq \alpha$. Multiplying both sides by $c(s)$ we get,

$$[U'(c(s)) - \lambda(s) f_c] c(s) \leq \alpha c(s).$$

Condition (8) requires

$$H(s) = U(c(s)) - \alpha c(s) - \lambda(s) f(W(s), c(s)) \geq 0.$$

Combining these two inequalities we see that

$$\lambda(s) \leq (U(c) - U'(c)c) / (f(W, c) - f_c(W, c)c) = m_H(W, c).$$

To prove that the specified conditions are sufficient for perfect smoothing to be optimal, assume the contrary. Since the conditions imply that consumption must be constant when positive, perfect smoothing could only fail because $s < T$. If consumption terminates at $s < T$, then $c = \frac{\bar{R}}{s} > c_H$. Since less restraint will be exercised over a shorter interval $W(s) > \hat{W} > 0$. This in turn has two implications. Since more willpower will be bequeathed to the alternative activities, $m'(\hat{W}) \geq m'(W(s))$. In addition, since $W(s) > 0$, (10) implies that $\lambda(s) = m'(W(s)) > 0$. By hypothesis, $s < T$. Then, as shown above, the first-order conditions imply $\lambda(s) = m_H(W(s), c)$.

²³Clearly \hat{W} depends on the initial levels of willpower and cake but we suppress this dependence for simplicity.

²⁴Recalling the case studied in Section 2.1, where $f(W, c) = g(c)$ and $m'(\cdot) = m$, this condition reduces to $m \leq (U(c_H) - c_H U'(c_H)) / (g(c_H) - c_H g'(c_H))$.

Since $m_H(\cdot, \cdot)$ is weakly increasing in the first argument and strictly increasing in the second argument, $m_H(W(s), c) > m_H(\hat{W}, c_H)$. Given the hypothesis that $s < T$ we therefore conclude that $m'(\hat{W}) \geq m'(W(s)) = \lambda(s) = m_H(W(s), c) > m_H(\hat{W}, c_H)$. But $m'(W(s)) > m_H(\hat{W}, c_H)$ violating $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$. ■

Thus Proposition 3 shows that when willpower has alternative uses, the necessary conditions for perfect smoothing are restrictive, but not impossible to meet. Limited willpower *largely* undoes the signature prediction of the standard theory of intertemporal choice. If willpower is depleted by restricting consumption in the right way, and if the marginal value of willpower allocated to other activities is not too high, then the consumer may still smooth perfectly.

3 Comparison with Alternative Models of Self-Control

The preceding sections were largely focused on how the predictions of a willpower model differ from those of a standard model of intertemporal choice. In this section we compare our willpower model with existing models of self-control.

We note first that limited willpower provides complementary explanations for the hallmark predictions of self-control models: a preference for commitment, profound procrastination, and apparent time-inconsistencies. Consider first a taste for commitment. As in each of the leading models of self-control, an agent in our model would strictly prefer to have his “cake” or paycheck doled out to him by a savings club.²⁵ For if the entire amount were available, resisting spending it would deplete his willpower and willpower has a scarcity value. Next consider profound procrastination. Like other models of self-control, limited willpower also explains extreme forms of procrastination. Suppose, for example, an agent must spend a fixed amount of time on an assignment before a certain deadline, but can allocate the remaining leisure time optimally.²⁶ Then if the agent is willpower constrained, he may enjoy leisure early in the program, and later work non-stop until the deadline.

Depletable willpower is also consistent with apparent intertemporal preference reversals. When

²⁵Benabou and Tirole (2004) is an exception, as in that model commitments may prevent the agent from obtaining valuable information about his willpower.

²⁶This is exactly the setup Fischer (2001) used to investigate procrastination but we have introduced a willpower depletion constraint. Without it, Fischer finds that the discount factor required to explain procrastination in a model with standard time discounting and additively-separable utility is unrealistic. Moreover, with standard time discounting, such procrastination is not something to be avoided; even if they could, agents would not seek commitments to better smooth their work effort. Our model can generate procrastination (zero consumption of leisure for the final phase of the planning horizon) for any discount rate, and agents in our model would, if possible, seek commitments to avoid it.

asked to choose between a smaller reward that will arrive immediately and a larger one that would arrive after some delay, an agent with limited willpower may choose the former. The willpower cost of resisting the immediate, smaller reward may outweigh the utility gain from receiving the later, larger reward. However, if asked to choose now one of these same two options set off in the temporal distance, the consumer may choose the later, larger prize. The choice now of a prize to be delivered in the future allows the consumer to commit irreversibly to an option. He then prefers the later, larger prize since willpower ceases to be required to resist the earlier, smaller prize. In this way a willpower-constrained agent may appear to behave in a time-inconsistent manner.

While it offers complementary explanations for a taste for commitment, preference reversals and profound procrastination, the willpower model also has implications that distinguish it from existing models of self-control.

The distinguishing feature of the limited-willpower model is the linkages that it generates between otherwise unrelated choices. We saw the implications of these linkages in the “no smoothing” results, domain-specific time preference, spillovers of self-control problems into demand for non-tempting goods, and poverty affecting patience. Perhaps most basically, we also saw an implication of these links as consumption of even non-addictive goods in the past had influence on future choices regarding seemingly unrelated goods. Thus, different from models of hyperbolic discounting (Laibson 1997, O’Donoghue and Rabin, 1999), and temptation costs (e.g., Gul and Pesendorfer 2001, 2004), keeping the current choice set constant, in our model the history of past choices matters for behavior.

By linking otherwise unrelated, and non-addictive self-control activities, and by explicitly modeling a willpower reserve, our model can be distinguished from models of cues like Laibson (2001) and Bernheim and Rangel (2004). And even if attention is restricted to the effects of past self-restraint in one activity on future self-restraint in the same activity, there remains a difference. In the cue models, past efforts at self-restraint make it *less* likely that one will enter a hot state and thereby lose all self-control. In our model, by contrast, the past exercise of self-restraint will, at least over the short term, make future self-control less likely.

Benabou and Tirole (2004) offer an alternative model of willpower. In their model, agents are quasi-hyperbolic discounters, but have limited knowledge of the degree of their present-bias (willpower for Benabou and Tirole). In the moment of greatest temptation agents know their willpower, but later can only infer their preferences from previous experience. To build self-reputation and show their future selves that they are not impatient, agents apply self discipline. The mechanism through which willpower works in their model is thus quite different from ours, and would in general lead to different predictions. One important difference is in how various self-control activities would be linked. In their model, an agent who successfully uses self-restraint

in one domain would become increasingly confident in his willpower and, thus, more likely to use self-restraint in another domain. In our model, the opposite is true, an agent who is depleted from using willpower in one situation will be less likely to use self-restraint in another seemingly unrelated situation.

Loewenstein and O'Donoghue (2005) and Fudenberg and Levine (2006) also have related models of costly self-control. In Loewenstein and O'Donoghue's framework the past exercise of self-restraint affects the direct cost of self-control in the future. Their model is quite general and they apply it both to intertemporal choice and to static choice under uncertainty. We view our paper as complementing theirs with a further examination of the implications of limited willpower for intertemporal choice and with greater emphasis on the allocation of willpower across competing demands for self-restraint. Fudenberg and Levine (2006) model a game between a long-run self and a sequence of short-run (myopic) selves. The short-run selves are both the "doer" selves (they "physically" make the choices) and the conduits of flow utility to the long-run self. At a cost, the long-run self may alter the preferences of the short-run self and thereby exercise self-control. By allowing cognitive loads to affect the cost of a given act of self-control, Fudenberg and Levine (2006) can accommodate the influence of such loads on the willingness to exercise self-control within a period. They also discuss how within period linkages between activities requiring self-control may emerge if self-control costs depend on multiple activities in a non-linear way. Our model is different in that it predicts *longer-term* effects, through the willpower constraint, of both shocks to willpower stocks and anticipated, alternative needs for self-restraint.²⁷

In addition, the willpower model's accommodation of increasing paths of consumption distinguishes it from the bulk of existing models of self-control in which agents are tempted to over-consume.²⁸ Negative time preference is usually viewed as unrelated to self-control and is attributed to utility from beliefs (anticipation). Models of anticipatory utility *can* accommodate increasing consumption paths, but cannot by themselves explain a taste for commitment, or profound procrastination, or links between seemingly unrelated acts of self-regulation. As we have shown, optimal willpower management can explain all of these phenomenon while, at the same time, inducing a preference for "saving the best for last."

Note that our analysis shows how willpower concerns may induce a preference for increasing consumption *in the absence of commitment*. Our willpower model predicts perfect smoothing of

²⁷There is also a conceptual difference between our approach and that taken by both Loewenstein and O'Donoghue (2005) and Fudenberg and Levine (2006) in that we model self-control costs as opportunity, rather than direct costs. See the end of this section for more on this issue.

²⁸The exception is Benabou and Tirole (2004). A complete discussion of this issue is left to the Technical Appendix available online at www.umich.edu/~dansily/research.

consumption if the consumer can commit to that path at any time prior to the moment consumption begins. The importance of commitment thus permits a simple way to distinguish the willpower motive from anticipatory utility in contributing to tastes for increasing consumption. To the extent that such tastes are revealed under commitment, it points to anticipatory utility from beliefs. To the extent that such tastes are *enhanced* by the absence of commitment, it points to willpower concerns.

In this regard, there is suggestive field evidence of substantial increases in consumption expenditure over short time intervals by unambiguously uncommitted consumers. Studies of high frequency data on individual expenditures in Britain indicate that consumers sometimes choose intervals of increasing consumption between the arrival of paychecks.²⁹ Both Kelly and Lanot (2002) and Huffman and Barenstein (2004) find that the non-durable expenditures of British consumers exhibit a U-shaped pattern over the course of the four weeks between paychecks. (See, for example, their Figures 1 and 4, respectively). Kelly and Lanot (2002) attribute the U-shaped pattern of consumption expenditure to uncertainty and precautionary saving within the pay period. Huffman and Barenstein (2004) attribute the interval of apparently increasing average consumption to measurement error.³⁰

Finally, we conclude this section by distinguishing our treatment of self-control as an endogenous shadow cost from the treatment in the previous literature as an exogenous direct cost.³¹ Presumably

²⁹There is also some, less statistically certain evidence of intervals of increasing consumption in high frequency data on daily expenditure by US consumers. See Stephens (2003), Tables 2 and 3 for total expenditure and Figures 4c-4g for non-durable expenditures.

³⁰The Family Expenditure Survey used in both studies asks respondents to record their expenditures in a diary for two weeks. The day upon which the diary begins is randomly assigned. Public use data aggregate each household's expenditures by week. Researchers can identify whether a respondent is paid monthly and the date his last paycheck arrived but not, precisely, the date of his next check. Therefore, diary weeks that begin 25 days or more after the last check arrived will typically include expenditures from at least one day after the next check has arrived. If consumption is higher on days just after the check arrives, this measurement error may explain the higher levels of consumption in diary weeks that began 25-30 days after the last check arrived (see Huffman and Barenstein, 2004, Figure 4). This error would not seem to explain, however, the increase in average expenditure between weeks starting 21 (perhaps 17) and 24 days after the last check. Neither can this measurement error explain why average expenditure is often higher in the third week after the arrival of the last check than it is in the second (Kelly and Lanot, 2002, Figure 1).

³¹In discussing an appropriate formulation to capture Baumeister's experiments, Loewenstein and O'Donoghue (2005) regard the exertion of willpower as generating disutility although Loewenstein (2000) regards the matter as an open question. Whether or not willpower depletion has a utility cost seems difficult to resolve empirically. What we have resolved analytically, however, is that including the stock of willpower (or its rate of change) in the utility function is not necessary to capture either the behaviors psychologists have documented in their laboratories or to explain a number of prominent "anomalies" behavioral economists have reported in the field. For these purposes, it is sufficient to append to the conventional formulation of intertemporal utility maximization the additional constraint

there are circumstances in which the classical model predicts well. It is, therefore, an attractive feature of any new model supplanting a classical one that it identifies the circumstances in which the classical results arise endogenously. If willpower depletion is appended as a constraint, there will be circumstances in which perfect smoothing will be optimal and consumption behavior predicted by the willpower model will be observationally equivalent to the one predicted by the classical model.

4 Extensions

Our approach in the previous sections has been to examine the implications of limited willpower in the simplest intertemporal problem. For the sake of simplicity we have also assumed that the agent always knows his stock of willpower with certainty and that willpower, if it is renewable, recovers at a rate unrelated to how it is exercised. In this section, we discuss the consequences of relaxing each assumption.

4.1 Depleting a Stock of Willpower of Unknown Size

Loewenstein (2000) points out that an agent may not be a good judge of the willpower he has in reserve. A similar issue has been analyzed by economists studying optimal consumption of a cake of unknown size.³² That analysis can easily be adapted to the case where the agent knows his cake size but not his initial reserves of willpower. As we show, along the optimal path the agent may discover either that he overestimated his self-control resources or that he has more willpower than he anticipated and has “a second wind.”

A two-state example will clarify the basic ideas. Suppose an agent has one of two initial stocks of willpower and assigns positive probability to each state. Assume that the agent has no alternative use of willpower and that even the larger stock of willpower is insufficient to implement perfect smoothing. As in subsection 2.1, assume that $f_W = 0$ so that consumption is constant when strictly positive.

If the agent knew that his willpower stock was low, he would conserve it by consuming at a rapid, constant rate and would deplete his cake quickly, after which he would have nothing left to consume until the end of the horizon. Denote as s_l the date when he would exhaust his cake if he was sure his initial willpower stock was low. Similarly, denote as $s_h > s_l$ the date when he would exhaust his cake if he was sure his initial willpower stock was high. Suppose, however, that the consumer could not observe his willpower reserves directly. Then, whatever consumption path he

that the consumption path chosen must not overexhaust the agent’s willpower.

³²See Gilbert (1979), among others.

chooses, he would obtain no information about his initial willpower stock until the date when his willpower would have run out if his initial stock had been small. At that point, he could infer the size of his initial reserves: if he loses control because he has no more willpower, then his initial reserves were small; if he retains control because he has more willpower, then his initial reserves were large.

Under uncertainty of this kind, it is optimal to consume at a constant rate intermediate between the alternative consumption rates he would choose under certainty. Since this would involve restraining his consumption more than is optimal if he were certain his willpower stock was low, he would run out of willpower sooner in the event that his initial stock was in fact low. If we denote as s_{ul} the date when he would run out under uncertainty if his initial stock of willpower was in fact low, then $s_{ul} < s_l$. In that event, lacking willpower reserves to control his consumption, he would exhaust the remaining cake at the rate \bar{c} and then would have nothing left to consume until the end of the horizon.

On the other hand, if at s_{ul} the agent discovered that he still had the ability to restrain consumption, then he would rationally conclude that his initial willpower stock had been high. Since he had been consuming at a faster rate than if he had known this state from the outset, he would find himself at s_{ul}^+ with more willpower and less cake. In this circumstance, the consumer would take advantage of this “second wind” by restricting his consumption below the constant rate he would have chosen if he had known from the outset of his high reserves of willpower.

4.2 Self-Control Builds Willpower Like a Muscle

Common intuition and some psychology experiments indicate that willpower may be like a muscle: controlling visceral urges depletes willpower over the short term, but the regular exercise of self-restraint may eventually build willpower. Most important for our purposes here, there is some evidence that willpower can be built up in one domain and then used to advantage in other arenas (Muraven *et al.*, 1999 and Muraven and Baumeister, 2000).³³

To incorporate the effects of buildable willpower, we introduce a third state variable, muscle, the level of which is denoted by $M(t)$. We augment our earlier model in two ways. First, we alter the transition equation for willpower (equation 4) to: $\dot{W}(t) = \gamma M(t) - f(W(t), c(t))$. As before, willpower is depleted by restraining consumption, $f(W(t), c(t))$, but this depletion is moderated by the service flow γ from the stock of muscle, $M(t)$. Since the stocks of willpower

³³In one experiment (Muraven *et al.*, 1999), subjects who participated in two-week self-control drills (regulating moods, improving posture, etc.) later showed significant increases in the length of time they would squeeze a handgrip relative to those who did not participate in the drills.

and muscle cannot jump, the only way to alter the rate of willpower depletion immediately is by altering contemporaneous consumption. However, current exercise helps develop future muscle and this muscle provides additional willpower at rate γ . If $\gamma = 0$, we recover our original transition equation. Second, we assume that the rate at which muscle develops (or deteriorates) is given by $\dot{M} = f(W(t), c(t)) - \sigma M(t)$. The idea is that exercising willpower today contributes to one's future muscle but the contributions decay.

The optimal consumption path in this muscle model shares some qualitative features with that in the model without muscle. First, the cake is entirely consumed. Second, *if willpower has no alternative uses*, perfect smoothing is optimal whenever it is feasible. In particular, for every initial level of muscle there is a willpower level $W_{\tilde{H}}$ above which the optimal path entails perfect smoothing. This is true because, for every initial stock of muscle there is an initial stock of willpower sufficiently large such that perfect smoothing is feasible and therefore optimal.³⁴ If we start with an initial level of willpower sufficient for perfect smoothing, decreases in that stock will eventually lead to a willpower level ($W_{\tilde{H}}$) where any further reduction in the initial stock of willpower will make the perfectly smooth path infeasible.

As in the model without muscle, on the optimal path the *sum* of the direct and indirect marginal benefits of increased consumption must remain equal to the marginal cost of that consumption (having a bit less cake for the future). But now, expanding consumption has a second marginal cost: besides depleting cake it reduces future muscle mass.

The opportunity to build muscle affects the optimal consumption path. For example, in the case where $f_W = 0$, it was optimal to consume at a constant rate until the cake was exhausted in the absence of muscle-building. In the case where muscle decays at a faster rate than it contributes to willpower ($\sigma \geq \gamma$), consumption is always increasing. Relative to the optimal path in the absence of muscle, the ability to build willpower through its exercise leads the agent to bear down at the beginning of the program in order to enjoy a greater willpower later. A more detailed analysis of this case is available in an online appendix.³⁵

5 Conclusion

This paper has explored the consequences of including in a conventional model of intertemporal choice a cognitive constraint consistent with a large body of experimental evidence. Specifically, we assumed that if an agent restricts his consumption of a tempting good, then exercising self-

³⁴Indeed, if the initial muscle level is large enough, the agent will be able to achieve perfect smoothing without any initial willpower.

³⁵The online appendix may be downloaded from www.umich.edu/~dansilv.

restraint depletes his finite stock of willpower – a resource useful for regulating urges of all kinds. This willpower constraint captures the common notion, consistent with laboratory experiments, that an individual has a limited, though positive, capacity to control his own impulsive behaviors.

To study how a willpower constraint affects choices over time, we introduced it in the simplest intertemporal model, where an agent decides how to consume a cake (or paycheck or stock of leisure time) over a fixed time period. Any willpower left over after regulating intertemporal consumption is used to control urges in other activities.

The model is simple but generates a rich set of implications. Most important, a willpower constrained consumer regards seemingly unrelated activities as linked because he uses the same cognitive resource to exercise self control in different activities. As a consequence, prior acts of self-restraint may affect an agent’s subsequent conduct because they reduce the stock of willpower he can use to regulate his behavior. The link between seemingly unrelated behaviors also implies that agents with limited willpower may appear to have domain-specific time preference – behaving as if completely myopic with regard to some choices and extremely forward-looking with regard to others. As important, time preference in one domain will change with the returns to self-control in another. Limited willpower also implies that self-control problems will affect demand for even nontempting goods, and that a lack of resources may cause impatience.

Limited willpower also has implications for the time-path of consumption. We saw that, because he draws on a single cognitive resource to regulate various urges, an agent may never smooth consumption, even if his initial willpower makes it feasible. In addition, an agent in our model may *increase* his consumption over time because exercising self control later, when his stock of willpower is reduced, requires more willpower than exercising the same self control earlier.

We considered two extensions of the intertemporal model. First, we considered the possibility that the decision maker is uncertain about his stock of willpower. We found that the agent would optimally consume at a higher rate as a hedge against the risk of running out of willpower and, if he learns that he has larger willpower reserves, experiences a “second wind” that permits further self-discipline later on the optimal path. Second, we investigated what would happen if current exercise of self control, while immediately depleting willpower, also builds willpower reserves in the future. We found that renewable willpower induces a time preference even when consumption would have been constant if willpower had been nonrenewable.

That willpower is scarce and depletable accords with both introspection and experiment. To explore how this constraint affects intertemporal behavior, we appended it to the canonical model of intertemporal choice. To our surprise, we found that the augmented model accounts *both* for prominent anomalies of intertemporal choice that have been the focus of the self-control literature *and* for other anomalies that are often treated as altogether separate phenomena requiring separate

models. Future research should clarify experimentally the form of the willpower depletion function. Given that taking account of a willpower constraint is so tractable, it should then be embedded in other economic models where self-control issues arise.

6 Appendix

Proof of Proposition 2: From (6) $\alpha(t)$ is a constant function, and with a slight abuse of notation we denote this constant as $\alpha \geq 0$. First assume $\bar{W} \geq W_H$. Consider the case where $\lambda(t) = 0$ for all $t \in [0, s]$. Since $U'(\cdot) > 0$, (3) requires $\alpha > 0$ and $c(t)$ constant. With a slight abuse of notation, we denote this constant as $c \geq 0$. Since $\alpha > 0$, (9) requires $R(s) = 0$. Since $R(0) = \bar{R} > 0$, then $c > 0$. From (8), $T - s \geq 0$. If $T - s > 0$, (8) would require $U(c) - U'(c)c = 0$. But since $U'(\cdot) > 0, U''(\cdot) < 0$, and $c > 0$, $U(c) - U'(c)c > 0$; hence, $T - s > 0$ must be ruled out. It follows that when $\lambda(t) = 0$ for all $t \in [0, s]$, then $s = T$ and $c = \frac{\bar{R}}{T}$. By definition of W_H , $W(T) \geq 0$. Thus (10) is satisfied. This proves the first statement in the proposition. Now assume $\bar{W} < W_H$. Then $\lambda(t) > 0$ for all $t \in [0, s]$, for suppose to the contrary that $\lambda(t) = 0$ for some $t \in [0, s]$. Equation (7) implies that $\lambda(t)$ is weakly decreasing and can be written as $\lambda(t) = \lambda(0) e^{\int_{n=0}^t f_W(W(n), c(n)) dn}$. Since $e^{\int_{n=0}^t f_W(W(n), c(n)) dn} > 0$ for all t , $\lambda(t) = 0$ for some $t \in [0, s]$, implies that $\lambda(t) = 0$ for all $t \in [0, s]$. But as seen above the conditions above then imply that $c = \frac{\bar{R}}{T}$ which is infeasible when $\bar{W} < W_H$. So if $\bar{W} < W_H$ then $\lambda(t) > 0$ for all $t \in [0, s]$, and by (10) $W(s) = 0$. To satisfy (3) with $U'(\cdot) > 0$ and $-\lambda f_c > 0$ requires $\alpha > 0$; and, again, since $\alpha > 0$, (9) requires that the cake is entirely consumed ($R(s) = 0$).

Figures




Figure 1: Optimal time (s) to exhaust the cake and optimal willpower bequest ($W(s)$) as a function of the return to willpower applied to alternative activities (m).

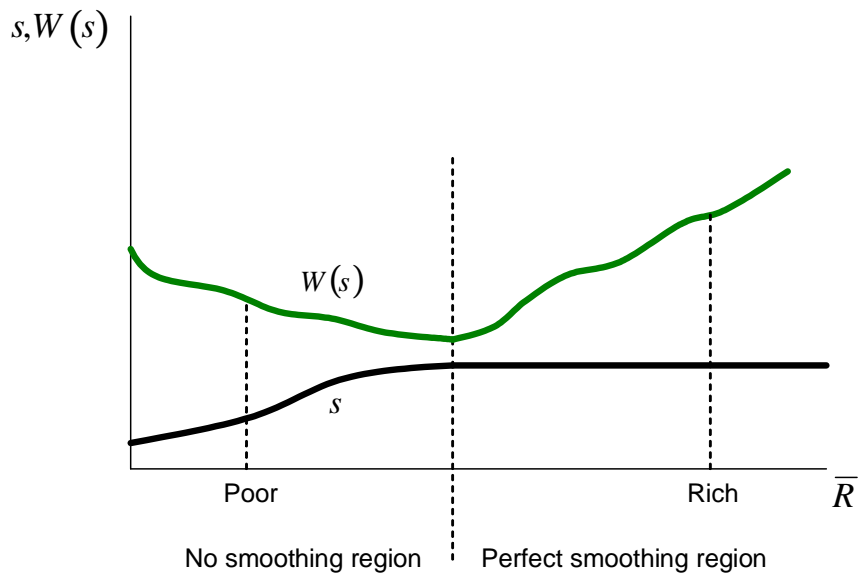


Figure 2: Optimal time time (s) to exhaust the cake and optimal willpower bequest ($W(s)$) as a function of the initial cake size (\bar{R}).

References

- [1] Baumeister, Roy; Heatherton, Todd and Tice, Dianne. *Losing Control: How and Why People Fail at Self-Regulation*. San Diego, CA: Academic Press, 1994.
- [2] Baumeister, Roy and Vohs, Kathleen. "Willpower, Choice, and Self-Control" in Loewenstein, George; Read, Daniel and Baumeister, Roy, eds. *Time and Decision*. New York: Russell Sage Foundation, 2003.
- [3] Benabou, Roland and Tirole, Jean. "Willpower and Personal Rules." *Journal of Political Economy*, 2004, 112(4), 848-87.
- [4] Bernheim, B. Douglas and Rangel, Antonio. "Addiction and Cue-Triggered Decision Processes." *American Economic Review*, 2004, 95(5), pp. 1558-1590.
- [5] Dewitte, Siegfried; Pandelaere, Mario; Briers, Barbara and Warlop, Luk. "Cognitive Load Has Negative After Effects on Consumer Decision Making." Mimeo, Katholieke Universiteit Leuven, 2005.

- [6] Faber, Ronald and Vohs, Kathleen. "To Buy or Not to Buy?: Self-Control and Self-Regulatory Failure in Purchasing Behavior" in Baumeister, Roy and Vohs, Kathleen, eds. *Handbook of Self-Regulation: Research, Theory, and Applications*. New York: Guilford Press, 2004.
- [7] Fischer, Carolyn. "Read This Paper Later: Procrastination with Time-Consistent Preferences." *Journal of Economic Behavior and Organization*, 2001, 46(3), pp. 249-69.
- [8] Fudenberg, Drew and Levine, David K. "A Dual Self Model of Impulse Control." *American Economic Review*, 2006, 96(5):1449-1476.
- [9] Gilbert, Richard "Optimal Depletion of an Uncertain Stock." *Review of Economic Studies*, 1979, 46 (1), pp. 47-58.
- [10] Gruber, Jonathan and Köszegi, Botond. "Is Addiction 'Rational?' Theory and Evidence." *Quarterly Journal of Economics*, (2001), 116(4), pp. 1261-1305.
- [11] Gul, Faruk and Pesendorfer, Wolfgang. "Temptation and Self-Control." *Econometrica*, 2001, 69(6), pp. 1403-35.
- [12] Gul, Faruk and Pesendorfer, Wolfgang. "Self-Control and the Theory of Consumption." *Econometrica*, 2004, 72(1), pp. 119-58.
- [13] Gul, Faruk and Pesendorfer, Wolfgang. "Harmful Addiction." *Review of Economic Studies*, 2007, 74(1), pp. 147-172.
- [14] Hinson, John, Jameson, Tina, and Whitney, Paul. "Impulsive Decision Making and Working Memory." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 2003, 29(2):298-306.
- [15] Huffman, David and Barenstein, Matias. "Riches to Rags Every Month? The Fall in Consumption Expenditures Between Paydays." IZA Discussion Paper No. 1430, December 2004.
- [16] Kelly, Clare and Lanot, Gauthier. "Consumption Patterns Over Pay Periods." Warwick Economic Research Papers No. 656, November 2002.
- [17] Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 1997, 112(2), pp.443-77.
- [18] Laibson, David. "A Cue-Theory of Consumption." *Quarterly Journal of Economics*, 2001, 116(1), pp.81-119.

- [19] Lawrance, Emily. "Poverty and the Rate of Time Preference: Evidence from Panel Data." *The Journal of Political Economy*, Vol. 99, No. 1, 54-77. Feb., 1991.
- [20] Loewenstein, George. "Willpower: A Decision-Theorists's Perspective." *Law and Philosophy*, 2000, 19, pp. 51-76.
- [21] Loewenstein, George and O'Donoghue, Ted. "Animal Spirits: Affective and Deliberative Processes in Economic Behavior." Mimeo, Carnegie Mellon University, 2005.
- [22] Muraven, Mark. "Mechanisms of Self-Control Failure: Motivation and Limited Resource." Ph.D. Dissertation, Case Western Reserve University, 1998.
- [23] Muraven, Mark and Baumeister, Roy. "Self-Regulation and Depletion of Limited Resources: Does Self-Control Resemble a Muscle?" *Psychological Bulletin*, 2000; 126(2); pps. 247-259.
- [24] Muraven, Mark; Baumeister, Roy and Tice, Dianne. "Longitudinal Improvement of Self-Regulation Through Practice: Building Self-Control Strength Through Repeated Exercise." *Journal of Social Psychology*, 1999, 139(4), pp. 446-457.
- [25] O'Donoghue, Ted and Rabin, Matthew. "Doing It Now or Later." *American Economic Review*, 1999, 89(1), pp. 103-124.
- [26] Shiv, Baba and Fedorikhin, Alexander. "Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making." *Journal of Consumer Research*, December 1999, 26(3):278-292.
- [27] Stephens, Melvin, Jr. "'3rd of tha Month': Do Social Security Recipients Smooth Consumption Between Checks." *American Economic Review*, 2003, 93(1):406-422.
- [28] Thaler, Richard and Shefrin, Hersh. "An Economic Theory of Self-Control," *Journal of Political Economy*, 1988, 89 392-406.
- [29] Vohs, Kathleen and Faber, Ronald. "Spent Resources: Self-Regulation and Impulse Buying." Under review, 2004.