HIER

Harvard Institute of Economic Research

Discussion Paper Number 2049

A Dual Self Model of Impulse Control

by

Drew Fudenberg and David K. Levine

November 2004

Harvard University
Cambridge, Massachusetts

A Dual Self Model of Impulse Control[1]

Drew Fudenberg and David K. Levine[2]

First Version: August 5, 2004

This Version: October 19, 2004

Abstract: We propose that a simple "dual-self" model gives a unified explanation for several empirical regularities, including the apparent time-inconsistency that has motivated models of hyperbolic discounting and Rabin's paradox of risk aversion in the large and small. The model also implies that self-control costs imply excess delay, as in the O'Donoghue and Rabin models of hyperbolic utility, and it explains experimental evidence that increased cognitive load makes temptations harder to resist. Finally, the reduced form of the base version of our model is consistent with the Gul-Pesendorfer axioms.

[1] We thank David Laibson and Matt Rabin for many years of instructive conversations on related topics, Stefano DellaVigna and Ulrike Malmendier for helpful discussion of their research, and Drazen Prelec for detailed comments on an earlier draft. We are grateful to NSF grants SES-01-12018, SES-03-14713, and SES-04-26199 for financial support.

[2] Department of Economics Harvard University, and UCLA/Federal Reserve Bank of Minneapolis.

"The idea of self-control is paradoxical unless it is assumed that the psyche contains more than one energy system, and that these energy systems have some degree of independence from each other."
(McIntosh [1969])

1. Introduction

This paper proposes a simple "dual-self" model that gives a unified explanation for a number of empirical regularities. This includes the apparent time inconsistency that has motivated economists' models of hyperbolic discounting: Faced with a choice between consuming some quantity today and a greater quantity tomorrow, some people will choose to consume the lesser quantity today. However, when these same individuals are faced with the choice between the same relative quantities a year from now and a year and a day from now, they choose to consume the greater quantity a year and a day from now.[3] The second regularity is Rabin's [2000] paradox of risk aversion in the large and small. The paradox is that the risk aversion experimental subjects show to very small gambles implies hugely unrealistic willingness to reject large but favorable gambles. In addition, the model explains the effect of cognitive load on self-control that is noted by Shiv and Fedorkin, and it predicts that increased costs of self-control lead to increased delay in stationary stopping-time problems, as in O'Donoghue and Rabin [2001].

Our theory proposes that many sorts of decision problems should be viewed as a game between a sequence of short-run impulsive selves and a long-run patient self. This is consistent with recent evidence from MRI studies such as McClure, Laibson, Loewenstein, and Cohen [2004] that show that short-term impulsive behavior is governed by different areas of the brain than long-term planned behavior. The findings of this research are reinforced also by introspection – we are all aware of the internal conflict when our "rational self" is faced with short-term indulgences that lead to bad long-term consequences. We argue that our theory explains a broad range of behavioral anomalies, and that it is a better fit for the modular structure of the brain than the hyperbolic model,

[3] The economics literature on hyperbolic discounting, following Strotz [1955] and Laibson [1997], uses the now-familiar (β,δ) form. This is called "quasi-hyperbolic discounting" in the psychology literature to distinguish it from the discounting function $f(t)=(1+\alpha t)^{\beta/\alpha}$, which actually is hyperbolic. See Prelec [2004] for a characterization of these functions in terms of "deceasing impatience."

which posits a game between multiple "selves," one in each period. Moreover, our model is analytically simpler than the hyperbolic discounting model: it always yields a unique equilibrium that can be calculated as the solution to a decision problem. The only model of hyperbolic preferences we know of that has similar properties is the Harris and Laibson [2004] model of instantaneous gratification in a consumption-savings problem. That model seems to us to be more complicated and specialized than our own.

In our model the patient long-run self and a sequence of myopic short-run selves share the same preferences over stage-game outcomes; they differ only in how they regard the future. We propose that actual decisions affecting current and future welfare are made by the short-run myopic self. However, the long-term self can manipulate these decisions by controlling the utility of the short-term myopic self. Specifically, we imagine that the short-run myopic self has "base-preferences" in the stage game that depend only on the outcome in the current stage. That is, the short-run players are completely myopic.[4] The stage game is played in two phases. In the first phase, the long-run self chooses the utility function of the myopic self. At some reduction in utility (for both selves – who share the same stage game utility function) the long-run self can choose preferences other than the "base preferences." Physiologically, we imagine that this may take place through the release of chemicals in the brain that determine "mood" and other variables relevant to the preferences of the impulsive myopic self. It is consistent also with the possibility that, with the expenditure of some effort, the long-run self simply takes control. If we wish to model the possibility of imperfection in control over the behavior of the short-run player, this choice of preferences can be stochastic.

In the second phase of the stage game, after the short-run player preferences have been chosen, the short-run player takes the final decision. Notice that we are not allowing the long-run self to precommit for the entire dynamic game – rather she begins each stage game facing the choice which preferences to give the myopic self – or equivalently, how much self-control to exert. Note also that while the hyperbolic discounting model emphasizes the conflict between present and future selves, we emphasize also that the long-run self has the same stage game preferences as the short-run self, and so wishes to serve the interests of future short-term selves.

[4] This is a very stark assumption; we hope to relax it in future research.

Games with long run versus short run players are relatively simple to analyze. This particular class is especially simple. Imposing a minimal perfection requirement that the short-run self must always play a best response, the long-run self implicitly controls the short-run self – albeit at some cost. Equilibria of this game are equivalent to the solution to an optimization problem. In this respect, the long-run versus short-run player model is more conservative than hyperbolic discounting, preserving many of the methods and insights of existing theory, as well as delivering strong predictions about behavior.

Our model is similar in spirit to that of Thaler and Shefrin [1981] (from whom we have taken the McIntosh quotation at the start of the paper.) Like them, we view our model as "providing a simple extension of orthodox models that permits [self-control behavior] to be viewed as rational." One difference is that their model is defined only for the consumption-savings problem we study in section 3, while we develop a more general model that can be applied to other situations. Also, we work with more precise specifications of the costs of self-control, and show how to reduce the game between the selves to a single decision problem. This makes the model analytically tractable, and enables us to make more precise predictions.

Although our point of departure is different, the reduced form of the dual self model is closely connected to the dynamic Gul and Pesendorfer [2002] model. They consider a single player who has preferences over choice sets that includes the desire to limit the available alternatives. Under a particular set of axioms, they show that the decision process can be represented by a utility function with a cost of self-control that is closely connected to ours. Besides presenting a different (and complementary) motivation for this reduced form, this paper advances the Gul-Pesendorfer program in several ways. Although the reduced form of the model leads to the similar decision problem, we have a concrete interpretation of preferences in terms of those of a myopic self, and as a result are able to bring both introspective and physiological evidence to bear on what those preferences might be.

The dual self model predicts the same behavioral anomalies that motivated the hyperbolic discounting model, albeit through a different mechanism. When dealing with decisions that effect only future options, the short-run self is indifferent, hence can be manipulated by the long run self at minimal cost. The long-run self, then, has two different sorts of mechanisms through which to change the behavior of future short-run selves. She can intervene directly in a future stage game by choosing an appropriate utility function, but to do so requires a substantial utility cost. Alternatively, in some settings it may be possible for the current short-run self to limit the alternatives available to the future short-run selves; manipulating these decisions has negligible cost. For example, in the classical choice problem that gave rise to the theory of hyperbolic discounting, faced with a decision between consuming something today or more tomorrow, the long-run decision maker may prefer to consume more tomorrow. But to manipulate the current short-run self into making that decision may be more costly than it is worth. On the other hand, when faced with making a commitment to the same decision involving future dates, the cost of manipulating the decision of the current short run self is negligible, since the current short-run self does not care about the future.[5]

For a less trivial application, we examine a simple one-person savings problem. We show that if the short-run self has access to all available wealth, the savings rate is reduced to keep the cost of self-control low. On the other hand, when wealth is kept in a bank account, and the short-run self that withdraws the money is different from the short-run self who (at a later time) spends the money, savings are exactly those predicted in the absence of self-control costs. However, the dual self model predicts that the propensity to spend out of unanticipated cash receipts is greater than out of unanticipated bank-account receipts. In particular, a sufficiently small unanticipated cash receipt will be spent in its entirety.

This high propensity to spend from cash receipts underlies our explanation of risk aversion in the large and small. Winnings from sufficiently small cash gambles are spent in their entirety, and so are evaluated by the short-run self's preferences, which are over consumption. This is consistent with experimental evidence showing that – when wealth is given by the amount of money on the table, and not by lifetime wealth – subject behavior is consistent with sensible expected utility functions. In the dual self model,

[5] We should point out that this discussion implicitly supposes that the short-run self is directly motivated by money, even though the only real consequence of earning money is the future consumption that it brings. This is consistent with evidence (such as Pavlov's bell) that the impulsive short-run self responds to learned behavioral cues in addition to direct stimulus. Modern physiological research is making progress in identifying some of the brain chemistry that reflects the response to these stimuli, see, for example, Haruno et al [2004]. Camerer, Lowenstein and Prelec [2000] say that "roughly speaking, it appears that similar brain circuitry—dopaminergic neurons in the midbrain—is active for a wide variety of rewarding experiences (including) money rewards." The conclusion speculates about some possible extensions of our model to learned cues.

when the stakes are large, self-restraint kicks in, part of the winnings will be saved and spread over the lifetime, and the global implications of the theory explain the paradox proposed by Rabin [2000]. Note that these results are in the opposite direction of Gul and Pesendorfer [2002, 2004], who do not consider environments with mechanisms such as banks that substitute for self-control.

We also apply the dual-self model to the study of procrastination and delay in a stationary stopping-time environment that is very similar to that of O'Donoghue and Rabin [2001]. Like them, we find that self-control costs lead to longer delays, but our model yields a unique prediction, in contrast to their finding of multiple equilibria. Our model also suggests some qualifications to the interpretations that DellaVigna and Malmendier [2003] give to their data on health-club memberships.

2. The Model

Time is discrete and potentially unbounded $t = 1,2,\ldots$. There is a fixed, time-invariant set of actions A for the short-run selves; this is assumed to be a compact subset of Euclidean space. There is also a measure space Y of states, and a set R of self-control actions for the long-run self, a compact convex subset of Euclidean space. The point $0 \in R$ is taken to mean that no self-control is used. A finite history of play $h \in H$ consists of the past states and actions, $h = (y_1, a_1, r_1, \ldots, y_t, a_t, r_t)$ or the null history 0. The length of the history is denoted by $t(h)$, the final state in h by $y(h)$. There is an initial state y_1. The probability of states at time $t+1$ is determined by the current state and action y_t, a_t according to the exogenous probability measure $\mu(y,a)$. Note that the long-run self's action r has no effect on the future state. The game is between the long-run self, whose pure strategies are maps from histories and the current game to self-control actions $\sigma_{LR} : H \times Y \to R$, and the sequence of short-run selves. Each short-run self plays in only one period, and observes the self-control action chosen by the long-run self prior to moving. Denote by H_t the set of t-length histories H_t. A strategy from the time-t short-run self is a map $\sigma_t : H_t \times Y \times R \to A$; we denote the collection of all of these strategies by σ_{SR}. The strategies together with the measure μ give rise to a measure π_t over histories of length t.

The utility of the short-run self is given by $u(y,r,a)$: the long-run player's current self-control action influences the short-run player's payoff. The utility of the long-run self is given by

$$U = \sum_{t=1}^{\infty} \delta^{t-1} \int u(y,r,a) d\pi_t(y(h)).$$

In this formulation, the self-control cost (that is, the difference between $u(y,r,a)$ and $u(y,0,a)$) is borne by both selves. However, since the short-run self cannot influence that cost, all that matters is the influence of self-control on the marginal incentives of the short-run self and thus on its decisions. In some cases, self-control simply takes the form of constraints on the set of actions that are feasible for the short-run self; in these cases we can think of the short run self as maximizing $u(y,0,a)$ given the constraints.

Since each move begins a proper subgame, the strategies are a *subgame perfect equilibrium* if each self's strategy is optimal following every history, and the short-run self's strategy is also optimal following the move of the long-run self.

Assumption 1 (Costly Self-Control): *If* $r \neq 0$ *then* $u(y,r,a) < u(y,0,a)$.

Assumption 2 (Unlimited Self-Control): *For all* y,a *there exists* r *such that for all* a', $u(y,r,a) \geq u(y,r,a')$.

Under the assumptions that self-control is costly and unlimited, we may define the cost of self-control

$$C(y,a) \equiv u(y,0,a) - \sup_{\{r \mid u(y,r,a) \geq u(y,r,\cdot)\}} u(y,r,a)$$

Assumption 3 (Continuity): $u(y,r,a)$ *is continuous in* r,a.

This assures that the supremum in the definition of C can be replaced with a maximum.

Assumption 4 (Limited Indifference): *for all* $a' \neq a$, *if* $u(y,r,a) \geq u(y,r,a')$ *then there exists a sequence* $r^n \to r$ *such that* $u(y,r^n,a) > u(y,r^n,a')$.

This means that when the short-run self is indifferent, the long-run self can break the tie for negligible cost.

Notice that from Assumptions 1 and 2, if $a = \arg\max(u(y,r,a))$ then $C(y,a) = 0$, and $C(y,a') > 0$ for $a' \neq a$. In addition by Assumption 3, $C(y,a)$ is continuous in a. Conversely, if we have given functions $u(y,0,a)$ and $C(y,a)$ satisfying these properties, then we can take $R = A$ and construct

$$u(y,r,a) = \begin{cases} u(y,0,r) - C(y,r) - \| r - a \| & a \mid u(y,0,a) \geq u(y,0,r) \\ u(y,0,a) - C(y,r) - \| r - a \| & a \mid u(y,0,a) < u(y,0,r) \end{cases}$$

which gives rise to the target cost of self-control function, while satisfying Assumptions 1-4.

By way of contrast, we consider also the following *reduced form* optimization problem, of choosing a strategy from histories and states to actions $\sigma_{SC} : H \times Y \to A$ to maximize the objective function

$$\sum_{t=1}^{\infty} \delta^{t-1} \int [\, u(y,0,a) - C(y,a) \,] \, d\pi_t(y(h)) \, .$$

Theorem 1 (Equivalence of Subgame Perfection to the Reduced Form): *A subgame perfect equilibrium exists if and only if a solution to the reduced form problem exists; if a solution exists then for every optimal* σ_{SC} *there are equilibrium strategies* σ_{LR}, σ_{SR} *such that* $\sigma_{SC} = \sigma_{LR} \circ \sigma_{SR}$ *and vice versa.*

Remark: We have not imposed sufficient assumptions on Y and μ to guarantee the existence of a solution to the optimization problem. If Y is finite, it is well known that this problem has a solution; however we wish to examine cases where Y is infinite, and although in our examples existence of a solution is unproblematic, it is complicated to give general conditions guaranteeing the existence of an optimum in the infinite case.

Finally, we wish to consider a tightly parameterized functional form for the cost of self-control, namely that it is proportional to the difference in utility between the best available action and that actually taken.

Assumption 5 (Linear Self-Control Cost): $C(y,a) = \gamma \left[\max_{a'} u(y,0,a') - u(y,0,a) \right]$, so that

$$
\begin{aligned}
U &= \sum_{t=1}^{\infty} \delta^{t-1} \int [u(y,0,a) - C(y,a)]\, d\pi_t(y(h)) \\
&= \sum_{t=1}^{\infty} \delta^{t-1} \int [(1+\gamma)u(y,0,a) - \gamma \max_{a'} u(y,0,a')]\, d\pi_t(y(h))
\end{aligned}
$$

Notice that the introduction of self-control costs into the optimization problem means that independence of irrelevant alternative generally fails. When Assumption 5 holds, for example, $\max_{a'} u(y,0,a')$ contributes to utility, even though the argmax that determines this maximum is generally irrelevant. On the other hand, under Assumption 5, the marginal cost of self-control does satisfy the independence of irrelevant alternatives. That is, improving the best available alternative lowers utility, but does not change the marginal cost of self-control. In this sense, Assumption 5 is conservative, maintaining as much of the independence of irrelevant alternatives as is consistent with an interesting theory of self-control. It is also important to note that self-control costs that satisfy Assumption 5 are shown by Gul and Pesendorfer [2001, 2002] to satisfy their axioms. Assumption 5 is also analytically convenient. On the other hand, Section 6 argues that some evidence is more consistent with strictly convex costs of control.

3. A Simple Savings Model

To start, consider the simple case of an infinite-lived consumer making a savings decision. The state $y \in \Re_+$ represents wealth, which may be divided between consumption and savings according to the action $a \in [0,1]$ representing the savings rate. Borrowing is not allowed. Savings are invested in an asset that returns wealth $y_{t+1} = R a_t y_t$ next period, there is no other source of income.

In each period of time, the base preference of the short-run self has logarithmic utility,[6]

$$u(y,0,a) = \log((1-a)y)$$

where we define $\log(0) = -\infty$.

The short-run self wishes to spend all available wealth on consumption. We assume a separable cost of self-control, so

---

[6] The appendix presents the extension to CES form of which logarithmic utility is a special case.

$C(y,a) = \gamma[\max_{a'} u(y,0,a') - u(y,0,a)] = \gamma(\ln(y) - \ln(1-a)y)) = -\gamma\ln(1-a)$.

The reduced form for the long-run self has preferences

$$\begin{aligned} U &= \sum_{t=1}^{\infty} \delta^{t-1} \int [u(y_t,0,a_t) - C(y_t,a_t)]\, d\pi_t(y(h)) \\ &= \sum_{t=1}^{\infty} \delta^{t-1} [(1+\gamma)\log((1-a_t)y_t) - \gamma\log(y_t)] \end{aligned}$$

The long-run self's problem is thus to maximize this function subject to the wealth equation $y_t = Ra_{t-1}y_{t-1}$. It is shown in the Appendix that there is a solution, and that the solution has a constant savings rate strictly between zero and one.[7] Thus we compute present value utility for constant savings rates, and maximize

$$\begin{aligned} U &= \sum_{t=1}^{\infty} \delta^{t-1} \Big[(1+\gamma)\Big(\log(1-a) + (t-1)\log Ra + \ln y_0\Big) - \gamma(t-1)\log(Ra) - \gamma\log(y_0)\Big] \\ &= \sum_{t=1}^{\infty} \delta^{t-1} \Big[(1+\gamma)\Big(\log(1-a)\Big) + (t-1)\log(Ra) + \log(y_0)\Big]; \end{aligned}$$

using $\sum_{t=1}^{\infty} t\delta^t = \frac{\delta}{(1-\delta)^2}$, this simplifies to

$$\begin{aligned} U &= \frac{\Big[(1+\gamma)\Big(\log(1-a)\Big) + \log(y_0)\Big]}{(1-\delta)} + \frac{\log(Ra)}{(1-\delta)^2} - \frac{\log(Ra)}{(1-\delta)} \\ &= \frac{\Big[(1+\gamma)\Big(\log(1-a)\Big) + \log(y_0)\Big]}{(1-\delta)} + \frac{\delta\log(Ra)}{(1-\delta)^2}. \end{aligned} \qquad (1)$$

To find the constant savings rate that maximizes U, we differentiate (1) to find

$$dU/da = \frac{-(1+\gamma)}{(1-a)(1-\delta)} + \frac{1}{a(1-\delta)^2} - \frac{1}{a(1-\delta)},$$

7 The reason a proof is required is that the objective function is not globally concave. We can extend the conclusion that savings are a constant fraction of wealth to the case where asset returns $\tilde{R}$ are stochastic and i.i.d. provided that there is probability 0 of 0 gross return. In the more general CES case studied in the Appendix, the solution given there remains unchanged provided we define $R^{1-\rho} = E\left(\tilde{R}^{1-\rho}\right)$.

so
$$a = \frac{\delta}{1+\gamma-\delta\gamma} \tag{2}$$

Since the solution must be interior, it must satisfy the first-order condition, and since there is a unique solution to the first order condition, this is the optimum

The comparative statics are immediate and intuitive: As γ increases, so self-control becomes more costly, the savings rate is reduced, to avoid the cost of self-control. As the long-run player becomes more patient, (as δ increases) this cost of future self control becomes more important, so the effect of γ increases, which tends to increase the difference between the savings rate at a fixed γ and that at $\gamma = 0$. (In particular, γ is irrelevant when $\delta = 0$, as the savings rate is 0 with or without costs of self-control.) However, increasing δ also increases the savings rate for any fixed γ, as is the case when $\gamma = 0$ and there is no self-control problem. This latter effect dominates, as total saving increases.

Note that when $\gamma = 0$, so there are no self-control costs, the optimum savings rate is $a^* = \delta$. In this case the agent's lifetime utility as a function of initial wealth y_0 is

$$\frac{\left[\log(1-\delta) + \log(y_0)\right]}{(1-\delta)} + \frac{\delta \log(R\delta)}{(1-\delta)^2} \tag{3}$$

we use this fact in the following section.

These results should be contrasted with predictions generated by models of hyperbolic discounting. As Harris and Laibson [2004] emphasize, in the usual discrete-time hyperbolic discounting model, consumption need not be monotone in wealth, even in a stationary infinite-horizon environment. In contrast, our conclusion of a constant savings rate obtains for any specification of the per-period utility, as it derives from the stationarity of the problem. Second, the hyperbolic model typically has multiple equilibria (Krusell and Smith [2000]), which complicates both its analysis and its empirical application.

In response, Harris and Laibson [2004] propose a continuous-time model of the consumption-savings problem, where the return on savings is a diffusion process. They show that the equilibrium is unique in the limit where individuals prefer gratification in the present discretely more than consumption in the only slightly delayed future.[8] Moreover, in our case of constant return on assets, their results show that consumption is a constant fraction of wealth if the discount factor is sufficiently close to 1. Thus the limit form of their model makes qualitatively similar predictions to ours; we feel that our approach is more general and more direct.

4. Simple Banking Model

In practice there are many ways of restraining the sort-run self besides the use of self-control: the obvious thing to do is make sure that the short-run self does not have access to resources that would represent a temptation. Here we consider the consequences of a simple model in which basic savings decisions are made in a context (the bank) where consumption temptations are not present. In the bank, the decision is made how much "pocket cash" to make available for spending when a consumption opportunity arises in the following period. Since savings decisions are made in the bank, with perfect foresight, the optimum without self-control can be implemented simply by rationing the short-run self. In this sense, the model has an equilibrium equivalent to a model without a self-control problem. However, the reaction of the consumer to unanticipated cash receipts (or simply uncertain cash receipts) is quite different than the reaction to anticipated receipts, or unanticipated bank account receipts: the propensity to consume out of a small unanticipated cash receipt is 100%, while the propensity to consume out of a similar amount of money received in the bank account (for example, a small capital gain on a stock) is small.

This wedge between the propensity to consume out of pocket cash and to consume out of bank cash has significant implications for "risk aversion in the large and small." Winnings from sufficiently small cash gambles are spent in their entirety, and so are evaluated by the short-run self's preferences, which are over consumption. This is consistent with experimental evidence showing that – when wealth is given by the amount of money on the table, and not by lifetime wealth – subject behavior is consistent with sensible expected utility functions. When the stakes are large, self-restraint kicks in,

[8] To do this, they show that equilibrium is characterized by the solution of a single-agent problem, where the agent's utility function is derived from the shadow values in the original problem. When base preferences are CRRA, the only difference between the derived utility function and that of a "fully rational" agent (an exponential discounter) is that the agent gets a utility boost at zero wealth,

part of the winnings will be saved and spread over the lifetime, so the model explains the paradox proposed by Rabin [2000]. Note that these results are in the opposite direction of Gul and Pesendorfer [2002, 2004], who do not consider environments with mechanisms such as banks that substitute for self-control.

The implication of the pocket versus bank cash model are very important in the interpretation of experimental results: in experiments the stakes are low, but individuals demonstrate substantial curvature in the utility function. Besides exhibiting risk aversion, when given the opportunity, for example, to engage in altruistic behavior, they generally do not make the minimum or maximum donation, but some amount in between. (Similar behavior is observed on the street: many people will make a positive donation to a homeless person, but few will empty their pockets of all cash.) If utility is viewed in terms of wealth, this type of behavior makes little sense, since the effect of a small donation on the marginal utility of wealth to either the donor or recipient is miniscule. Viewed in terms of pocket cash, the relevant point of comparison when there is a wedge due to the rationing of cash to the short-run self, this behavior makes perfect sense.

Formally, we augment the simple saving model by supposing that each period consists of two subperiods, the "bank" subperiod and the "nightclub" subperiod. During the "bank" subperiod, consumption is not possible, and wealth y_t is divided between savings s_t, which remains in the bank, and cash x_t which is carried to the nightclub. In the nightclub consumption $0 \leq c_t \leq x_t$ is determined, with $x_t - c_t$ returned to the bank at the end of the period. Wealth next period is just $y_{t+1} = R(s_t + x_t - c_t)$. For compatibility, we take the discount factor between the two consecutive nightclub periods to be δ; preferences continue to have the logarithmic form.[9] First, as in the simple savings problem, we consider the perfect foresight problem in which savings are the only source of income. Since no consumption is possible at the bank, the long-run self gets to call the shots; and the long-run self can implement a^*, the optimum of the problem without self-control, simply by choosing pocket cash $x_t = (1 - a^*)y_t$ to be the desired consumption. We will see shortly that the short-run self will in fact spend all the pocket cash; that having solved the optimum without self-control, the long-run self does not in fact wish to exert self-control at the nightclub.

[9] The appendix provide the parallel computations for the CES case.

Now we turn to the problem of stochastic cash receipts (or losses). That is, we suppose that at the nightclub in the first period there is a small probability the agent will be offered a choice between several lotteries. Let $\tilde{z}_1$ be the chosen lottery. If the choices are themselves drawn in an i.i.d. fashion, this will also result in a stationary savings rate that is slightly different from the a^* computed above, but if the probability that a non-trivial choice is drawn is small, the savings rate will be very close to a^*. We find it easier to consider the limit where the probability of drawing the gamble is zero, and we do not need to engage in an elaborate computation to find a savings rate close to but not exactly equal to a^*.

For the agent to evaluate a lottery choice $\tilde{z}_1$, he needs to consider how he would behave conditional on each possible realization z_1. The short-run self is constrained to consume $c_1 \leq x_1 + z_1$. Next period wealth is given by

$$y_2 = R(s_1 + x_1 + z_1 - c_1) = R(y_1 + z_1 - c_1)$$

The utility of the long-run self starting in period 2 is given by the solution of the problem without self control, as in equation (3):

$$U_2(y_2) = \frac{1}{(1-\delta)}\left(\log(1-\delta) + \log(y_2) + \frac{\delta}{1-\delta}\log(R\delta)\right)$$

The utility of both selves in the first period is $(1+\gamma)\ln(c_1) - \gamma\ln(x_1 + z_1)$, and so the overall objective of the long-run self is to maximize

$$\begin{aligned}&(1+\gamma)\log(c_1) - \gamma\log(x_1 + z_1)\\&+\frac{\delta}{(1-\delta)}\left(\log(1-\delta) + \log(R(y_1 + z_1 - c_1)) + \frac{\delta}{1-\delta}\log(R\delta)\right)\end{aligned} \tag{4}$$

The first order condition for optimal consumption is

$$\frac{1+\gamma}{c_1} = \frac{\delta}{(1-\delta)(y_1 + z_1 - c_1)}$$

so

$$c_1\delta = (1-\delta)(y_1 + z_1 - c_1) + \gamma(1-\delta)(y_1 + z_1 - c_1)$$

and

$$c_1 = \frac{(1-\delta)(1+\gamma)(y_1+z_1)}{\delta+(1+\gamma)(1-\delta)} = \left(1-\frac{\delta}{\delta+(1+\gamma)(1-\delta)}\right)(y_1+z_1) \qquad (5)$$
$$\equiv (1-B)(y_1+z_1)$$

Note that when $\gamma = 0$, (5) simplifies to $c_1 = (1-\delta)(y_1+z_1)$, as it should. If the solution c_1 satisfies the constraint $c_1 \le x_1 + z_1$ it represents the optimum; otherwise the optimum is to consume all pocket cash, $c_1 = x_1 + z_1$. Because x_1 is the solution for $\gamma = 0$, we know that $x_1 = (1-\delta)y_1$. Thus $c_1 \le x_1 + z_1$ if $z_1 < z_1^*$, where the critical value of z_1^* is derived from

$$\left(1-\frac{\delta}{\delta+(1+\gamma)(1-\delta)}\right)(y_1+z_1^*) = (1-\delta)y_1 + z_1.$$

This yields

$$(1-B)(y_1+z_1^* = (1-\delta)y_1 + z_1^*$$
$$z_1^* = (\delta / B - 1)y_1$$
$$z_1^* = \gamma(1-\delta)y_1$$

Since $B < \delta$ when $\gamma > 0$ we see that for a range of positive z_1 it is in fact optimal to spend the entire amount of pocket cash x_1. Note also that when $\gamma = 0$, so there is no self-control problem, so $z_1^* = 0$: it is never optimal to spend all of the increment to wealth.

The above establishes

Theorem 2: If $z_1 < z_1^*$, overall utility is

$$\log(x_1 + z_1) + \frac{\delta}{(1-\delta)}\left(\log(1-\delta) + \log(R(y_1 - x_1)) + \frac{\delta}{1-\delta}\log(R\delta)\right) \qquad (6)$$

If $z_1 > z^*$ utility is

$$
\begin{aligned}
&(1+\gamma)\log\left(\frac{(1-\delta)(1-\gamma)}{1+\gamma(1-\delta)}(y_1+z_1)\right)-\gamma\log(x_1+z_1) \\
&+\frac{\delta}{(1-\delta)}\left(\log(1-\delta)+\log\left(\frac{R\delta}{1+\gamma(1-\delta)}(y_1+z_1)\right)+\frac{\delta}{1-\delta}\log(R\delta)\right)
\end{aligned} \tag{7}
$$

To analyze risk aversion, imagine that $\tilde{z}_1 = \overline{z} + \sigma\varepsilon_1$, where ε_1 has zero mean and unit variance, and suppose that σ is very small. Now consider the usual conceptual experiment of comparing a lottery with it certainty equivalent. For $\overline{z} < z^*$ overall payoff is given by (6). Thus relative risk aversion is constant and equal to ρ, where wealth is $w = x_1 + \overline{z}_1$ so risk is measured relative to pocket cash. On the other hand, for $\overline{z} > z^*$, the utility function (7) is the difference between two others, one of which exhibits constant relative risk aversion relative to wealth $y_1 + \overline{z}$, the other of which exhibits constant risk aversion relative to pocket cash $x_1 + \overline{z}$. When γ is small, the former dominates, and to a good approximation for large gambles risk aversion is relative to wealth, while for small gambles it is relative to pocket cash.[10]

We can see this effect graphically in the case of the Rabin Paradox. In Rabin [2000] it is shown with respect to the standard theory that

> "Suppose we knew a risk-averse person turns down 50-50 lose $100/gain $105 bets for any lifetime wealth level less than $350,000, but knew nothing about the degree of her risk aversion for wealth levels above $350,000. Then we know that from an initial wealth level of $340,000 the person will turn down a 50-50 bet of losing $4,000 and gaining $635,670."

The point being of course that many people will turn down the small bet, but no one would turn down the second. In our model, however, we can easily explain these facts, with, say, logarithmic utility.

The first bet is most sensibly interpreted as a pocket cash gamble – the experiments with real monetary choices in which subjects exhibit similar degrees of risk aversion over similar stakes certainly are. Moreover, if the agent is not carrying $100 in cash, then there may be a transaction cost in the loss state reflecting the necessity of

[10] The former also dominates as δ approaches $1/R^{1-\rho}$, which is relevant when $\rho \leq 1$.

finding a cash machine or bank. The easiest calculations are for the case where the gain \$105 is smaller than the threshold z^*. In this case, logarithmic utility requires the rejection of the gamble if pocket cash x_1 is \$2100 or less. That is, $.5\log(2100-100)+.5\log(2100+105)\approx\log(2100)$. If in fact the gain of \$105 is to be smaller than the threshold z^*, we require $\gamma \geq 105/x_1$, so if $x_1 = 2100$ is pocket cash, we need a $\gamma > .05$, while for[11] $x_1 = 300$, γ must be of at least 0.35. However, these calculations are quite robust to the presence of the threshold. For example, if pocket cash is \$300 and $\gamma = 0.05$, and the wealth corresponding to \$300 pocket cash is \$300,000, then the favorable state of \$105 will be well over the threshold of \$15.[12] Nevertheless, a computation shows that the gamble should be rejected, and indeed is not close to the margin. Indeed, the disutility of the \$100 loss relative to pocket cash of \$300 is so large, that even a very flat utility for gains is not enough to offset it, that even if we bound the utility of gains by replacing the logarithmic utility with its tangent above \$300, not only should this gamble be rejected, but even a gamble of lose \$100, win \$110 should be rejected.

Turning to the large stakes gamble, unless pocket cash is at least \$4,000, the second gamble must be for bank cash; for bank cash, the relevant parameter is wealth, not pocket cash. It is easy to check that if wealth is at least \$4,026, then the second gamble will always be accepted. So, for example, an individual with pocket cash of \$2100, $\gamma = 0.05$ and wealth of more than \$4,026 will reject the small gamble and take the large one, as will an individual with pocket cash of \$300, $\gamma = 0.05$ and wealth equal to the rather more plausible \$300,000.

We should acknowledge that other models can yield these results. In particular, in this specific case, the hyperbolic discounting model yields the identical prediction about bank savings in the first period and second periods, and thus about the response to unanticipated cash shock. To see this, note that our model of response to an unanticipated shock is to maximize the utility function

$$(1+\gamma)\log c_1 - \gamma\log(x_1+z_1) + \delta U_2(R(y_1+z_1-c_1)).$$

---

[11] The usual daily limit in the U.S. for ATM withdrawals is \$300.

[12] The relationship between pocket cash and wealth depends on δ and hence on the period length.

Denoting the hyperbolic discount factor by β, the hyperbolic discounting model says that the response to an unanticipated shock is to maximize the utility function

$$\log c_1 + \delta\beta U_2(R(y_1 + z_1 - c_1)\,.$$

In both cases, the utility function U_2 is the utility function derived by solving the unconstrained problem, which is the same in the two cases and equal to the utility function of an agent without self-control problems ($\beta = 1$ or $\gamma = 0$.) Since $x_1 + z_1$ is not a decision variable at this "nightclub" stage of the problem, we see that if $\beta = 1/(1+\gamma)$ the two objective functions differ only by a linear transformation, and so necessarily yield the same preferences over lotteries at the nightclub stage,

The analysis so far has supposed that cash is only available at the banking stage. If the agent, when banking, anticipates the availability of $300 from an ATM during the nightclub stage, it is optimal to reduce pocket cash by this amount. Of course if the goal is to have pocket cash less than $300, then self-restraint will be necessary in the presence of cash machines. Note that this explains why we find cash machines where impulse purchases are possible: where lottery tickets are sold, for example. In equilibrium, few if any additional overall sales are induced by the presence of these machines, since their presence is anticipated, but of course the competitor who fails to have one will have few sales. So one consequence of the dual self-model is that we may see an inefficiently great number of cash machines.[13]

Credit cards and checks also pose complications in applying the theory, as for many people the future consequences of using credit cards and checks can be significantly different than the expenditure of cash. That is, it is one thing to withdraw the usual amount of money from the bank, spend it all on the nightclub and skip lunch the next day. It is something else to use a credit card at the nightclub, which, in addition to the reduction of utility from lower future consumption, may result also in angry future recriminations with one's spouse, or in the case of college students, with the parents who pay the credit card bills. So for many people it is optimal to exercise a greater degree of self-control with respect to non-anonymous expenditures such as checks and credit cards,

[13] Of course businesses engage in a variety of methods to induce impulse purchasing, for example a car salesman's offer to let the purchaser drive away in the car right now.

than it is with anonymous expenditures such as cash. This conclusion is consistent with the finding of Wertenbroch, Soman, and Nunes [2002] that individuals who are purchasing a good for immediate enjoyment have a greater propensity to pay by cash, check or debit card than by credit card.

The implications of the theory for experiments are ambiguous and complicated. On the one hand the theory explains why we see substantial risk aversion in the laboratory. On the other hand the theory also predicts a high degree of idiosyncrasy in that risk aversion. It will depend, for example, on such factors as how much cash the subjects are carrying with them, the convenience of nearby cash machines and the like.

Finally, we should point out that even without a self-control problem, fear of theft can also lead agents to impose binding constraints on their ability to draw against wealth in nightclub periods, and so predicts that unanticipated losses must be absorbed from consumption. This fear-of-theft model predicts that unanticipated gains will be treated the same regardless of whether they are received in cash or in the bank, while the self-control model does not. It is true that the treatment of losses that is relatively more important for resolving the paradox of high risk aversion for small stakes gambles, but for choices between gambles that have only gains, (the usual laboratory case) the "fear of theft" model predicts little risk aversion, where the dual self and hyperbolic discounting model predict that risk aversion will continue to be substantial.

5. Procrastination and Delay

We can also use the dual-self model to study procrastination and delay. Consider the following model: Every period $t = 1, 2, \ldots$ the short-run self must either take an action or wait. Waiting allows the self to enjoy a leisure activity that yields a stochastic amount of utility x_t, and the problem repeats next period; we suppose that the x_t are i.i.d. with fixed and known cumulative distribution function *P* and associated density *p* on the interval $[\underline{x}, \overline{x}]$. Taking the action ends the game and results in a flow of utility v beginning next period, and so gives a present value of

$$V = \frac{\delta}{1-\delta} v .$$

We assume that $\delta v > \delta\mu + \underline{x}$, where $\mu = Ex$; so that the payoff of the strategy "never act" is less than the payoff of the strategy "act in the first period if x is equal or near to $\underline{x}$;" when the reverse inequality holds, an agent without self-control costs will never act. We also assume that $\delta v < \overline{x}$, so that an agent without self-control would choose to delay when x is close to $\overline{x}$.

Except for the use of the dual-self model, this setup is very similar to that of O'Donoghue and Rabin [2001], who consider hyperbolic preferences. We compare the two models after deriving our conclusions.

Because the current value of x has a monotone effect on the payoff to waiting, and no effect on the payoff to doing it now, the optimal solution is a cutoff rule: If $x \geq x^*$ then wait, and if $x < x^*$ take the action. The expected present value tomorrow if the action is not taken today and the cutoff rule is used is given by

$$W = P(x*)\left(-\gamma E(x \mid x < x^*) + V\right) + (1 - P(x^*))(E(x \mid x > x^*) + \delta W).$$

Theorem 3 *$\underline{x} \leq x^* < \overline{x}$; if $\delta v > \delta\mu + (1+\gamma)\underline{x}$ then $x^* > \underline{x}$.*

Proof: Suppose that the optimum is at $x^* = \overline{x}$, and that $x_t = \overline{x}$. Then doing it now gives payoff of $\delta V - \gamma\overline{x}$, while waiting one period and then conforming to the presumed optimal rule gives $\overline{x} + \delta^2 V - \delta\gamma\mu$, because the agent is certain to act next period. Thus waiting is optimal unless $\delta(1-\delta)V \geq \overline{x} + \gamma(\overline{x} - \mu)$, but this contradicts $\delta v < \overline{x}$. Now suppose $x^* = \underline{x}$ and $x_t = \underline{x}$. Conforming to the strategy yields payoff $\underline{x} + \delta\mu/(1-\delta)$, while acting yields $\delta V - \gamma\underline{x}$.

☑

If the optimal cutoff x^* is in the interior of $[\underline{x}, \overline{x}]$, optimality implies that the agent is indifferent between waiting and acting when $x = x^*$: $-\gamma x^* + V = x^* + \delta W$, so

$$x^*(1+\gamma) = V - \delta W =$$

$$\frac{(1-\delta)V}{(1-\delta(1-P(x^*)))} - \delta\left(\frac{\left(P(x^*)(-\gamma E(x \mid x < x^*)) + (1-P(x^*))E(x > x^*)\right)}{(1-\delta(1-P(x^*))}\right) \qquad (8)$$

When utility of leisure is uniform on $[\underline{x}, \overline{x}]$, the optimal cutoff is always interior, and substituting into (8) shows that

$$x^*(1+\gamma) = \frac{(1-\delta)V(\overline{x}-\underline{x})}{\left(\overline{x}-\underline{x}-\delta(\overline{x}-x^*)\right)} - .5\delta\frac{-\gamma(x^{*2}-\underline{x}^2)+(\overline{x}^2-x^{*2})}{\left(\overline{x}-\underline{x}-\delta(\overline{x}-x^*)\right)},$$

which can be simplified to

$$x^*(1+\gamma)(\overline{x}-\underline{x}-\delta\overline{x}) + .5\delta(1+\gamma)x^{*2} = (1-\delta)V(\overline{x}-\underline{x}) + .5\delta\left(\gamma\underline{x}^2-\overline{x}^2\right).$$

and explicitly solved to find

$$x^* = \frac{-(\overline{x}(1-\delta)-\underline{x}) + \sqrt{(\overline{x}(1-\delta)-\underline{x})^2 + 2\delta^2\left(v+.5\left(\gamma\underline{x}^2-\overline{x}^2\right)\right)/(1+\gamma)}}{\delta}$$

Computation shows that when $\overline{x}$ rises and $d\underline{x} = -d\overline{x}$, so that the mean is constant, the cut-off x^* falls and delay increases, provided that δ is large enough that $\overline{x}(1-\delta) < \underline{x}$. Thus the familiar conclusion that the option value of waiting is increased by a mean-preserving spread extends to the dual-self model, at least when δ is not too small.[14]

Several other features of this solution deserve note. For simplicity focus on the case where $\underline{x} = 0, \overline{x} = 1$. First the cutoff is increasing in V and decreasing in γ. That is, the greater the reward relative to the cost, and the smaller the cost of self-control, the less the delay. Note also that even if $\gamma = 0$ there is delay, due to the familiar "option value" of waiting. The effect of γ is easiest to see in the limit case where $\delta = 1$. Here if $\gamma = 0$ and $\delta = 1$ we find that $x^* = \sqrt{2v-.5}$ provided that[15] $v > .5$. In this case the expected waiting time until the action is taken is

$$\frac{1}{1-x^*} = 1 + \frac{\sqrt{2v-.5}}{1-\sqrt{2v-.5}}.$$

---

[14] For small values of δ we hit the corner solution $x^* = \underline{x}$; which again falls with mean-preserving spreads; we believe that the monotonicity also holds for intermediate values.

[15] When $v < .5$ the project is not worth doing, as its flow return is less than the expected value of waiting.

As γ increases, x^* falls, and expected waiting time increases to

$$1+\frac{\sqrt{2v-.5}}{\sqrt{1+\gamma}-\sqrt{2v-.5}};$$

the difference in waiting times of

$$\frac{\sqrt{2v-.5}(\sqrt{1+\gamma}-1)}{(1-\sqrt{2v-.5})(1-\sqrt{(2v-.5)/(1+\gamma)})}$$

is the procrastination due to costly self control.

O'Donoghue and Rabin [2001] say that the agent "procrastinates" if he never acts even though there is an action that is worth doing *given* his hyperbolic discounting of future returns. They also show that their model typically has multiple equilibria, unlike ours. The equilibria are cyclic, with intervals of length T between "action dates." We view these cyclic equilibria as artificial and unappealing. Moreover, because they rely on long chains of backwards induction, the results of Fudenberg, Kreps and Levine [1988] show that they are not robust to even small payoff uncertainty. Despite the presence of multiple equilibria, O'Donoghue and Rabin can show that sophisticated agents (that is those who know their own hyperbolic parameter β) never procrastinate, although they may postpone acting for a few periods.

DellaVigna and Malmendier [2003] report some calibrations of the O'Donoghue-Rabin model to data on delay in canceling health club memberships, which they attribute to a combination of hyperbolic preferences and "lack of sophistication," meaning that consumers misperceive their own hyperbolic parameter and thus incorrectly forecast their health club usage.[16] Our model suggests several qualifications to their analysis. First of all, as is standard in models of timing, it is not in general optimal for the agent to act whenever he is indifferent between acting now or not at all, as there is an "option value"

[16] Sophisticated, "low β" agents who have correct perceptions about the costs and benefits of the club would correctly forecast that they would rarely attend but take a long time to cancel, while agents who misperceive their β would expect to exercise a lot. DellaVigna and Malmendier also show that agents choose monthly or annual plans with no per-visit charge when it would be cheaper to pay per visit. The use of prepayment as a commitment device is a consequence of both the hyperbolic and dual-self models.

in waiting.[17] Second, while there is some evidence that agents do not have perfect knowledge about themselves, we expect them to have more information about things that they have had more chances to observe.[18] Thus it seems natural to assume that the misperceptions about the short-run disutility and long-run benefits of going to the health club are larger than misperceptions about their own impulsiveness, and of course these former misperceptions can also explain the excessive delay.

We can also compare the dual-self model to the deterministic, finite horizon model of O'Donoghue and Rabin [1999], which allows non-stationary costs. In their Example 1, there are 4 periods to do a report, with costs 3, 5, 8, 13, and value v. Applying our model to this problem, we see that in the final period there is no self-control problem, so the payoff to acting is $v - 13$. In the next to last period the short-run self gets 0 from waiting, -8 from acting, so the utility if doing now is $v - 8(1+\gamma)$; in the previous periods the payoffs are $v - 5(1+\gamma), v - 3(1+\gamma)$ respectively. Thus the solution in our model is: if $\gamma < 10/3$ do it at the start, else do it at the end, and never do in intermediate periods. This is different than their solution, even for the case of a sophisticated agent, as the sophisticated agent acts in the second and in the fourth periods. Basically this equilibrium corresponds to one of the unappealing cyclic equilibria in their infinite horizon model; with an odd number of periods, the equilibrium is for the sophisticated agent to act in the first, third and fifth periods.[19]

---

[17] This factor is also present in the O'Donoghue-Rabin model, but the discussion of cancellation in DellaVigna and Malmendier [2003] seems to use a deterministic specification for the costs of cancellation. Also, the calibration measures cancellation lag by the number of full months between the last attendance and contract termination for users who hold a monthly contract at the time of termination. This is a conservative estimate *if* the customer knows that she will not attend in the future by the end of the month that included the customer's last visit, but otherwise may exaggerate the amount of delay.

[18] Bodner and Prelec [2003] and Benabou and Tirole [2004] build on the idea of imperfect self-knowledge to develop models of "self-signalling' that they use to explain the use of "personal rules." These models assume that the agent is uncertain of only one thing. In Benabou and Tirole [2004], for example, the agent knows the distribution of costs but does not know his hyperbolic parameter β. Both Bodner and Prelec [2003] and Benabou and Tirole [2004] analyze Bayesian equilibria of their models, which raises the question of whether a plausible non-equilibrium learning process would lead agents to learn the strategy of their "other selves" without learning the underlying value of β. Dekel et al [2004] analyze this issue in games between multiple agents; their results show that when β if fixed over time, assumptions that allow the agent not to learn β are typically too weak to justify restricting attention to equilibrium.

[19] The multiple-selves version of the delay game is continuous at infinity, so from Fudenberg and Levine [1983] every limit of finite-horizon subgame-perfect equilibria is a subgame-perfect equilibrium in the infinite horizon.

6. Cognitive Load and Self Control

Shiv and Fedorikhin [1999] report on the following experiment. Subjects were asked to memorize either a two- or a seven-digit number, and then walk to a table with a choice of two deserts, namely chocolate cake and fruit salad. Subjects would then pick a ticket for one of the deserts, and go to report both the number and their choice in a second room. In one treatment, actual samples of the deserts were on the table, and in a second treatment, the deserts were represented by photographs. The authors' hypothesize that subjects will face a self-control problem with respect to the cake, in the sense that it will have higher emotional appeal but be less desirable from the "cognitive" viewpoint; that the subject's reaction is more likely to be determined by the emotional ("affective") reaction when cognitive resources are constrained by the need to remember the longer number, and that this effect will be greater when faced with the actual deserts than with their pictures.[20] The experimental results confirm these predictions. Specifically, when faced with the real deserts, subjects who were asked to remember the seven-digit number chose cake 63% of the time, while subjects given the easier two-digit number chose cake 41% of the time, and this difference was statistically significant. In contrast, when faced with the pictures of the deserts, the choices were 45% and 42% respectively, and the difference was not significant.

The finding that the increasing the cognitive load increases responsiveness to temptation can be easily captured in our model by the assumption that the cost of self-control is higher when the long-run self – which we identify with cognitive processing – has other demands on its resources. Here are two ways of formalizing this. First, we can assume that the cognitive center has a fixed amount D of cognitive resources, and that the resources required for self-control are proportional to the short-run foregone utility. Then when other cognitive tasks consume d resources, the utility of the short-run self can be reduced by at most $D-d$; greater self-control is simply infeasible. Notice that this formulation is consistent with our earlier reduced form, provided that the resource constraint does not bind; in this case it is consistent also with the Gul-Pesendorfer axioms.

[20] They base this last hypothesis on the work of Lowenstein [1996].

Alternatively, we can suppose that the cognitive center does not face a fixed resource constraint, but instead has increasing marginal cost. Let d be the resources required for cognitive tasks such as short-term memory, let s be the resources required for self control, and let $g(d+s)$ be the disutility of the resources used in the two tasks, with the resources required for self-control proportional to the foregone utility, and thus equal to γg. This reduces to our base model if $g(d) = \gamma d$, but we now want to assume that cost is convex, that is, $g' > 0$ and $g'' > 0$.[21]

Let u^c and u^f be the short-run utilities of chocolate and fruit respectively, with $u^c > u^f$. Then $s(fruit) = u^c - u^f$, and the overall utility functions are $U(chocolate) = u^c - g(d)$ and $U(fruit) = u^f - g(d + u^c - u^f)$. This implies that increases in d make chocolate comparatively more attractive.

In the hyperbolic model, there is a sequence of far-sighted individuals with time-inconsistent preferences parameterized by β, so the most obvious way to explain the Shiv and Fedorikhin result is to assume that β decreases when the cognitive center has other tasks. This is roughly analogous to our proposal, but to us it seems more natural and direct to assume that cognitive load uses up self-control resources.

The effect of substituting photographs for the actual deserts shows the importance of cues and framing: The evidence supports Lowenstein's theory that vividness influences the effect of temptation. This raises the conjecture that the use of "rules of thumb" like "only have sweets at dinner" aids self-control buy reducing the vividness of the temptation; confirming this would require comparing brain scans of agents who use these rules to other agents who do not.

7. Conclusion and Discussion

Our resolution of the Rabin paradox shows how the dual self model captures "framing" effects. In our model it makes a difference whether an unanticipated payment is received on the floor of a casino or the lobby of a bank. Cues are obviously the key to understanding framing. The dual self theory implies that it is the attention span of the short-run self that is relevant for determining what constitutes a "situation" – the most difficult modeling issue in confronting these types of issues. This suggests that one might

[21] The model with convex costs does not satisfy assumption 5, so it does not satisfy the Gul-Pesendorfer axioms. It would be interesting to know how to relax those axioms to include this specification.

be able to use experimental and physiological data to determine what the relevant frames are. The dual self theory would then enable us to paste information about the motivation of the myopic self into the broader context in which real decision making takes place.

Although stimulus-response models seem to be a good framework for understanding how cues arise, the standard forms of these models seem to be a poor fit for many aspects of human cognition.[22] For example, faced with no observations people will respond differently depending on their prior knowledge of a situation. More strikingly, they have been seen to learn by "figuring things out" without any external stimulus at all. In future research, we hope to use the dual-self model to reconcile these two theories, while at the same time providing a theory of how cues arise.

This paper has focused on the case of "sophisticated" agents who are aware of their own self-control costs. Many papers on hyperbolic utility also consider the case of "naïve" agents, who have a current self-control problem but incorrectly forecast that they will not have such problems in the future. One could extend our model to allow for this by allowing agents to misperceive the future value of γ; this remains a topic for future research, as does enriching the dual-self model to allow for phenomena such as habit formation and various sources of systematically biased forecasts. All of these would add descriptive realism in some settings.

---

[22] By the "standard model" here we mean that reinforcements are applied directly to actions. Stimulus-response dynamics can be defined on much larger spaces of sequences of actions, hypothetical reinforcements, etc.; at that level of generality they encompass a much larger set of phenomena.

References

Benabou, R. and J. Tirole [2004]: "Willpower and Personal Rules, *Journal of Political Economy,* forthcoming.

Bodner, R. and D. Prelec [2003]: "The Diagnostic Value of Actions in a Self-Signaling Model, in I. Brocas and J. Carillo, (eds.), *The Psychology of Economic Decisions, Vol. 1*, Oxford University Press, Oxford.

Camerer, C., G. Lowenstein, and D. Prelec [2000]: "Why Economists Need Brains," http://www.hss.clatech.edu/~camerer/scanjecon11final.doc

Dekel, E., D. Fudenberg, and D. K. Levine [2004]: "Learning to Play Bayesian Games," *Games and Economic Behavior*, **46**: 282-303.

DellaVigna, S. and U. Malmendier, [2003]: "Overestimating Self-Control: Evidence From the Health Club Industry," Stanford GSB Research Paper 1800.

DellaVigna, S. and U. Malmendier [2003]: "Contract Design and Self Control: Theory and Evidence," http://emlab.berkeley.edu/users/sdellavi/wp/self_control_jan04.pdf.

Fudenberg, D., D. Kreps and D. K. Levine [1988]: "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, **44**: 354-380

Fudenberg, D. and D. K. Levine [1983]: "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games," *Journal of Economic Theory*, **31**: 251-258.

Fudenberg, D. and D. K. Levine [1998]: *Theory of Learning in Games*, MIT Press, Cambridge MA.

Gul, F. and W. Pesendorfer [2001]: "Temptation and Self Control," *Econometrica*, **69**: 1403-1436.

Gul, F. and W. Pesendorfer [2002]: "Self Control and the Theory of Consumption," forthcoming, *Econometrica*.

Gul, F. and W. Pesendorfer [2004]: "Self Control, Revealed Preference and Consumption Choice," forthcoming, *Review of Economic Dynamics*.

Harris, C. and D. Laibson [2004]: "Instantaneous Gratification," mimeo.

Haruno, M. , T. Kuroda, K. Doya, K. Toyama, M. Kimura, K. Samejima, H. Imamizu, and M. Kawato [2004]: "A Neural Correlate of Reward-Based Behavioral Learning in Caudate Nucleus: A Functional Magnetic Resonance Imaging Study of a Stochastic Decision Task," *The Journal of Neuroscience*, **24**:1660-1665.

Krusell, P. and T. Smith. [2000]: "Consumption-Savings Decisions with Quasi-Geometric Discounting," *Econometrica*.

Laibson, D. [1997]. "Golden eggs and hyperbolic discounting." *Quarterly Journal of Economics, 112*, 443-477.

Lowenstein. G. [1996] "Out of Control: Visceral Influences on Behavior," *Organizational Behavior and Human Decision Process,* 65,272-292.

McItnosh, D. [1969] *The Foundations of Human Society,* U. Chicago Press, Chicago

McClure, S.M., Laibson, D.I., Loewenstein, G., Cohen, J.D. [2004]: "Separate neural systems value immediate and delayed monetary rewards," *Science*, forthcoming

O'Donoghue, T. and M. Rabin [1999]: "Doing It Now or Later," *American Economic Review* **89**, 103-124.

O'Donoghue, T. and M. Rabin [2001]: "Choice and Procrastination" *Quarterly Journal of Economics*, 121-160.

Prelec, D. [2004] "Decreasing impatience:A criterion for non-stationary time preference and 'hyperbolic' discounting," mimeo.

Rabin, M., [2000]: "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*, 1281-1292.

Shiv, B. and A. Fedorikhin [1999]: "Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making," *Journal of Consumer Research*, 26, 72-89.

Strotz, R. H. [1955]. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies, 23*, 165-180.

Thaler, R. and H.M. Shefrin [1981] "An Economic Theory of Self-Control," *Journal of Political Economy,* 89, 392-406.

Wertenbroch, K., D. Soman, and J. Nunes [2002]: "Debt Aversion and Self-Control: Consumer Self-Management of Liquidity Constraints." Mimeo.

Appendix

I. We first give a result for general per-period utility functions in the simple savings model. Consider the problem of maximizing

$$U(\vec{a}) = \sum_{t=1}^{\infty} \delta^{t-1} \left[(1+\gamma)u((1-a_t)y_t) - \gamma u(y_t)\right]$$

over all feasible plans $\vec{a}$, i.e. plans that satisfy $a_t \in [0,1]$ and the wealth equation $y_t = R a_{t-1} y_{t-1}$. We suppose that u is non-decreasing and continuous on $(0,\infty)$; we do not require continuity on $[0,\infty)$ because we want to allow for the logarithmic case where $u(0) = \lim_{c \to 0} u(c) = -\infty$. Let $\bar{U}$ be the supremum in this problem.

Proposition A.1: *Suppose* $R > 1$

(A.1) $$\sum_{t=1}^{\infty} \delta^{t-1} u(y_0 R^t) < \infty .$$

Then: (i) For any feasible plan the sum defining *U* has a well defined value in the sense that either the sum converges absolutely or converges to $-\infty$.

(ii) The supremum $\bar{U}$ of the feasible values satisfies $-\infty < \bar{U} < \infty$

(iii) If feasible $\vec{a}^n \to \vec{a}^*$ in the product topology then $\vec{a}^*$ is feasible. If in addition $U(\vec{a}^n) \to \bar{U}$ then $U(\vec{a}^*) = \bar{U}$.

(iv) An optimal plan exists. That is, there is a feasible plan that attains $\bar{U}$.

Proof: (i) For any sequence $(\vec{a}, \vec{y})$ with $y_t = R a_{t-1} y_{t-1}$, let

$$\chi_+(x) = \begin{cases} x & x > 0 \\ 0 & x \le 0 \end{cases}$$

and $\chi_-(x) = -\chi_+(-x)$. We can write any finite sum as the sum of negative and positive parts.

$$\sum_{t=1}^{T} \delta^{t-1} \left[(1+\gamma) u((1-a_t) y_t) - \gamma u(y_t) \right] =$$
$$\sum_{t=1}^{T} \delta^{t-1} \chi_{+} \left[(1+\gamma) u((1-a_t) y_t) - \gamma u(y_t) \right] +$$
$$\sum_{t=1}^{T} \delta^{t-1} \chi_{-} \left[(1+\gamma) u((1-a_t) y_t) - \gamma u(y_t) \right]$$

The positive part of the sum is summable from (A.1), since

$$(1+\gamma) u((1-a_t) y_t) - \gamma u(y_t) \leq u((1-a_t) y_t) \leq u(y_0 R^t).$$

The negative part is monotone decreasing in *T,* so it either converges absolutely or converges to $-\infty$. In the former case the entire sum converges absolutely; in the latter case the sum converges to $-\infty$.

(ii) Part (i) already shows that $\bar{U} < \infty$. To see that $\bar{U} > -\infty$, note that it is feasible to set $a_t = 1/R$ for all *t*, and that for $R > 1$ this plan yields a finite value.

(iii) Consider a sequence of feasible plans $\vec{a}^n \to \vec{a}^*$. Because the constraints are period by period and closed, it is clear that $\vec{a}^*$ satisfies the constraints, so it is feasible. Now suppose in addition that $U(\vec{a}^n) \to \bar{U}$. Choose any $\varepsilon > 0$ and pick n large enough that $\bar{U} - U(\vec{a}^n) < \varepsilon / 2$. If we now pick τ such that

$$\sum_{t=\tau+1}^{\infty} \delta^{t-1} \left| u(R^t y_0) \right| < \varepsilon / 2,$$

we know that

$$\bar{U} - \sum_{t=1}^{\tau} \delta^{t-1} \left[(1+\gamma) u((1-a_t^n) y_t^n) - \gamma u(y_t^n) \right] \leq \varepsilon.$$

Since $\bar{U}$ is finite, and payoffs in the first τ periods are bounded above by $u(y_0 R^t)$, each term in this summation is bounded below (by $\bar{U} - \varepsilon - T u(y_0 R^t)$). Since per-period payoffs are continuous at any (a, y) with $a > 0$, $\vec{a}^n \to \vec{a}^*$, and $y_t^n \to y_t^*$, it follows that

$$\bar{U} - \sum_{t=1}^{\tau} \delta^{t-1} \left[(1+\gamma) u((1-a_t^*) y_t^*) - \gamma u(y_t^*) \right] \leq \varepsilon.$$

Since this is true for any $\varepsilon > 0$ and we know that $U(\vec{a}^*) \leq \bar{U}$, we conclude that $U(a^*) = \bar{U}$.

(iv) Now consider a feasible sequence $(\vec{a}^n, \vec{y}^n)$ with $U(\vec{a}^n) \to \bar{U}$. Each savings rate a_t must lie in the compact interval $[0,1]$ and each y_t must lie in the compact interval $[0, R^t y_0]$, so the sequence $(\vec{a}^n, \vec{y}^n)$ has an accumulation point $(\vec{a}^*, \vec{y}^*)$ in the product topology. This accumulation point is a maximum by part (iii).

☑

II. Now we specialize to the CRRA utility functions

$$u(c) = \frac{(c)^{1-\rho} - 1}{1-\rho}$$

and $u(c) = \ln(c)$, which corresponds to the case $\rho = 1$. Assuming $\delta < R^{\rho-1}$ implies

$$\sum_{t=1}^{\infty} \delta^{t-1} u(y_0 R^t) < \infty .$$

It follows from Proposition A.1 that an optimum $\vec{a}^*$ exists.

Proposition A.2: *With CRRA utility a stationary optimum with* $a_t = a$ *exists.*

Proof: Suppose that $\vec{a}^*$ is an optimal plan. By homogeneity of the objective function, and the fact that plans are defined in terms of savings rates, $\vec{a}^*$ is also an optimal plan starting in period 2 (for any initial condition). Note that the plan $\vec{a}^2 = (a_1^*, a_1^*, a_2^*, a_3^*, \ldots)$ yields wealth in period 2 of $a_1^* R y_1$, and let $\bar{U}(y_1)$ denote the maximized utility when starting in the second period with wealth y_1. Then

$$U(\vec{a}^2) = (1+\gamma)u((1-a_1^*)y_0) - \gamma u(y_0) + \delta \bar{U}(a_1^* R y_0) = \bar{U}$$

where the first equality follows because $\vec{a}^*$ is optimal from period 2 on, and the second equality because $\vec{a}^*$ is optimal from the first period. Proceeding in this way we can construct sequence of feasible plans $\vec{a}^n = (a_1^*, a_1^*, \ldots, a_1^*, a_2^*, a_3^*, \ldots)$ that play a_1^* for the first n periods such that $U(\vec{a}^n) = U(a^*) = \bar{U}$. Clearly $\vec{a}^n$ converges in the product topology to the plan of choosing the fixed savings rate a_1^*. Hence it follow from Proposition A.1 (iii) that this limiting plan is feasible and gives utility $\bar{U}$; that is, it is optimal.

☑

III. In the CES case, then, it is sufficient to compute the present value utility from a fixed savings rate a, and maximize over savings rates. We have present value utility

$$U = \frac{y_0^{1-\rho}}{1-\rho} \sum_{t=1}^{\infty} (\delta (Ra)^{1-\rho})^{t-1} \left[(1+\gamma)(1-a)^{1-\rho} - \gamma \right] - \frac{1}{(1-\delta)(1-\rho)}$$
$$= \frac{y_0^{1-\rho}}{1-\rho} \frac{\left[(1+\gamma)(1-a)^{1-\rho} - \gamma \right]}{1-\delta (Ra)^{1-\rho}} - \frac{1}{(1-\delta)(1-\rho)}$$

Since the optimal savings rate cannot be 0 or 1, we may differentiate with respect to the saving rate to find

$$dU / da =$$
$$y_0^{1-\rho} (1-a)^{-\rho} \frac{[(1+\gamma)(1-a) - \gamma (1-a)^{\rho}] \delta R (Ra)^{-\rho} - (1+\gamma) \left(1 - \delta (Ra)^{1-\rho} \right)}{\left(1 - \delta (Ra)^{1-\rho} \right)^2}$$

which gives necessary condition for an optimum[23]

$$(1+\gamma) a^{\rho} = R^{1-\rho} \delta ((1+\gamma) - \gamma (1-a)^{\rho}).$$

When $\gamma = 0$ we get the usual solution $a^* = R^{(1-\rho)/\rho} \delta^{1/\rho}$. Thus we can rewrite the first order condition as

$$(a / a^*)^{\rho} = ((1+\gamma) - \gamma (1-a)^{\rho}) / (1+\gamma).$$

Turning to the simple banking model, utility starting in the second period is the $\gamma = 0$ solution

[23] We do not know if the first-order condition has a unique solution, except in the logarithmic case.

$$U_2(y_2) = \frac{y_2^{1-\rho}}{1-\rho}\frac{(1-a^*)^{1-\rho}}{1-\delta(Ra^*)^{1-\rho}} - \frac{1}{(1-\delta)(1-\rho)}$$
$$= \frac{y_2^{1-\rho}}{1-\rho}\frac{1}{\left(1-a^*\right)^{\rho}} - \frac{1}{(1-\delta)(1-\rho)}$$
$$= \frac{y_2^{1-\rho}}{1-\rho}\frac{1}{\left(1-\delta^{1/\rho}R^{(1-\rho)/\rho}\right)^{\rho}} - \frac{1}{(1-\delta)(1-\rho)}$$

The utility of both selves in the first period is

$$(1+\gamma)\frac{(c_1)^{1-\rho}-1}{1-\rho} - \gamma\frac{(x_1+z_1)^{1-\rho}-1}{1-\rho},$$

and so the overall objective of the long-run self is to maximize

$$(1+\gamma)\frac{(c_1)^{1-\rho}-1}{1-\rho} - \gamma\frac{(x_1+z_1)^{1-\rho}-1}{1-\rho}$$
$$+\frac{(R(y_1+z_1-c_1))^{1-\rho}}{1-\rho}\frac{\delta}{\left(1-\delta^{1/\rho}R^{(1-\rho)/\rho}\right)^{\rho}} - \frac{\delta}{(1-\delta)(1-\rho)}$$

The first order condition for optimal consumption is

$$\frac{c_1}{R(y_1+z_1-c_1)} = \frac{(1+\gamma)^{1/\rho}\left(1-\delta^{1/\rho}R^{(1-\rho)/\rho}\right)}{(\delta R)^{1/\rho}}.$$

If there are one or more solutions that satisfy the constraint $c_1 \le x_1 + z_1$ then one of them represents the optimum; otherwise the optimum is to consume all pocket cash, $c_1 = x_1 + z_1$.

Note that x_1 is the solution for $\gamma = 0$, so it satisfies

$$\frac{x_1}{R(y_1-x_1)} = \frac{\left(1-\delta^{1/\rho}R^{(1-\rho)/\rho}\right)}{R^{1/\rho}}.$$

Thus we can write the first order condition as

$$\frac{c_1}{y_1+z_1-c_1} = (1+\gamma)^{1/\rho}\frac{x_1}{y_1-x_1}$$

or

$$
\begin{aligned}
c_1 &= \frac{(1+\gamma)^{1/\rho} x_1}{y_1 - x_1 + (1+\gamma)^{1/\rho} x_1}(y_1 + z_1) \\
&= \frac{1 + \left[(1+\gamma)^{1/\rho} - 1\right]}{1 + \left[(1+\gamma)^{1/\rho} - 1\right](1 - a^*)}(1 - a^*)(y_1 + z_1) \\
&= B(1 - a^*)(y_1 + z_1)
\end{aligned}
$$