

Willpower and the Optimal Control of Visceral Urges

Emre Ozdenoren, Stephen Salant, and Dan Silverman

Willpower and the Optimal Control of Visceral Urges*

Emre Ozdenoren[†]

Stephen Salant[‡]

Dan Silverman[§]

This Version: June 16, 2010

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 over time an endowment of a tempting and storable consumption good or “cake”. We assume that restraining consumption below the most tempting feasible rate requires willpower. Any willpower not used to regulate consumption may be valuable in controlling other urges. Willpower thus links otherwise unrelated behaviors requiring self-control. An agent with limited willpower will display apparent domain-specific time preference. Such an agent will almost never perfectly smooth his consumption, even when it is feasible to do so. Whether the agent relaxes control of his consumption over time (as experimental psychologists predict) or tightens it (as most behavioral theories predict) depends in our model on the net effect of two analytically distinct and opposing forces.

*We thank Roy Baumeister, Gérard Gaudet, and Miles Kimball for their important contributions to this paper. We also thank Roland Benabou, Dan Benjamin, Doug Bernheim, Drew Fudenberg, Yoram Halevy, Kai-Uwe Kuhn, David Laibson, David Levine, Robert Mendelsohn, Muriel Niederle, Andrew Postlewaite, Matthew Rabin, Lones Smith, Joel Sobel, Kathleen Vohs, Susan Woodward, and Itzhak Zilcha as well as participants in seminars at Berkeley, Harvard, IAS/Princeton, SITE, Calgary, Central Florida, Michigan, Milano-Bicocca, Montreal, NYU, Penn, Pittsburgh, and Santa Barbara for their helpful comments on a previous draft. Finally, we want to thank George-Marios Angeletos and the three anonymous referees for helpful suggestions that led to the current version of the paper. Silverman gratefully acknowledges the support of the Institute for Advanced Study.

[†]Department of Economics, London Business School, Regent’s Park, London, NW1 4SA.

[‡]Department of Economics, University of Michigan, and Nonresident Fellow, Resources for the Future, 611 Tappan St., Ann Arbor, MI 48109-1220.

[§]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 (Laibson 1997; O’Donoghue and Rabin 1999), temptation costs (Gul and Pesendorfer 2001, 2004), and conflicts between selves or systems (Thaler and Shefrin 1988; Bernheim and Rangel 2004; Fudenberg and Levine 2006). These models are consistent with a great deal of experimental evidence, and they have been fruitfully applied to 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 economics 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*.

Common behaviors indicate that the ability to self-regulate may be a limited resource. 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 staying up late to watch TV while gorging on junk food. Dieters can often maintain their discipline for short periods but find such self-restraint unsustainable over the long term.¹ Households can often stick to a budget for a while but end fall back into old spending patterns and end up well short of their saving goals.

Experimental psychology (Baumeister et al. 1994; Baumeister and Vohs 2003) has gone beyond such anecdotes and 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; the remainder serve 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.² Subjects who participate in the first phase display substantially less endurance in the second phase. This apparent link between the exercise of self-

¹Kessler (2009) offers especially vivid accounts of this pattern.

²For example, in the first phase subjects have been asked not to eat tempting foods, not to drink when thirsty, and to inhibit automated/habitual behaviors such as reading the subtitles of film. The self-control tasks typically differ in the two phases of these experiments, although in Vohs and Heatherton (2000) tempting food is used in both phases.

control in one activity and later self-discipline in another has been observed repeatedly, under many different manipulations and measures of self-regulation (Baumeister and Vohs 2003; Vohs and Faber 2004). Moreover, subjects exhibit foresight, performing less well on the initial self-control task if informed at the outset that they will be asked to complete an additional task requiring self-control at the end (see Muraven et al. 2006 and the references therein). Two experiments (Vohs and Faber 2004; Dewitte et al. 2005) have shown that willpower depletion and prior cognitive loads affect subsequent *economic* behavior.³ Although individual experiments have their weaknesses, we believe that overall this collection of experimental findings reinforces the intuitive notion of willpower as a depletable cognitive resource.

Self-control may be exercised for a variety of reasons. In this paper we investigate the use of willpower to limit immediate consumption of a good so that more of that good can be enjoyed later. To do so, we add a willpower constraint to an especially simple model of intertemporal choice: the canonical cake-eating problem. In our formulation, moderating consumption of tempting goods—including leisure—requires willpower; and the greater restraint the consumer exercises, the faster his willpower erodes.⁴ We also take account of other activities besides intertemporal consumption that require willpower: cramming for and taking exams, training for performances, maintaining a diet, or sustaining mental focus over long periods.

Our model of limited willpower is simple, but 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 yet, at the same time, appear to be myopic when it comes to the long-term consequences of food and drink. His time preference would thus appear *domain-specific*.

Limited willpower also 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 forgoing this option and

³In 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. (2005) should be distinguished from research showing the effects of cognitive loads on *contemporaneous* choices. In these latter studies, respondents were more likely to choose cake over fruit (Shiv and Fedorikhin 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.

⁴Experiments indicate that willpower may also be drained by inhibiting the automatic responses of subjects or DAN increasing the cognitive loads that subjects must bear. Our model accommodates these causes of depletion either as anticipated uses of willpower that follow consumption choices or as a shock to the initial willpower stock.

using more willpower to regulate alternative activities. In addition, the Baumeister experiments suggest that a person with limited willpower will relax his self control over time and thus consume food or leisure at an increasing rate—a pattern consistent with decisions to postpone consumption or get the hard part of a project out of the way at the outset. The thrust of previous economic models of self-control is quite different: agents are predicted to *reduce* their consumption over time. Either outcome is possible in our model, and we can distinguish analytically the opposing forces that determine whether consumption rises or falls.

In an extension to our basic model, we consider the case where, consistent with some psychology experiments (Muraven et al. 1999; Muraven and Baumeister 2000), the ability to self-regulate is like a muscle; that is, avoiding temptations depletes willpower over the short term, but the regular exercise of such restraint can eventually build willpower. This extended model shows how the benefits of building willpower through its exercise generates an incentive to apply more self-restraint at the outset. We find that muscle building causes optimal consumption to rise with time even when basic willpower concerns alone would imply that consumption would be constant.

The rest of the paper proceeds as follows. Section 2 introduces and analyzes our baseline model. Section 3 discusses how our theory relates to the pertinent literature. Section 4 considers an extension where the exercise of self-control depletes willpower over the short term but builds it over the longer term. Section 5 concludes.

2 A Model of Limited Willpower

To investigate the consequences of limited willpower in an especially simple setting, we consider the canonical cake-eating problem in which a consumer chooses his consumption path $c(t)$ to maximize his discounted utility $U(c(t))$ over a finite horizon. We assume that $U(0) = 0$ and, for $c \in (0, \bar{c})$, that $U(c)$ is strictly concave, strictly increasing, and twice differentiable. We interpret the consumption rate \bar{c} as reflecting either (i) a physical limit on the rate of consumption,⁵ or (ii) the rate of consumption resulting in the highest utility flow ($U(\bar{c}) \geq U(c)$ for $c > \bar{c}$), in which case consuming at a faster rate than \bar{c} is never optimal.⁶ The analysis that follows permits either interpretation. Our focus is on non-addictive goods such as those used in experiments by Baumeister and others so we assume that U is independent of prior consumption. We denote the stock of cake remaining at time t as $R(t)$ and assume that $R(0)$ is given. There is no return to saving in the

⁵This interpretation is especially appropriate if the consumption good is leisure time. If, for example, time is measured in hours, then leisure cannot be consumed faster than one hour per hour: $c \in [0, 1]$.

⁶Even if consuming at a rate faster than \bar{c} were feasible, it would never be chosen because it saves no willpower and, by shortening the time to exhaustion of the cake without increasing the utility flow, it reduces overall utility.

model; therefore, at time t the cake declines at rate $c(t)$.

2.1 Willpower Depletion

We depart from the canonical model by assuming that the consumption good is tempting and that the agent, initially endowed with willpower stock $W(0)$, may deplete his willpower when he limits his rate of consumption. More precisely, we assume that the rate of willpower depletion at t depends on the stock of remaining willpower, $W(t)$, and the current rate of consumption, $c(t)$. We also assume that the depletion function, denoted $f(W, c)$, is nonnegative and that $f(W, c) = 0$ for $c \geq \bar{c}$. Finally, we assume that f is decreasing in W , strictly decreasing and strictly convex for $c \in [0, \bar{c}]$, and that $f_{cW} \geq 0$; when no cake remains, the rate of willpower depletion is zero.

These assumptions follow, in large part, from a natural view of the determinants of temptation costs, a view with foundations established in the literature on self-control and dual selves (Gul and Pesendorfer 2001; Fudenberg and Levine 2006).⁷ In this view, the temptation (willpower) costs associated with consuming at rate c are increasing in the utility forgone by choosing c instead of the most tempting feasible rate. Thus, in our setting, no willpower is expended if the consumer chooses rate \bar{c} , and no willpower is expended if no cake remains to be consumed. Similarly, if the agent is committed to consume only at rate c (i.e., no other rate is feasible) then, once again, no willpower is depleted. More generally, given that utility is strictly increasing in consumption for $c \in [0, \bar{c}]$, whenever cake remains the rate of willpower depletion will, as we have assumed, strictly decline with c . Assuming, in addition, that f is strictly convex in c , captures the natural idea that increases in the rate of consumption save less willpower when the rate is closer to \bar{c} .

Adopting this view also implies—given our continuous-time formulation—that the utility forgone by restraining the rate of consumption does not depend on the size of the remaining cake. This follows because, in continuous time, as long as some cake remains ($R(t) > 0$), it is generally feasible to consume at any rate below \bar{c} . Hence, if cake remains and there are no external restrictions on the menu of consumption rates, then the rate of willpower depletion at time t depends only on the rate of consumption and not the size of the cake.⁸

⁷This argument is formalized in the Appendix.

⁸In a discrete-time formulation of our model, the utility forgone by restraining consumption—and thus the rate of willpower depletion—*would* depend on the size of the cake. So long as the remaining cake were less than \bar{c} , the most tempting choice at any time would be to consume the entire cake. This is also why, in Gul and Pesendorfer’s (2004) discrete-time model, the most tempting choice in any period is consumption of all remaining wealth. If Gul and Pesendorfer’s model were, like ours, formulated in continuous time, then the most tempting choice at any moment would also be the most tempting feasible rate of consumption (a flow), not the entirety of remaining wealth (a stock). However, even in continuous time, there is a sense in which richer people, with their greater consumption opportunities

Later in this section we shall clarify both the motivation for and the implications of our other substantive behavioral assumption: that the rate of willpower depletion is weakly decreasing in W .

2.2 Alternative Uses of Willpower

We allow for the possibility that any willpower that remains after moderating intertemporal consumption may be used to restrain urges in other activities. We regard this willpower “bequest” not as directly generating utility, but rather as altering the feasible set of choices in other activities and thereby *indirectly* determining the utility generated from those activities.⁹ We use $m(\cdot)$ to denote this indirect utility function. To illustrate, suppose the intertemporal consumption problem is how a student allocates time between leisure and studying a requisite 6 hours for a test to be given the next day. His time horizon is 24 hours, during which he will consume 18 hours of leisure, his cake. A student with no willpower would party until 6 hours before the exam, when he would begin to study. If the student could precommit, he would recognize that nonstop partying is “too much of a good thing” and would instead spread his studying evenly throughout the day, studying a bit, then partying a while, then studying a bit, and so on until the exam. While this plan would maximize his enjoyment of leisure, in the absence of commitment it would require enormous self-restraint and might leave the student with too little willpower to perform well on the exam itself. In this example, the student not only enjoys leisure but values success on the exam (the alternative activity) which in turn depends on his having enough willpower in reserve to maintain his focus and avoid careless mistakes. We also allow $m(\cdot)$ to be a constant; i.e., willpower’s value might derive solely from its use in moderating intertemporal consumption. More generally, we view the bequest function as a stand-in for the value of willpower applied to subsequent activities in which the agent might benefit from self-control.

2.3 The Cake-Eating Problem

The agent cannot choose a consumption path that results in negative willpower. From the set of consumption paths that maintain nonnegative stocks of both willpower and cake, the consumer (more cake), endure greater temptation. If, for example, a rich person duplicates the *entire* consumption path of a poor person who at some point exhausts his (smaller) cake, then the rich person will have depleted more of his willpower. After he has finished his cake, the poor person requires no willpower to consume nothing, yet the rich person must use his willpower to resist temptation and set his consumption to zero.

⁹None of our results depends on the units in which willpower is measured. The reason is simple: for any way of measuring willpower, we identify the utility-maximizing path among those that the agent has sufficient self-control to implement. Since using a different willpower metric would identify the same set of decision paths and since utility depends only on these paths and never on the willpower number ascribed to them, our invariance result follows.

chooses the path that maximizes the sum of his discounted utility of consumption and the value of the willpower bequeathed to the regulation of his alternative activities.

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

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

Here ρ is the subjective rate of time discount and $s = \sup\{t \in [0, T] : R(t) > 0\}$.¹⁰ For the bulk of our analysis, we interpret the relevant time horizon as short (a day, several days, or a week), and so we set $\rho = 0$. Variations in the rate of consumption are thus attributable to the effects of limited willpower. Note that the function describing changes in the stock of willpower $\left(\dot{W}(t)\right)$ jumps to zero when the cake is exhausted.

As discussed in what follows, 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 first investigate a particularly simple case where the willpower constraint induces linkages but not time preference; then we move on to the more general case where the willpower constraint can also induce time preference.¹¹

2.4 Induced Linkages

We begin our analysis by studying optimal choice in the case where the rate of willpower depletion depends only on the level of consumption and not on the willpower stock. Formally, we assume that $f(W(t), c(t)) = g(c(t))$. Moreover, we assume that additional willpower devoted to alternative activities has a constant marginal value ($m'(W(t)) \equiv m$). Again we note that m might equal zero, in which case a willpower bequest has no value.

Under these assumptions, it is optimal (i) to exhaust the cake and (ii) to consume it at a constant rate as long as cake remains. To see why this is true, suppose the contrary. If it is optimal

¹⁰We 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.

¹¹Because 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 about the equivalent problem are provided in the Appendix.

to leave some cake unconsumed, then over some time interval the agent must consume at a rate slower than \bar{c} . This follows because we have assumed that the cake is too small to consume at the maximal rate over the entire time horizon ($\bar{R} < \bar{c}T$). But then an increase in the consumption rate over part of that interval would increase overall utility and would at the same time save willpower. Hence, the original program is not optimal and the optimal program must exhaust the cake. Now suppose it were optimal for the agent to vary his consumption rate before the cake is exhausted. Because utility is strictly concave in tempting consumption, it follows that marginally reducing larger consumption while marginally increasing the smaller consumption 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, we are assuming that willpower depletion is independent of the willpower stock and is a decreasing function of consumption. This means that the proposed perturbation will actually make more willpower available for use elsewhere: increasing the smaller consumption releases more willpower than is depleted by reducing the larger consumption by an offsetting amount, and no changes would occur in willpower depletion at other times since consumption would be unchanged then. So if the original program is feasible, then the perturbed program will also be feasible but will yield higher utility. Hence, in the optimal program 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 must be zero. A perturbation that further smooths consumption in this situation has a decidedly different effect than the one described previously. In this situation, marginally reducing consumption in the phase when it is positive and marginally increasing it in the phase when it is zero depletes willpower in *both phases*. This occurs because no willpower is depleted when no cake is available but, once cake becomes available in the second phase, consuming a small amount requires additional willpower. As the analysis to follow will clarify, even when feasible 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, a better strategy is to complete consumption before time T and save more willpower for other purposes. We will refer to the case where the agent stretches consumption to time T as the “perfect-smoothing regime”, and the alternative as the “no-smoothing regime”.

¹²See note 24 for a formal proof of these arguments.

2.4.1 Simple Comparative-Statics

In the rest of Section 2.4, we assume that the agent has enough cake or willpower to implement perfect smoothing if he chooses.¹³ Then the payoff from consuming at constant rate $c \in [0, \bar{c}]$ and exhausting the cake at time $s \in [s_{\min}, T]$ where $s_{\min} = \bar{R}/\bar{c}$, is $sU(c) + m[\bar{W} - sg(c)]$, for $s = \bar{R}/c$. Provisionally substituting \bar{R}/s for c we obtain a continuous and strictly concave objective function of one variable, s . The agent's problem is equivalent to choosing the time to exhaust the cake, s , to solve:

$$\max_{s \in [s_{\min}, T]} sU\left(\frac{\bar{R}}{s}\right) + m[\bar{W} - sg\left(\frac{\bar{R}}{s}\right)].$$

It is straightforward to verify that an optimum exists and, since the maximand is strictly concave, any solution to the first-order condition is optimal. The optimum may occur at $s < T$ or at $s = T$.

No-smoothing Regime When no smoothing is optimal, the first-order condition¹⁴ implies that

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

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 equation (2) implicitly determines the rate of constant consumption (c) as a function of the exogenous marginal value (m) of bequeathed willpower; in the no-smoothing regime, therefore, the rate of consumption does not depend on either of the other exogenous variables (\bar{W} and \bar{R}). If an agent in this regime begins with larger initial willpower, then he optimally consumes at the same rate and, with a cake of unchanged size, exhausts 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

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

¹⁴It may also be optimal to consume at rate \bar{c} until the cake is exhausted at s_{\min} . This happens if and only if the left-hand side of (1) is weakly negative at $s = s_{\min}$. Similarly, it may be optimal, when m is sufficiently small, not to smooth perfectly and yet carry no willpower over into alternative activities. Each of these cases is less interesting and so, in light of space constraints, we do not discuss them.

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.¹⁵

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 .¹⁶ Hence, the optimal rate of consumption would increase in response. Given that the cake size is unchanged, it would be exhausted sooner—leaving more willpower for the subsequent activity.

Perfect-smoothing Regime In contrast, when perfect-smoothing is optimal, behavior does not change 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 so would 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. Given instead a larger cake, he would consume it at a faster rate, exhaust it at an unchanged date ($s = T$), and then would have more willpower left over for the alternative activities.¹⁷ Comparative statics results for each regime are summarized in Table 1. Next we turn to the behavioral implications of these results.

2.4.2 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

¹⁵Indeed, as observed in the previous note, for m sufficiently small it is possible, that the consumer will choose to carry no willpower over into alternative activities.

¹⁶Totally differentiating the first-order condition that determines 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 because $g(\cdot)$ is positive, strictly decreasing, and strictly convex while U is strictly concave.

¹⁷Agents with very large cakes may consume at a rate very close to \bar{c} for the entire horizon, and so have nearly all of their initial willpower stock left for self-control in alternative activities. The model thus predicts that the very rich will appear especially disciplined in non-consumption activities requiring self-control (maintaining exercise regimes, preparing for performances or tests). With great wealth, it is often unnecessary to resist tempting consumption.

	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 Statics Results in Each Region.** 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}). Results apply on the interior of each region.

the return to the alternative activities is high enough that the agent is in the no-smoothing regime.¹⁸ If this agent had more willpower then, rather than using any more to restrain tempting consumption, 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 his initial stock of willpower.

2.4.3 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 to postpone gratification in one set of activities and yet be profoundly myopic about choices in another, even though he discounts time at a single rate. 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 that 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, then an increase in m will result in the agent consuming his tempting cake faster, exhausting it more quickly, and hence having more willpower for work related activities (see Figure 1 and the top rows in Table 1). Thus, the agent would appear more myopic when it comes to consumption choices but at the same time more forward-looking when it comes to work.

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

¹⁸To see which regime is optimal for given parameters m and \bar{R} , set $s = T$ in the left hand side of (1). If the resulting expression 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, our model of limited willpower makes distinctive predictions about changes in self-restraint exercised in one activity as a function of the return to exerting self-restraint in another.

2.5 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.¹⁹

2.5.1 The Time Path of Consumption

To begin, consider the intuitive trade-offs revealed by the first-order conditions characterizing optimal choice. When consumption is interior ($c(t) \in (0, \bar{c})$), the first-order condition (A1) is

$$\underbrace{U'(c(t))}_{\text{direct marginal benefit}} + \underbrace{(-f_c)\lambda(t)}_{\text{indirect marginal benefit}} = \underbrace{\alpha(t)}_{\text{marginal cost}}, \quad (3)$$

where $\lambda(t)$ is the current shadow value of additional willpower at t and $\alpha(t)$ is the current shadow value of additional cake then. Equation (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: the increase in utility at time t from consuming more at t . The marginal cost of consuming at a faster rate at t , $\alpha(t)$, is also standard: the utility forgone because the additional cake consumed at t cannot be consumed at another time.

What is distinctive in (3) is that increasing consumption at t also has an *indirect* marginal benefit ($-f_c\lambda(t) > 0$) which arises from two factors. First, in raising the rate of consumption at t , the agent depletes willpower more slowly at t ($-f_c(W(t), c(t))$). Second, because $f_W < 0$, the value $\lambda(t)$ of each unit of willpower saved at t reflects an unfolding sequence of subsequent willpower savings that cumulatively *magnify* the initial savings.²⁰ We refer to the first of these factors as the “immediate indirect benefit” and the second as the “magnifier”.

¹⁹See the Appendix for details of this analysis.

²⁰Using the differential equation (A5), which describes the evolution of $\lambda(t)$ together with the endpoint condition

Figure 2 depicts the determinants of optimal consumption at t . The marginal benefit schedule (the *sum* of the direct and indirect marginal benefits) is downward sloping because $U(c)$ and $-f(W, c)$ are, by assumption, strictly concave in consumption at t . The marginal cost curve $\alpha(t)$ is horizontal because the marginal value of additional cake at an instant is unaffected by the rate of consumption at that instant. Optimal consumption occurs where the two curves intersect.²¹

As time elapses, the marginal cost curve remains fixed (see equation (A4)) because additional cake is equally valuable whenever it arrives; its mere availability prior to consumption provides no services, and its future arrival can be anticipated by consuming more in advance. Hence, optimal consumption changes over time if and only if the marginal benefit curve shifts. If the marginal benefit curve shifts outward (resp. inward) over time, it is optimal to increase (resp. decrease) consumption. The direct component of the marginal benefit schedule, $U'(c)$, does not change with time, so shifts in the indirect marginal benefit determine whether it is optimal to increase or decrease consumption over time.

If the stock of remaining reserves has no effect on the rate of willpower depletion ($f_W = 0$), then neither factor varies over time and it is optimal to consume at a constant rate while cake remains.²² We provisionally adopted this assumption to simplify the exposition in Section 2.4.

The Possibility of Negative Time Preference In the case where regulating consumption depletes willpower more quickly when willpower reserves are smaller ($f_W < 0$), each factor determining the indirect marginal benefit changes over time. The *immediate* indirect benefit of consuming at a faster rate increases as time elapses because, with time, willpower reserves are depleted by use. At lower levels of willpower, consuming marginally faster releases more willpower. However, this effect is magnified less because the magnification occurs over a shorter interval of time.

Whether consumption rises or falls over time depends on which these two factors dominates. It is instructive first to identify the boundary case where consumption is constant even though

(A8), we can write the indirect marginal benefit in equation (3) as the product of two factors:

$$(-f_c)\lambda(t) = -f_c(W, c) \cdot \left[m'(W(s)) \exp \left(\int_{x=t}^{\min(s, T)} -f_W(W(x), c(x)) dx \right) \right],$$

where for clarity we have chosen as the endpoint the case where $W(s) > 0$. If no willpower were carried over into alternative activities, $m'(W(s))$ would be replaced by a different constant, but this would not affect our discussion.

²¹The figure also illustrates the two corner cases. If $c = 0$ is optimal, then the marginal cost curve must weakly exceed the marginal benefit curve at $c = 0$; if $c = \bar{c}$ is optimal, then the marginal benefit curve must weakly exceed the marginal cost curve at $c = \bar{c}$.

²²To see that $c(t)$ must be constant for $t \leq s$, observe that (A4) implies that $\alpha(t)$ is constant. If $f_W \equiv 0$, then $f_c(W(t), c(t))$ does not vary with $W(t)$ and, given (A5), $\lambda(t)$ is constant. Since we are assuming that $\rho = 0$, the solution to (A1) will be the same for all $t \leq s$.

both factors vary over time. For this case to occur, the magnifier must decline, in percentage terms, by exactly as much as the immediate indirect benefit increases. The magnifier declines in percentage terms by $-\dot{\lambda}/\lambda = -f_W(W, c) > 0$. When consumption is constant, the immediate indirect benefit increases in percentage terms by $f_{cW}\dot{W}/f_c = -f_{cW}f/f_c > 0$. When consumption is locally constant these two percentage changes must be equal. When consumption strictly increases (strictly decreases) locally, the percentage increase in the immediate benefit must be strictly larger (strictly smaller) than the percentage reduction in the magnifier.

To summarize, the following proposition characterizes the determinants of the (local) rate of change in consumption.

Proposition 1 *Let $c(t) \in (0, \bar{c})$ for some $t \leq s$. Then $\dot{c} \gtrless 0$ as $\lambda \frac{d}{dW} \frac{d}{dc} \ln f(W, c) \gtrless 0$.*

Proof. Because $\alpha(t)$ is constant over time, differentiating equation (3) with respect to t implies $[U''(c) - \lambda f_{cc}]\dot{c} + \{\dot{\lambda}(-f_c) + \lambda f_{cW}(-\dot{W})\} = 0$. Since $[U''(c) - \lambda f_{cc}] < 0$, it follows that $\dot{c} \gtrless 0$ if and only if $-f_c \lambda \{(f_{cW}/-f_c)(-\dot{W}) + (\dot{\lambda}/\lambda)\} \gtrless 0$. Now using (A5) and (A2) to replace $\dot{\lambda}/\lambda$ and $-\dot{W}$, we obtain $\dot{c} \gtrless 0$ if and only if $\lambda \{f f_{cW} - f_c f_W\} \gtrless 0$. The result then follows because $\lambda \{f f_{cW} - f_c f_W\} \gtrless 0$ as $\lambda \frac{d}{dW} \frac{d}{dc} \ln f(W, c) \gtrless 0$. ■

This proposition implies that, when interior, consumption is constant in only three circumstances: (i) when willpower is not scarce ($\lambda = 0$), (ii) when the rate of willpower depletion is independent of the stock of willpower ($f_W = 0$), or (iii) when, despite neither of the preceding circumstances, the indirect marginal benefit of a marginal increase in consumption is constant.

Proposition 1 can also be interpreted as stating that, when willpower is scarce, consumption strictly increases whenever the logarithm of the depletion function has a strictly positive cross partial derivative.²³ Knowing this, it is straightforward to construct examples where the consumption path rises or falls monotonically before jumping to zero. For instance, consider the following function of willpower depletion:

$$f(W, c) = v(c)A + K(W)\Gamma(v(c)),$$

where $v(c) \equiv U(\bar{c}) - U(c)$. We assume that $K(\cdot) > 0$, $K'(\cdot) < 0$, $\Gamma(\cdot) > 0$, $\Gamma''(\cdot) > 0$, and that $K(\bar{W})\Gamma'(0) + A > 0$. Note that the last condition is always satisfied when $A \geq 0$ and that it will be violated only when A is too negative. These conditions together ensure that $f(W, c)$ satisfies all the assumptions imposed earlier (i.e., $f(W, \bar{c}) = 0$, $f_c < 0$, $f_W \leq 0$, $f_{cc} \geq 0$, $f_{cW} \geq 0$). Differentiating then yields

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

²³The reader may recognize this condition as log-supermodularity of the depletion function.

Since $\Gamma'' > 0$, it follows that $v(c)\Gamma'(v(c)) - \Gamma(v(c)) > 0$. Therefore, the sign of the expression just displayed is the same as the sign of A . So, whenever $\lambda > 0$, by Proposition 1 we have

$$\text{sign } \dot{c} = \text{sign } A.$$

This example demonstrates how limited willpower may induce either positive or negative time preference.

Increasing paths of consumption have previously been attributed to the effects of habit or to utility from beliefs (anticipation).²⁴ Here we see how problems of self-control may also generate incentives to “save the best for last” or to “get the hard part out of the way first”.

Even though our model is flexible enough to explain both positive and negative time preference revealed by consumption choices, it can still 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 render the net effect ambiguous. With sufficient information about preferences, however, unambiguous predictions can be derived. Proposition 1 shows that, given sufficient information about the willpower depletion technology, our model makes strong predictions about qualitative features of the consumption path. It is important to note that choice experiments can provide the requisite information. For example, it is straightforward to design such experiments that would identify whether the rate of willpower depletion depends on the stock of remaining willpower.²⁵

2.5.2 *Almost Never Smooth Even When Feasible*

In Section 2.4 we showed that perfect smoothing, the signature prediction of the standard model, may not occur when $f_W = 0$. Here 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. But if willpower has alternative uses then,

²⁴See Loewenstein and Prelec (1993) for a review.

²⁵Consider the following variation on Baumeister’s two-phase experiments. In the first phase, subjects perform a willpower-depleting activity (A); in the second phase, the same subjects perform another willpower-depleting activity (B). The treatment group performs A before B , the control group performs B before A . All subjects are informed ahead of time about the two activities and their order. If the level of willpower does not affect its rate of depletion then the two groups should, on average, perform equally well on a given activity. If, instead, depletion is anticipated to be more rapid when reserves are lower, then subjects should restrain themselves more on a given activity when it occurs first.

in *almost* all circumstances, an agent 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 so 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 $c(t) = \bar{R}/T$ for $t \in [0, T]$ is feasible, and 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 $c^*(t) = \bar{R}/T$ 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 also $\lambda(t) > 0$ for all $t \in [0, s]$.*

Proof. See the Appendix. ■

Corollary 1 *When $\rho = m(W) = 0$, $\bar{W} < W_H$, and $c(t) \in (0, \bar{c})$, consumption 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 it is interior ($c(t) \in (0, \bar{c})$). Yet we have shown that this description does not preclude intervals of zero consumption. Suppose the agent must spend a fixed amount of time to complete a project, and suppose 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, and 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; 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 that devoting more willpower to controlling alternative urges is always beneficial, and we 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, first suppose that it is optimal to use some amount of willpower to regulate consumption and the remaining amount to regulate other urges. Then it follows that willpower cannot be advantageously re-allocated between the two uses ($\lambda(s) = m'(W(s))$ by condition (A8)). However, we have assumed

that additional willpower devoted to alternative activities is valuable ($m'(\cdot) > 0$); hence $\lambda(s) > 0$. Since this shadow value weakly declines over time (condition (A5)) 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 Section 2.4 that, even if consumption is flat when positive, perfect smoothing still may not occur. Instead it may be optimal instead to consume at a constant but rapid rate, thereby exhausting the cake before T and saving more willpower for alternative activities. Intuitively, this would occur if additional willpower devoted to regulating other urges were sufficiently valuable.

How valuable must willpower be in alternative uses in order for perfect smoothing to fail even though optimal consumption is flat whenever positive? For perfect smoothing to be optimal it must be that, when $c(t) = \bar{R}/T$ 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 shall prove that this latter marginal value can be bounded above by

$$m_H(W, c) = \frac{U(c) - U'(c)c}{f(W, c) - f_c(W, c)c}.^{26}$$

Thus, using \hat{W} to denote the willpower available at T if perfect smoothing has been implemented, and c_H to denote the rate of consumption under perfect smoothing, if $m'(\hat{W}) > m_H(\hat{W}, c_H)$ then perfect smoothing never occurs.²⁷

The following proposition summarizes the necessary conditions derived above and also 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 (a) $\frac{d}{dW} \frac{d}{dc} \ln f(W, c) = 0$ for $c = c_H$ and $W \in [\hat{W}, \bar{W}]$ and (b) $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$. 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 if $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$, then perfect smoothing must occur.²⁸*

Proof. We have already shown 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 we need only derive that $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$. If any willpower is used to restrain consumption, then optimality of this decision requires that $m'(W(s)) \leq \lambda(s)$. Next we use the first-order conditions to establish an

²⁷It is clear that \hat{W} depends on the initial levels of willpower and cake, but we suppress this dependence for simplicity.

²⁸Recalling the case studied in Section 2.4, where $f(W, c) = g(c)$ and $m'(\cdot) = m$, this condition reduces to $m \leq (U(c_H) - c_H U'(c_H)) / (g(c_H) - c_H g'(c_H))$.

upper bound on $\lambda(s)$. Condition (A1) implies that $U'(c(s)) - \lambda(s) f_c \leq \alpha$. Multiplying both sides by $c(s)$ then yields

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

Condition (A6) requires that

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

Combining these two inequalities we obtain

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

To prove that the specified conditions are sufficient for perfect smoothing to be optimal, assume the contrary. The conditions imply that consumption must be constant when positive, so perfect smoothing could only fail because $s < T$. If consumption terminates at $s < T$, then $c = \bar{R}/s > c_H$. Because less restraint will be exercised over a shorter interval $W(s) > \hat{W} > 0$. This in turn has two implications. First, since more willpower will be bequeathed to alternative activities, $m'(\hat{W}) \geq m'(W(s))$. Second, since $W(s) > 0$, (A8) implies that $\lambda(s) = m'(W(s)) > 0$. By hypothesis, $s < T$. Hence, as shown previously, the first-order conditions imply that $\lambda(s) = m_H(W(s), c)$. Because $m_H(\cdot, \cdot)$ 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 conclude that $m'(\hat{W}) \geq m'(W(s)) = \lambda(s) = m_H(W(s), c) > m_H(\hat{W}, c_H)$. But $m'(W(s)) > m_H(\hat{W}, c_H)$, which contradicts $m'(\hat{W}) \leq m_H(\hat{W}, c_H)$. ■

Thus Proposition 3 shows that, when willpower has alternative uses, the necessary conditions for perfect smoothing are restrictive but are 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 a certain 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 Other Models of Self-Control

The preceding sections focused on how the willpower model differs from a standard model of intertemporal choice. In this section we compare our model with existing models of self-control. We first remark that limited willpower provides complementary explanations for the hallmark predictions of existing self-control models: a preference for commitment, profound procrastination, and apparent time inconsistencies. Consider commitment. As in most 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, for if the entire amount were available then not spending it would deplete his valuable willpower. Limited willpower also explains extreme forms of procrastination. Suppose, for example, that our agent must spend a fixed amount of time on an assignment before a certain deadline but can allocate the remaining leisure time optimally. If he is willpower constrained, then he may enjoy leisure early in the program and later work nonstop until the deadline.

Depletable willpower is also consistent with intertemporal preference reversals. When asked to choose between a smaller reward that will arrive immediately and a larger one that will arrive after some delay, an agent with limited willpower may choose the former. The willpower cost of resisting the sooner, smaller reward may outweigh the utility gain from receiving the later, larger one. However, if asked to choose now when both options are temporally distant, the consumer may switch and choose the later, larger reward. He then prefers the later, larger prize since, with commitment, willpower ceases to be required to resist the sooner, smaller one. In this way a willpower-constrained agent may appear to behave in a time-inconsistent manner.

While the willpower model offers complementary explanations of a taste for commitment, profound procrastination, and preference reversals, it also has implications that distinguish it from existing models of self-control. A distinctive 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 and in the domain-specific time preference. The most basic implication of these links was the influence of past consumption of a nonaddictive, storable good on the future consumption of that good. Thus, unlike in models of hyperbolic discounting (Laibson 1997, O'Donoghue and Rabin, 1999), and temptation costs (Gul and Pesendorfer 2001, 2004), if the current choice set (remaining cake and feasible rates of consumption) remains constant then in our model the history of past choices does matter for behavior.

In addition, by linking otherwise unrelated and nonaddictive self-control activities, and by explicitly modeling a willpower reserve, our model can be distinguished from models of cues (Laibson 2001; 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. But in our model, the past exercise of self-restraint will, at least over the short term, make future self-control less likely.

Benabou and Tirole (2004) offer a model of willpower where agents are quasi-hyperbolic discounters with limited knowledge of the degree of their present-bias (the authors' interpretation of willpower). Because agents can infer their own preferences from previous choices, they have an incentive to build self-reputation by resisting temptation. The basic mechanisms in their model are thus quite different from ours and lead to different predictions. One difference is that, 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 situation.

Loewenstein and O’Donoghue (2005) and Fudenberg and Levine (2006) also offer models of costly self-control. In Loewenstein and O’Donoghue’s quite general model, the past exercise of self-restraint affects the future, direct cost of self-control. Their model applies both to intertemporal choice and to static choice under uncertainty. We view our paper as complementing theirs with 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 make all “physical” choices and function as the conduits of flow utility to the long-run self. At a cost, the long-run self may alter the preferences of a short-run self and thus exercise self-control. This model can accommodate the influence of cognitive loads on an agent’s willingness to exercise self-control within a period and its authors discuss how within period linkages between activities requiring self-control may emerge if self-control costs depend on multiple activities in a nonlinear way. Our model differs in that it predicts longer-term effects of shocks to willpower stocks and also of anticipated alternative needs for self-restraint.

In addition, our 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 overconsume.²⁹ Negative time preference is usually viewed as unrelated to self-control and is often attributed (see e.g. Loewenstein 1987) to utility from beliefs (anticipation). Models of anticipatory utility can accommodate increasing consumption paths, but they cannot by themselves explain a taste for commitment, 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 also inducing a preference for “saving the best for last”.

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.³⁰ There are presu-

²⁹A complete discussion of this issue is left to the online appendix at www.umich.edu/~dansilv/research.

³⁰When the willpower remaining after regulating consumption has value, our model is formally analogous to one with direct utility cost because resisting temptation involves (indirect) disutility. (We thank an anonymous referee for pointing this out.) Note, however, that resisting temptation may be costly even if there is no indirect utility from a bequest (see Proposition 2 where $\lambda(t) > 0$). It is important to recall that we interpret the bequest function as giving the value of willpower applied to all subsequent activities in which the agent might benefit from self-control. Thus, consistent with the Baumeister experiments, we predict that a consumer who exerted greater self-control in the consumption problem would exhibit less self-control in subsequent activities.

ably circumstances in which the classical model predicts well. It is therefore an attractive feature of any new model aspiring to supplant a classical one that it identifies the circumstances under which the classical results arise endogenously. If willpower depletion is appended as a constraint, then there will be circumstances where perfect smoothing is optimal; in such cases, consumption behavior predicted by the willpower model will be observationally equivalent to behavior predicted by the classical model.

4 Self-Control Builds Willpower Like a Muscle

Common intuition suggests that exercising self-control is like exercising muscles: after a short while fatigue sets in (willpower depletion) and resolve weakens; but such exercise has a longer-run payoff because it builds future capacity to resist temptation (willpower renewal). Experiments seem to confirm this notion; indeed, the evidence suggests that the willpower built in one domain can be used in another domain (Muraven et al., 1999; Muraven and Baumeister 2000).³¹

Here we return to the special case of Section 2.4 and ask whether it would still be optimal to maintain constant control over consumption of a tempting storable good if exercising self-control today augmented the ability to resist temptation in the future. To do so, we introduce a third state variable, the stock of “muscle” denoted by $M(t)$. Studying behavior from some date ($t = 0$) onward, we assume that the initial stocks of cake, willpower, and muscle are specified; we also assume that the stock of muscle provides an inflow to the willpower stock that can be drawn on in the future. Specifically, we change the transition equation for willpower (equation (A2)) to: $\dot{W}(t) = \gamma M(t) - f(W(t), c(t))$. Restraining consumption depletes willpower as before, but now the self-control muscle provides a flow of willpower in the opposite direction. Thus, given the same stock of willpower and the same resistance to tempting consumption, willpower depletes more slowly the larger is one’s self-control muscle.

The self-control muscle is built by prior exercise of restraint and otherwise weakens with time. Specifically, we assume that the rate at which muscle develops (or deteriorates) is given by $\dot{M} = f(W(t), c(t)) - \sigma M(t)$, so it grows with the exercise of willpower and otherwise decays at rate σ . Thus, if an agent never exercises self-control then his muscle stock will decay exponentially but will remain strictly positive. We shall assume that willpower and muscle are not used in subsequent activities, so the shadow price of additional muscle at time s is zero ($\pi(s) = 0$).

The optimal consumption path in this model shares some qualitative features with that in our basic model without muscle. First, in both formulations, the cake is entirely consumed. Second,

³¹For example, in Muraven et al. (1999), subjects who participated in two-week self-control drills (regulating moods, improving posture) later showed significant increases in the length of time they would squeeze a handgrip.

if willpower has no alternative uses, then perfect smoothing is optimal whenever it is feasible. In particular, for every initial level of muscle there is a willpower level $W_{\bar{H}}$ (possibly zero³²) above which the optimal path entails perfect smoothing.

Recall the case from Section 2.4 in which our assumption that $f_W = 0$ implied it was optimal to restrain consumption of the tempting good equally over time until the cake was exhausted. We now show that the opportunity to build up one's willpower muscle may lead instead to increasing paths of consumption. First we write the Hamiltonian for this problem:

$$H(c(t), R(t), W(t), t, \alpha(t), \lambda(t), \pi(t)) = e^{-\rho t} U(c(t)) - \alpha(t) c(t) + \lambda(t) (\gamma M(t) - f(W(t), c(t))) + \pi(t) (f(W(t), c(t)) - \sigma M(t)).$$

Suppose that consumption is interior and consider the following first-order condition, which is analogous to equation (A1):

$$\underbrace{U'(c(t))}_{\text{direct marginal benefit}} + \underbrace{(-f_c)(\lambda(t) - \pi(t))}_{\text{indirect marginal benefit}} = \underbrace{\alpha(t)}_{\text{marginal cost}}. \quad (4)$$

Equation (4) shows consuming a bit faster at t generates the usual *direct* marginal benefit $U'(c(t))$ and marginal cost $\alpha(t)$; and as before, both are stationary. Therefore, once again, whether optimal consumption is increasing or decreasing locally at t depends on whether the indirect marginal benefit is increasing or decreasing over time.

By consuming faster at t , the agent depletes willpower more slowly at t ($-f_c(W(t), c(t))$). Having additional willpower at t is worth $\lambda(t)$ utils per unit of willpower saved. If $f_W = 0$, then additional willpower is equally valuable whenever it is acquired. Hence, the value of willpower for future use, $\lambda(t)$, is constant over time (there is no magnifier effect).

Although consuming faster at t yields the indirect benefit of freeing up valuable willpower, the net indirect benefit is smaller. The reason is that, in relaxing his self-control, the agent slows his muscle building. Any loss of muscle at t has current and future consequences because a smaller muscle delivers a smaller willpower flow. Relaxing self-discipline at t costs $\pi(t)$ utils per unit of willpower saved at t .

Unlike additional willpower, which is equally valuable whenever acquired, additional muscle is more valuable early in the program because it can then generate a willpower inflow over a longer interval; additional muscle at time s is worthless. The *net* indirect marginal benefit of consuming marginally faster at t is $(-f_c)(\lambda(t) - \pi(t))$ utils. If the value of additional muscle declines monotonically, then the indirect benefit of consuming at a faster rate rises over time and

³²If the initial muscle level is large enough, then the agent will be able to achieve perfect smoothing with $W_{\bar{H}} = 0$.

hence consumption rises. To see this formally, observe that:

$$\dot{\pi}(t) = -\gamma\lambda(t) + \sigma\pi(t).$$

Since $f_W = 0$ it follows that $\lambda(t)$ is constant. Assume that willpower is scarce, and let $\lambda(t) = \lambda > 0$. Now we can solve for $\pi(t)$ to obtain

$$\pi(t) = \lambda \frac{\gamma}{\sigma} \left(1 - e^{\sigma(t-s)}\right), \quad (5)$$

where we have used the fact that muscle has no value at s ($\pi(s) = 0$). Since $t \in [0, s]$, it follows that $\pi(t)$ is decreasing.

Finally, we show that if $f_W = 0$ then consumption must *increase* over time. To see this, take the derivative of equation (4) to obtain

$$\dot{c} = \frac{\dot{\pi}(-f_c)}{U'' - (\lambda(t) - \pi(t))f_{cc}}.$$

The numerator is negative since $\dot{\pi} < 0$ and $-f_c > 0$; the denominator is negative since we are assuming $c(t) \in (0, \bar{c})$ and the second-order condition holds because $c(t)$ maximizes the Hamiltonian. Therefore consumption must increase over time. Relative to the optimal path in the absence of muscle, the ability to build willpower through the exercise of self-control leads the agent to resist temptation strongly at the outset and to relax this restraint over time.³³

5 Conclusion

This paper explored the consequences of including, in a conventional model of intertemporal choice, a cognitive constraint that is consistent with a large body of experimental evidence. Specifically, we assumed that when an agent restricts his consumption of a tempting good, this exercise of self-restraint depletes his finite stock of willpower – a resource useful for regulating urges of all kinds. The model thus captures a common use of the term *willpower*: individuals must draw on a limited well of self-restraint in order to control their impulsive behaviors.

To our surprise, we found that this willpower model accounts not only for prominent anomalies of intertemporal choice, which have been the focus of the self-control literature, but also for other anomalies that are often treated as separate phenomena requiring separate models. Given our results and that taking account of a willpower constraint proved to be so tractable, future research should work to embed a willpower constraint into other economic models where self-control issues

³³If $f_W < 0$, then more intricate consumption patterns arise when there is opportunity to build muscle. A more detailed analysis of this case is available in the online appendix at www.umich.edu/~dansilv/research.

arise. Such modeling would draw out the specific implications of limited but manageable self-control for decisions about work effort, financial planning, unhealthy food consumption, and sexual behavior. As in other areas of the self-control literature, it would also be useful to study the optimal responses of managers, firms, and governments to workers, consumers, and citizens with limited willpower. Drawing out these implications of limited willpower and evaluating them with data will help calibrate the empirical relevance of willpower to economic decision making.

Our analysis suggests two especially important components of future willpower research. First, it will be critical to clarify, through experiment and observational study, the form of the willpower depletion (and replenishment) function and thereby further sharpen the predictions of models like ours. Second, when evaluating the economic importance of a willpower constraint, it will be essential to understand the cumulative consequences of a long sequence of decisions made using temporarily limited, but replenishing, willpower. Whether willpower limitations importantly affect, for example, human capital accumulation, retirement wealth, or health in old age will depend on whether decision makers can adequately anticipate and manage their self-control resources over longer horizons. Therefore, further research should therefore identify the conditions under which the consequences of limited willpower found in our simple setting would be amplified, diminished, or qualitatively changed as consumers face long sequences of linked decisions that require self-control.

6 Appendix

6.1 Motivating the Functional Form of Willpower Depletion

Here we offer a simple model of willpower depletion to motivate the functional form of f we assumed in the main text. Suppose willpower depletion at time t is given by the function $h(W(t), y(t))$, where $W(t)$ is the willpower stock and $y(t)$ is the utility forgone by not consuming at the most tempting feasible rate at t . More precisely, $y(t) = \sup_{x \in B(t)} U(x) - U(c(t))$, where $B(t)$ is the set of feasible consumption rates and $c(t) \in B(t)$ denotes the actual rate of consumption at t . We refer to $B(t)$ as the *menu* at t , and this menu may reflect commitments to limit the rate of consumption at t . We make two primary assumptions about the willpower depletion function.

Assumption 1 *Willpower depletion is a nonnegative, strictly increasing, and convex function of the utility forgone by not consuming at the most tempting feasible rate; and if the agent consumes at the most tempting feasible rate then there is no depletion of willpower. That is, $h(W(t), y(t))$ is nonnegative, strictly increasing, and convex in $y(t)$. Furthermore, $h(W(t), 0) = 0$ for all $W(t) \geq 0$.*

It follows from Assumption 1 that, given $B(t)$, the more the agent restrains his consumption,

the faster his willpower is depleted. It also follows that, if an agent can choose his consumption rate either from a given menu or from a proper subset of that menu, then any consumption rate available on both menus will involve weakly less willpower depletion when it is chosen from the smaller menu. If the most tempting choice from the larger menu isn't on the smaller one, then the agent will strictly prefer to choose a given consumption rate from the smaller menu whenever willpower has scarcity value. In the extreme, the agent prefers commitment to a menu of just one rate over choosing that rate from a larger menu. If no cake remains then Assumption 1 implies that, since the only feasible rate of consumption is zero, consuming nothing depletes no willpower.

The implications of our second substantive assumption were seen in Section 2.5.

Assumption 2 *The same act of self-restraint results in (weakly) faster depletion of willpower if the agent's reserves of willpower are lower. In other words, $h(W(t), y(t))$ is decreasing in $W(t)$.*

In a continuous-time cake-eating problem, if cake remains then it is feasible to consume at any finite rate. This means that, if cake remains, then the utility forgone by restraining consumption is independent of cake size. Hence, if cake remains and the menu is unrestricted, then willpower depletion at t is simply $h(W(t), U(\bar{c}) - U(c(t)))$. This result indicates that, when cake remains, willpower depletion depends only on the current rate of consumption and the willpower stock:

$$f(W, c) = h(W, U(\bar{c}) - U(c)).$$

Our assumptions on h and U imply that f is nonnegative, $f(W, c) = 0$ for $c \geq \bar{c}$; also, f is decreasing in W , strictly decreasing and strictly convex for $c \in [0, \bar{c}]$, and $f_{cW} \geq 0$. When no cake remains, Assumption 1 simply implies that willpower depletion is zero. As the cake shrinks to zero, willpower depletion is strictly positive for zero consumption because $f(W, 0)$ is strictly positive for all $R > 0$. However, when there is no cake remaining ($R = 0$), willpower depletion is zero for zero consumption.

6.2 Details on the Cake-Eating Problem

Problem (P1) from Section 2.3 is nonstandard because willpower depletion is discontinuous. To circumvent this difficulty, we examine a related and standard “free endpoint” problem (P2), which has the same solution as problem (P1). The consumer chooses $c(t) \in [0, \bar{c}]$ to maximize

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

In problem (P2), the agent chooses both an optimal consumption path and a date $s \leq T$ beyond which depletion of each stock ceases ($\dot{W}(t) = \dot{R}(t) = 0$ for $t \in (s, T]$). Problem (P2) therefore allows all the paths that are feasible in (P1) and evaluates each as (P1) does. However, (P2) also allows other 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 must show that these additional paths are never optimal in (P2). This claim is intuitively obvious because such paths are improved by consuming marginally more cake in the neighborhood of s . The additional consumption is feasible since more cake is available; moreover, consuming it would relax the willpower constraint.³⁴

For problem (P2), 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 refer to it as $H(t)$ when no confusion can arise. The first-order conditions include:

$$c(t) \in \begin{cases} \{0\} \\ (0, \bar{c}) \\ \bar{c} \end{cases} \quad \text{and} \quad , e^{-\rho t} U'(c(t)) - \alpha(t) - \lambda(t) f_c \begin{cases} \leq 0, \\ = 0, \\ \geq 0, \end{cases} \quad (\text{A1})$$

$$\dot{W}(t) = -f, \quad (\text{A2})$$

$$\dot{R}(t) = -c, \quad (\text{A3})$$

$$\dot{\alpha}(t) = 0, \quad (\text{A4})$$

$$\dot{\lambda}(t) = \lambda(t) f_W, \quad (\text{A5})$$

$$T - s \geq 0, H(s) - \rho e^{-\rho s} m(W(s)) \geq 0 \text{ and c.s.}, \quad (\text{A6})$$

$$R(s) \geq 0, \alpha(s) \geq 0 \text{ and c.s.}, \quad (\text{A7})$$

$$W(s) \geq 0, \lambda(s) - m'(W(s)) \geq 0 \text{ and c.s.} \quad (\text{A8})$$

In order to isolate the effects of the willpower constraint, we assume that there is no discounting ($\rho = 0$).

Proof of Proposition 2 From (A4), $\alpha(t)$ is a constant denoted, with abuse of notation, as $\alpha \geq 0$. First assume that $\bar{W} \geq W_H$ and consider the case where $\lambda(t) = 0$ for all $t \in [0, s]$. Since $U'(\cdot)$ is strictly positive and stationary, consumption will be constant. Condition (A1) admits three possibilities: (i) $c(t) = 0$ for all t and $\alpha > 0$; (ii) $c(t) = \bar{c}$ for all t and $\alpha \geq 0$; or (iii) $c(t) = c \in (0, \bar{c})$ and $\alpha > 0$ for all t . Other necessary conditions permit us to eliminate the first two possibilities. By (A6), either $c(s) > 0$ and $s = T$ or $c(s) = 0$ and $s \leq T$. We can rule out (i) because it would

³⁴The equivalence of (P1) and (P2) is proved in the online appendix at www.umich.edu/~dansilv/research.

imply that $R(s) = \bar{R} > 0$ and $\alpha > 0$, a violation of (A7). Hence we can rule out (ii) since $\bar{R} < \bar{c}T$ and $c(t) = \bar{c}$ for $t \in [0, T]$ implies that $R(T) < 0$, another violation of (A7). Therefore, (iii) must hold. Since $\alpha > 0$, (A7) implies that $R(T) = 0$ and $c(t) = c \in (0, \bar{c}) = R/T$. By definition of W_H , we have $W(T) \geq 0$. Thus (A8) is satisfied. This proves the first statement in the proposition.

Now assume that $\bar{W} < W_H$. Then $\lambda(t) > 0$ for all $t \in [0, s]$, as we now show. Suppose to the contrary that $\lambda(t) = 0$ for some $t \in [0, s]$. Equation (A5) implies that $\lambda(t)$ is weakly decreasing and can be written as $\lambda(t) = \lambda(0) \exp \left[\int_{n=0}^t f_W(W(n), c(n)) dn \right]$. Since $\exp \left[\int_{n=0}^t f_W(W(n), c(n)) dn \right] > 0$ 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 already show, the preceding conditions imply that $c = \bar{R}/T$ which is infeasible when $\bar{W} < W_H$. Formally, it would result in $W(T) < 0$, violating (A8). So if $\bar{W} < W_H$ then $\lambda(t) > 0$ for all $t \in [0, s]$ and, by (A8), $W(s) = 0$. To satisfy (A1) with $U'(\cdot) > 0$ and $-\lambda f_c > 0$ requires either (i) $\alpha > 0$ or (ii) $\alpha = 0$ but $c(t) = \bar{c}$ for all $t \in [0, s]$. But (ii) cannot occur because no willpower would be used and yet this case requires that the willpower stock be drawn down from a strictly positive level ($\bar{W} > 0$) to $W(s) = 0$. Hence (i) must obtain and since $\alpha > 0$, (A7) requires that the cake is entirely consumed ($R(s) = 0$). In sum, if $\bar{W} < W_H$ then both stocks will be depleted and $\lambda(t) > 0$ for all $t \in [0, s]$, as was to be proved.

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] Fudenberg, Drew and Levine, David K. "A Dual Self Model of Impulse Control." *American Economic Review*, 2006, 96(5):1449-1476.

- [7] Gul, Faruk and Pesendorfer, Wolfgang. "Temptation and Self-Control." *Econometrica*, 2001, 69(6), pp. 1403-35.
- [8] Gul, Faruk and Pesendorfer, Wolfgang. "Self-Control and the Theory of Consumption." *Econometrica*, 2004, 72(1), pp. 119-58.
- [9] 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.
- [10] Kessler, David (2009). *The End of Overeating: Taking Control of the Insatiable American Appetite*. New York: Rodale Books.
- [11] Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 1997, 112(2), pp.443-77.
- [12] Laibson, David. "A Cue-Theory of Consumption." *Quarterly Journal of Economics*, 2001, 116(1), pp.81-119.
- [13] Loewenstein, George. "Anticipation and the Valuation of Delayed Consumption." *Economic Journal*, (1987), 97, 666-684.
- [14] Loewenstein, George and O'Donoghue, Ted. "Animal Spirits: Affective and Deliberative Processes in Economic Behavior." Mimeo, Carnegie Mellon University, 2005.
- [15] Loewenstein, George and Prelec, Drazen. "Preferences for Sequences of Outcomes." *Psychological Review*, 1993, 100(1),91-108.
- [16] 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.
- [17] 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.
- [18] Muraven, Mark; Shmueli, Dikla and Burkley, Edward "Conserving Self-Control Strength. *Journal of Personality and Social Psychology*, 2006, 91(3), pp.. 524-537.
- [19] O'Donoghue, Ted and Rabin, Matthew. "Doing It Now or Later." *American Economic Review*, 1999, 89(1), pp. 103-124.

- [20] 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.
- [21] Thaler, Richard and Shefrin, Hersh. "An Economic Theory of Self-Control," *Journal of Political Economy*, 1988, 89 392-406.
- [22] Vohs, Kathleen and Faber, Ronald. "Spent Resources: Self-Regulatory Resources Availability Affects Impulse Buying," *Journal of Consumer Research*, March 2007, 33 537-547.
- [23] Vohs, Kathleen and Heatherton, Todd. "Self-Regulatory Failure: A Resource-Depletion Approach." *Psychological Science* 2000, 11(3) 249-254.

Figures

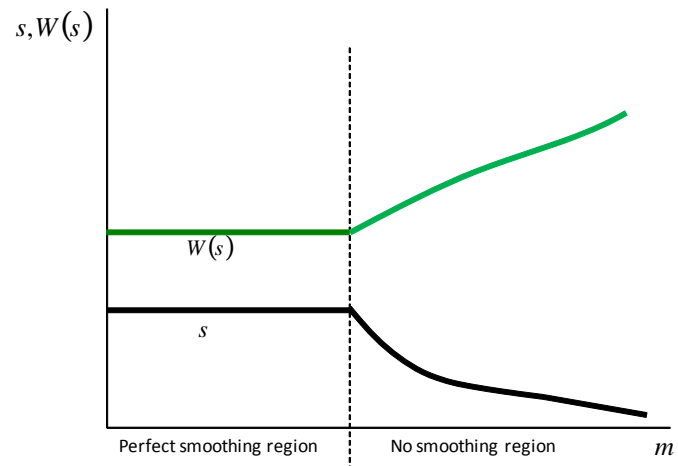


Figure 1: Optimal time (s) to exhaust the cake and optimal willpower bequest ($W(s)$) as a function of the return to willpower applied to alternative activities (m).




Figure 2: Determination of the optimal rate of consumption at time t