



A Dual-Self Model of Impulse Control

Citation

Fudenberg, Drew, and David K. Levine. 2006. A dual-self model of impulse control. *American Economic Review* 96(5): 1449-1476.

Published Version

<http://dx.doi.org/10.1257/aer.96.5.1449>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3196335>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

A Dual Self Model of Impulse Control¹

Drew Fudenberg and David K. Levine²

First Version: August 5, 2004

This Version: April 5, 2006

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 quasi-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 quasi-hyperbolic utility, and it explains experimental evidence that increased cognitive load makes temptations harder to resist. The base version of our model is consistent with the Gul-Pesendorfer axioms, but we argue that these axioms must be relaxed to account for the effect of cognitive load.

“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])

This paper argues that a simple “dual-self” model gives a unified explanation for a number of empirical regularities related to self-control problems and a value for commitment in decision problems. One of these regularities is the apparent time inconsistency that has motivated economists’ models of quasi-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 A 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 provides a possible explanation of 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 suggests that short-term impulsive behavior is associated with different areas of the brain than long-term planned behavior.⁴ We argue that our dual-selves model explains a broad range of behavioral anomalies, and that it is a better (though still crude⁵) fit for the modular structure of the brain than the quasi-hyperbolic model, which posits a game between multiple “selves,” one in each period. Moreover, the dual-selves model is analytically simpler than the quasi-hyperbolic discounting model, as the equilibria of the model can be calculated as the solution to a decision problem. In addition, in standard economic applications where the quasi-hyperbolic model has multiple equilibria, the dual-self equilibrium is unique. While multiple equilibria seem a good way of describing some situations, the unique equilibria of the dual-self model are an advantage to the extent that it can explain the empirical facts just as well as models with less precise predictions.⁶

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. Specifically, we imagine that the short-run myopic self has “baseline preferences” in the stage game that depends only on the outcome in the current stage. That is, the short-run players are completely myopic.⁷

The stage game is played in two phases. In the first phase, the long-run self chooses a self-control action that influences the utility function of the myopic self. That is, at some reduction in utility (for both selves) the long-run self can choose preferences other than the baseline preferences. 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. It is important that we do not allow the long-run self to precommit for the entire dynamic game. Instead, she begins each stage game facing the choice of which preferences to give the myopic self, or equivalently, how much self-control to exert. Note also that while the quasi-hyperbolic discounting model emphasizes the conflict between present and future selves, we emphasize 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. Indeed, while we find the language of multiple “selves” to be suggestive, the model can equally well be interpreted as describing the behavior of a single “self” whose overall behavior is determined by the interaction of two subsystems.⁸

Games with long-run versus short-run players are relatively simple to analyze, and 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. For this reason the equilibria of the 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 quasi-hyperbolic discounting, preserving many of the methods and insights of existing theory, as well as delivering strong predictions about behavior.

The dual self model predicts a preference for commitment, just as the quasi-hyperbolic model does. 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 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 may incur 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. Finally, in some cases the long-run self may be willing to incur short run costs to reduce the future cost of self control.

As an 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, and so winnings from sufficiently small cash gambles are evaluated by the short-run self's preferences, which are over consumption. These preferences are more risk-averse than the preferences over long-term consumption that are used to evaluate large gambles, so this "cash effect" provides an explanation of Rabin's [2000] paradox of risk aversion in the small and in the large.⁹

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.

Our final application of the model is to the effect of cognitive load on self-control, as shown in the experiments of Shiv and Fedorikhin [1999] and Ward and Mann [2000]. We use this evidence to motivate the assumption that the cost of self-control is convex, as opposed to

linear. We then show that agents with non-linear cost of self-control can violate the axiom “independence of irrelevant alternatives” when choosing from menus, as well as the version of the independence axiom used by Gul and Pesendorfer [2001].

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 2, 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. Independent work by O'Donoghue and Lowenstein [2004] describes a similar but more general model, with less focus on tractability and applications.¹⁰ Bernheim and Rangel [2004], and Benahib and Bisin [2004] consider multi-period models where a long-run self is only sometimes in control, either because it is unable to take control (Bernheim and Rangel) or chooses not to do so (Benahib and Bisin). We discuss these papers further in the conclusion. Section 5 discusses Miao [2005], who applies the dual-self model to a variant of the waiting-time problem we analyze in that section.

Although our point of departure is different, the reduced form of the dual self model is closely connected to the representations derived in Gul and Pesendorfer [2001,2004], and Dekel, Lipman, and Rustichini [2005]. These papers consider a single player who has preferences over choice sets that include the desire to limit the available alternatives. Under various axioms over choices over menus of lotteries, they show that the decision process can be represented by a

utility function with a cost of self-control.¹¹ The reduced form of our model leads to a similar decision problem; our interpretation of preferences over menus as arising from a conflict between two selves or systems lets us bring both introspective and physiological evidence to bear on what those preferences might be. This leads us to a model that is more restrictive in some ways, and less restrictive in others; we discuss the relationship in detail in Section 6. Krusell, Kuruscu and Smith [2005] examine a variation on the infinite-horizon Gul-Pesendorfer model in the setting of a consumption-savings problem. Within this setting, they consider a wider range of preferences; we say more about this in Section 2.

1. The Model

Time is discrete and potentially unbounded, $t = 1, 2, \dots$. There is a fixed, time-and history invariant set of actions A for the short-run selves; this is assumed to be a closed subset of a finite-dimensional Euclidean space.¹² We let \mathbf{A} denote the set of all probability measures on A .¹³ To encode the effects of past history on current and future payoff possibilities we use a second closed subset of a Euclidean space; we denote this space by Y and let \mathbf{Y} denote the probability measures on Y . (As an example, in a consumption-savings application the state will correspond to wealth, and the action will be the savings rate.) Finally there is a set R of self-control actions for the long-run self; R is a closed subset of Euclidean space, and \mathbf{R} are the probability measures on R . 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, $h = (y_1, r_1, a_1, \dots, y_t, r_t, a_t)$ or the null history 0 . The initial state is y_1 , the final state in h is $y(h)$, and the length of the history is denoted by $t(h)$. The probability distribution over states at time $t + 1$ depends on the time- t state and action y_t, a_t according to the exogenous probability measure $\mu(y_t, a_t)$, so that

$\mu(y_t, a_t)[Y']$ denotes the probability of the set Y' at time $t + 1$. We make the technical assumption that for fixed measurable $Y' \subseteq \mathbf{Y}$, $\mu(\cdot, \cdot)[Y']$ is a measurable function of y and a .

Thus the short-run self is both the short-run utility assessor and the “doer”; all interactions with the outside world are handled by the short-run self. The long-run self’s action r has no direct effect on the future state and serves only to influence the actions of the short-run player through its effect on the short-run player’s payoff function $u(y, r, a)$; ¹⁴ this rules out the possibility that the state directly encodes the agent’s past use of self-control, as in a model where the agent can lower the cost of self control by repeatedly exercising it. The model is agnostic about the exact form of the “influences actions” r ; the model’s observable implications concern the actions a .

We now analyze the game whose players are the long-run self, 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. The mixed strategies of the long-run self are maps from histories and the current state to self-control actions, $\sigma_{LR} : H \times Y \rightarrow \mathbf{R}$. Denote by H_t the set of t -length histories. A strategy for the time- t short-run self is a map $\sigma_t : H_{t-1} \times Y \times R \rightarrow \mathbf{A}$; we denote the collection of all of these strategies by σ_{SR} . In both cases we impose the technical restriction that for every measurable subset $R' \subseteq \mathbf{R}, A' \subseteq \mathbf{A}$ the functions $\sigma_{LR}(\cdot, \cdot)[A'], \sigma_t(\cdot, \cdot)[R']$ are measurable. This means that the strategies together with the measure μ give rise to measures π_t over histories of length t for every t ; we suppress the dependence of these measures on the strategies to lighten notation.

The utility of the long-run self is thus given by

$$U_{LR}(\sigma_{LR}, \sigma_{SR}) = \sum_{t=1}^{\infty} \delta^{t-1} \int u(y_t(h), r_t(h), a_t(h)) d\pi_t(h).$$

In this formulation, the self-control cost (that is, the difference between $u(y, 0, a)$ and $u(y, r, a)$) 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.

At this point, we have not imposed sufficient assumptions to guarantee that discounted expected utility is well-defined, as the discounted sum defining U could fail to converge, and the expected value of utility at a given state might be infinite. The following assumption ensures that the objective function of the long-run self is well behaved.

Assumption 0 (Upper Bound on Utility Growth): *For all initial conditions*

$$\sum_{t=1}^{\infty} \delta^{t-1} \int \max\{0, u(h)\} d\pi_t(h) < \infty.$$

This requires that the largest values of utility not grow in expected value faster than the discount factor, which allows us to interchange the order of the summation and the integral.

Now that we have a game between the long-run self and sequence of short-run selves, we need to choose an appropriate concept of equilibrium. We refer to a strategy profile in which the short-run self optimizes following every history as *SR-perfect*. We shall only be interested in SR-perfect Nash equilibria, which are profiles such that the short-run players optimize following every history, and the long-run player anticipates how the short-run player will react and plans accordingly. Note that we do not impose subgame perfection on the long-run self, both because it is not necessary for our results, and because it is not as compelling and robust as the requirement of SR-perfection. The point is that each short-run self plays once, and plays a game in which expectations of long-run self play do not matter: each short-run self faces a simple static decision problem. So while subgame perfection with its long chains of backwards induction is

not always plausible, and is not robust to small changes in the information structure, the assumption that the short-run self optimizes in every subgame is both robust and compelling.

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

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

If we did not impose assumption 2, the self-control cost of some actions in some states could be infinite; this would bring the model closer to the addiction model of Bernheim and Rangel.

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.

Assumptions 1 and 3 imply that the cost function is continuous and has the following property:

Property 1: (Strict Cost of Self-Control) *If $a \in \arg \max_{a'} (u(y, 0, a'))$ then $C(y, a) = 0$, and $C(y, a) > 0$ for $a \notin \arg \max_{a'} (u(y, 0, a'))$.*

Conversely, given continuous functions $u(y, 0, a)$ and $C(y, a)$ satisfying Property 1, we can extend that utility function to a function $u(y, r, a)$ that generates C and satisfies Assumptions 1-

3. For example, if we set $R = A$, the function

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

has the requisite properties.

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

This is a joint restriction on the utility function u and the set of feasible self-control actions R ; it implies that when the short-run self is indifferent, the long-run self has an action that can break the tie for negligible cost.

We next consider a *reduced-form* optimization problem that has only a single player, and omits the self-control variable r . Let $H^{AY} = \{(y_1, a_1, \dots, y_t, a_t)\}_t$ be the set of *reduced* histories, consisting of states and actions only. Choosing a strategy from reduced histories and states to actions, $\sigma_{RF} : H^{AY} \times Y \rightarrow \mathbf{A}$ induces distributions π_t^{RF} over histories of length t for every t . We assume that that $\sigma_{RF}(\cdot, \cdot)[A^t]$ is measurable in y and a for all measurable A^t .¹⁵ The reduced form optimization problem is to maximize the objective function

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

Our goal is to establish equivalence between SR-perfect Nash equilibria of the game and solutions to the reduced-form optimization problem. To do so, we must first establish a link

between strategy profiles in the game and strategies in the optimization problem. A given strategy profile $(\sigma_{LR}, \sigma_{SR})$ gives rise to a stochastic process on the space of states, actions and self-control actions, characterized by the conditional probabilities. This in turn gives rise to a stochastic process on the space of states and actions only. That stochastic process has associated with it well-defined regular conditional probability measures,¹⁶ and these conditional probabilities define the strategies σ_{RF} . When such a construction is possible, we say that the profile $(\sigma_{LR}, \sigma_{SR})$ is equivalent to σ_{RF} ; we have just shown that every profile has at least one equivalent reduced form strategy.¹⁷

Theorem 1 (Equivalence of Subgame Perfection to the Reduced Form): *Under Assumptions 1-4, every SR-perfect Nash equilibrium profile is equivalent to a solution to the reduced form optimization problem and conversely.*

Proof: See the Appendix.

Remark 1: 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.

Remark 2: In the economic applications we consider, the solution to the reduced-form optimization problem is unique for generic parameter values. This is always true if the reduced form problem is strictly concave, but a problem that is concave in the absence of self-control costs can fail to be concave when self-control costs are significant.

While Assumptions 1-4 are sufficient for the equivalence result in Theorem 1, they are too general to be of much use in applications, so our next order of business is to specialize the model in a way that still lets it cover the intended applications. In particular, at this point we have not made any substantive assumptions about how the utility function depends on the state. Thus even though the self-control action r does not have a direct impact on the state, the action a played today can change future preferences (through the state variable y) in such a way that the cost of self-control is reduced in the future. Such situations are ruled out by the next Assumption, which says that the state influences the cost of self-control only through the current utility possibilities.

Assumption 5 (Opportunity Based Cost of Self Control): *If a is less satisfying for the short-run self in state y than in state y' in the sense that $u(y,0,a) \leq u(y',0,a)$, yet state y has greater temptation than state y' in the sense that $\max_{a'} u(y,0,a') \geq \max_{a'} u(y',0,a')$, then it is more costly to choose a in state y , that is, $C(y,a) \geq C(y',a)$.*

Because the assumption 5's conditions must hold for all combinations of states and actions, it implies that the cost of self-control depends only on the utility of the best foregone option $\max_{a'} u(y,0,a')$ and the utility of the option chosen $u(y,0,a)$. Thus assumption 5 rules out situations in which some states might make self-control more difficult without having any effect on the utility possibilities at $r=0$. We explain how to adapt the model to handle this in section 5, where we discuss how cognitive load influences self-control. Until then, however, we will maintain Assumption 5 and assume this possibility away. Note also that Assumption 5 rules out cases where the long-run self is uncertain which action the short-run self will find most tempting.

For this reason, the assumption is most plausible when there is a very short time interval between the choice of r and the choice of a .

Adding Assumption 5 to Assumptions 1-3 implies that there is a continuous function \tilde{C} such that $C(y, a) = \tilde{C}(u(y, 0, a), \max_{a'} u(y, 0, a'))$, where \tilde{C} is decreasing in its first argument (the realized utility), increasing in the second (the “temptation”), and $\tilde{C}(u, u) = 0$. Conversely, given continuous functions $u(y, 0, a)$ and \tilde{C} that have these properties, we can use the associated cost function C to generate a utility function that satisfies Assumptions 1-3 and 5. Moreover, we show in Section 6 that under these assumptions the induced preferences over menus satisfy the set-betweenness axiom of Gul and Pesendorfer.

In many of our applications, it is convenient to strengthen Assumption 5 by requiring that c is linear:

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

This assumption provides a tractable and 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, where both utilities are evaluated at $r = 0$. Under Assumption 5’, improving the best available alternative does not change the marginal cost of self-control. Moreover, under Assumption 5’ preferences satisfy the classical property *independence of irrelevant alternatives*. In this sense, these assumptions are conservative, maintaining as much of the standard model as is consistent with an interesting theory of self-control. Moreover, as we show in section 6, under Assumption 5’ our model is consistent with the Gul and Pesendorfer axioms, essentially because it implies the independence axiom of Dekel, Lipman and Rustichini [1999]. However we argue in Sections 5 and 6 that both Assumption 5’ and the independence

axiom may be too strong, as they rule out the idea that self-control is a limited resource. They are also not well motivated in settings where some uncertainty will be resolved after the action of the short-run self. To accommodate these aspects of behavior, we do not need to revert to the full generality of Assumption 5. Instead, we will use a functional form with one more parameter, namely

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

Because we have assumed that the long-run self cares only about the utility of the short-run self, Assumptions 1-5 are in some ways stronger than those of Gul and Pesendorfer. Their axioms allow the cost of self-control to depend on $\max_a v(y, a)$, where the utility function v need not be the utility function used to evaluate choices. At the same time, without the linearity assumption our assumptions are in some ways less restrictive than the axioms of Gul and Pesendorfer and Dekel, Lipman, and Rustichini. We discuss the connection between our model and these papers in some detail in Section 6.

2. A Simple Savings Model

To start, consider the simple case of an infinite-lived consumer making a savings decision. The state $y \in \mathfrak{R}_+$ 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 $y_{t+1} = Ra_t y_t$ next period, there is no other source of income.¹⁸

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

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

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

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

$$C(y, a) = \gamma (\log(y) - \log(1 - a)y)) = -\gamma \log(1 - a).$$

The reduced form for the long-run self has preferences

$$U_{RF} = \sum_{t=1}^{\infty} \delta^{t-1} [(1 + \gamma) \log((1 - a_t)y_t) - \gamma \log(y_t)].$$

The long-run self's problem is thus to maximize this function subject to the wealth equation $y_t = Ra_{t-1}y_{t-1}$.¹⁹ We study the technical details of the problem for the more general CRRA case in the Supplementary Appendix. There we show that there is a unique solution, and that the solution has a constant savings rate strictly between zero and one.²⁰ Thus we compute present value utility for constant savings rates, and maximize

$$\begin{aligned} & \sum_{t=1}^{\infty} \delta^{t-1} [(1 + \gamma) (\log(1 - a) + (t - 1) \log Ra + \log y_0) - \gamma(t - 1) \log(Ra) - \gamma \log(y_0)] \\ &= \frac{(1 + \gamma) (\log(1 - a) + \log(y_0))}{(1 - \delta)} + \frac{\delta \log(Ra)}{(1 - \delta)^2}. \end{aligned} \quad (1)$$

From the first-order conditions we can then compute that²¹

$$a = \frac{\delta}{1 + \gamma - \delta\gamma}. \quad (2)$$

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

$$\frac{[\log(1 - \delta) + \log(y_0)]}{(1 - \delta)} + \frac{\delta \log(R\delta)}{(1 - \delta)^2}; \quad (3)$$

we use this fact in the following section.

To summarize, the dual-self model has a constant savings rate for both logarithmic and CRRA utility, and the savings rate is the solution to a first-order condition; the solution is particularly simple in the case of logarithmic utility. In contrast, as Harris and Laibson [2004] emphasize, consumption need not be monotone in wealth in the usual discrete-time quasi-hyperbolic model, even in a stationary infinite-horizon environment. Moreover, the quasi-hyperbolic model typically has multiple equilibria (Krusell and Smith [2003]), 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.²² 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 the dual-self approach is more general and more direct.

3. Banking, Commitment, and Risk Aversion

In practice there are many ways of restraining the short-run self besides the use of self-control: the obvious thing to do is to make sure that the short-run self does not have access to resources that would represent a temptation. This leads to a value for commitment, and can also explain some of the “preference reversals” discussed in the literature. We will shortly develop a more complex model where commitment occurs via a cash-in-advance constraint, but as a first illustration consider a three-period model where saving is impossible, and baseline consumption is constant. Suppose that at period 1 the consumer is offered a choice between two possible increases over base-line consumption, either an increase of 1 unit in period 2 or 2 units in period 3. Since neither option increases short-run consumption, the short-run self is indifferent, and the consumer will choose as if she were fully rational, that is, as if her self-control cost were equal to 0. On the other hand, offered the choice between a unit now and 2 units in period 2, the consumer does face a self-control cost in choosing the delayed payoff. Hence for a range of parameter values the consumer would choose the larger, later increment in the first decision and the smaller, sooner one in the second.

Before developing the banking model, we need to address a key unresolved issue in the behavioral literature, namely the correct way to model how agents view money rewards. This is important, both for our paper, and for the interpretation of empirical studies of preference

reversal in humans that examine not consumption choices but monetary payoffs.²³ Unless current consumption is liquidity-constrained, this evidence raises a puzzle for both the dual-self and quasi-hyperbolic models: since people cannot literally consume currency, why do they act as if current monetary rewards are tempting? In brief it seems that the short-run self treats money as a cue for an immediate reward even though the only real consequence of earning money is in the future,²⁴ but that observation leaves open the questions of exactly which financial rewards we should expect agents to view as tempting, and what other sorts of deferred rewards will be treated in the same way. The conclusion speculates about some possible extensions of our model to learned cues, but explaining when and why money is viewed as consumption is beyond the scope of this paper. Instead, we will develop a model where agents get utility only from consumption, and so are subject to self-control costs only when consumption is possible.

In particular, we develop a simple model in which basic savings decisions are made in a 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.²⁵ 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. Thus the baseline, deterministic version of the model has an equilibrium equivalent to a model without a self-control problem. However, the consumer’s response to unanticipated cash receipts is quite different than that to anticipated receipts, or to 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. When the stakes are large, self-restraint kicks in, part of the winnings will be saved and spread over the lifetime. This leads to less risk-averse preferences, so the model explains the paradox proposed by Rabin [2000].

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, which is 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 $y_{t+1} = R(s_t + x_t - c_t)$. The discount factor between the two consecutive

nightclub periods (which is where consumption occurs) is δ ; preferences continue to have the logarithmic form.²⁶ First, 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 $a^* = \delta$, the optimum of the problem without self-control, simply by choosing pocket cash $x_t = (1 - a^*)y_t$ to be the desired consumption. The short-run self will then spend all the pocket cash; because the optimum can be obtained without incurring self-control costs, the long-run self does not in fact wish to exert self-control at the nightclub.

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. If the lotteries 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^* . To avoid the small correction terms that would be generated when the lottery has a small but non-zero probability, we will consider the limit where the probability of drawing the gamble is exactly zero; this simplifies the calculations but is not important for the qualitative features of our conclusions.

For the agent to evaluate a lottery choice \tilde{z}_1 , he needs to consider how he would behave conditional on each of its possible realizations 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

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

The first order condition for optimal consumption in the nightclub can be solved to find

$$\begin{aligned} c_1^* &= \frac{(1-\delta)(1+\gamma)(y_1 + z_1)}{\delta + (1+\gamma)(1-\delta)} \\ &= \left(1 - \frac{\delta}{\delta + (1+\gamma)(1-\delta)} \right) (y_1 + z_1) \\ &\equiv (1-B)(y_1 + z_1) \end{aligned} \quad (5)$$

Note that when $\gamma = 0$, (5) simplifies to $c_1^* = (1-\delta)(y_1 + z_1)$, as it should. If $c_1^* \leq x_1 + z_1$, the agent will consume c_1^* ; 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 $c_1^* \leq x_1 + z_1 = (1-\delta)y_1 + z_1$ if $z_1 < z_1^*$, where z_1^* satisfies

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

so that $z_1^* = \gamma(1 - \delta)y_1$. Note that when $\gamma = 0$, so there is no self-control problem, so $z_1^* = 0$: in the absence of self-control costs it is never optimal to spend all of the increment to wealth.

The above establishes

Theorem 2: *If the agent receives an unanticipated reward $z_1 < z_1^*$ in the nightclub, overall utility²⁷ is*

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

If $z_1 > z_1^$ utility is*

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

To analyze risk aversion, imagine that $\tilde{z}_1 = \bar{z} + \sigma\varepsilon_1$, 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 its certainty equivalent. For small σ and $\bar{z} < z_1^*$, the agent's overall payoff is to a very good approximation given by (6). Thus relative risk aversion is constant and equal to 1, where wealth is $w = x_1 + \bar{z}_1$ so risk is measured relative to pocket cash. On the other hand, for $\bar{z} > z_1^*$, the utility function (7) is the difference between two others, one of which exhibits constant relative risk aversion relative to wealth $y_1 + \bar{z}$, the other of which exhibits

constant risk aversion relative to pocket cash $x_1 + \bar{z}$. 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.²⁸ Note in particular the asymmetry between losses, for which risk aversion is always relative to pocket cash, and sufficiently large gains for which it is relative to wealth.

We can see this effect graphically in the case of Rabin's [2000] paradox of risk aversion in the small and in the large:

“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.”

This is paradoxical in the sense that many people will turn down the small bet, but few if any would turn down the second.

We can easily explain these facts in our model using 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 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 finding a cash machine or bank.

The easiest calculations are for the case where the gain \$105 is smaller than the threshold z_1^* . In this case, logarithmic utility requires the rejection of the gamble if pocket cash

x_1 is \$2100 or less.²⁹ In order for \$105 to be smaller than the threshold, we require $\gamma \geq 105/x_1$, so if $x_1 = 2100$ is pocket cash, we need a $\gamma > .05$, while for $x_1 = 300$, γ must be of at least 0.35.³⁰ However, the conclusion that the gamble should be rejected also applies in some cases where the favorable state is well over the threshold. For example, if pocket cash is \$300, $\gamma = 0.05$, and wealth is \$300,000, then the favorable state of \$105 will be well over the threshold is \$15, but a computation shows that the gamble should be rejected, and in fact it is not close to the margin.³¹ 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. 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.

Sheffrin and Thaler [1988] argue that people use mental accounting to mitigate self-control problems, and that the propensity to spend out of the “current income account” is near 1. Our model provides an explicit micro-foundation for their argument and links it to the issue of risk aversion; we believe that a similar banking model could be used to explain the fact that subjects appear to be “more patient” in temporal trade-offs that involve larger sums of money, a fact that is hard to explain in models without some way of smoothing consumption over time.^{32,33}

We should point out that other models can yield these results. In particular, in this specific case, the quasi-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 the quasi-hyperbolic model with a bank has an equilibrium that is equivalent to the solution without self-control problems.³⁴ Our model of response to an unanticipated shock is to maximize the utility function

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

Denoting the quasi-hyperbolic discount factor by β , the quasi-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, and if the target pocket cash holding is \$300 or more, then it is optimal to reduce pocket cash by the cash machine limit of \$300. Of course if the target pocket cash is 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.³⁵

The implications of the theory for experiments are ambiguous and complicated. The theory explains why we see substantial risk aversion in experiments with immediate money payoffs. The theory also predicts that people will be less risk averse about gambles that resolve in the future, provided they do not resolve in the “night club” state, but that agents will be more risk averse when faced with future gambles that will be paid in cash and resolved when the agent is in a “nightclub” environment.³⁶ We are fairly confident that this is true about gambles a year or more away, but rather less confident about gambles that pay off in a few days or a few weeks. This brings us back to the previously mentioned fact that the psychology of money is not well understood. Unfortunately, the theory 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.

Credit cards and checks also pose complications in applying the theory, as for some 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, 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, as neither debit card statements nor checks show what was purchased, and debit cards may only show the location of the machine and not the name of its operator.

We should emphasize that this version of the nightclub model is very stark in both its assumptions and in the prediction of 100% propensity to consume out of small amount of cash. The explanation of the Rabin paradox, however, is driven by the fact that the propensity to consumer out of cash is higher than the propensity to consume at the bank. For example, if there were several nightclub period before returning to the bank, in the earlier periods some self-control would be used, and the marginal propensity to consume out of cash would be less than 100%. Nevertheless, the marginal propensity to consume out of small amounts of cash would be higher than if the money was received at the bank, and we conjecture that the qualitative features of the analysis would be preserved, though a careful analysis of this extension would require additional work.

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. Thus for choices between gambles that have only small

gains, the “fear of theft” model predicts little risk aversion, while the dual self predicts that risk aversion will be substantial.

4. Procrastination and Delay

O’Donoghue and Rabin [1999], [2001] use quasi-hyperbolic preferences to explain how self-control problems can lead to procrastination and delay. We can use the dual-self model to make similar predictions. Specifically, consider the following model: Every period $t = 1, 2, \dots$ 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 , whose value is known at the start of that period; think for example that the leisure activity is playing outside and its utility depends on the weather. 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 $[\underline{x}, \bar{x}]$, and mean $\mu = Ex$. Taking the action ends the game, and results in a flow of utility v beginning next period, and so gives a present value of

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

If the agent waits, the problem repeats in the next period.³⁷

Except for the use of the dual-self model, this model is very similar to that of O’Donoghue and Rabin [2001], who consider quasi-hyperbolic preferences in a stationary, a and deterministic environment.³⁸ We compare the 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 maximum utility that the agent can attain in any period is x_t ,

which is the payoff to waiting, while doing it now requires foregoing x_t . Hence waiting incurs no self-control cost, and acting has self-control cost of γx_t .

Let $W(x, \gamma)$ denote the value starting tomorrow when using cut-off x , $x^*(\gamma) = \arg \max_x W(x, \gamma)$, and $W^* = W(x^*(\gamma), \gamma)$, the value of waiting today when the leisure activity is worth x is $x + \delta W^*$, and the value of acting is $-\gamma x + \delta V$. The expected present value tomorrow if the action is not taken today and cutoff rule x^* is used in the future is then given by

$$W(x^*, \gamma) = P(x^*)(-\gamma E(x | x < x^*) + \delta V) + (1 - P(x^*))(E(x | x > x^*) + \delta W^*),$$

so that

$$W^* = \frac{P(x^*)(-\gamma E(x | x < x^*) + \delta V) + (1 - P(x^*))E(x | x > x^*)}{1 - \delta + \delta P(x^*)}. \quad (8)$$

We assume that when the opportunity cost is at its lowest possible value \underline{x} , the present value of acting is greater than the present value of waiting forever:

$$\delta V > \underline{x} + \frac{\delta \mu}{1 - \delta}, \text{ or equivalently, } v > \left(\frac{1 - \delta}{\delta} \right) \underline{x} + \mu.$$

When the reverse inequality holds, an agent without self-control costs will never act. For the time being we also assume that $\delta v < \bar{x}$, so that an agent without no cost of self-control would choose to delay when x is close to \bar{x} . This combination of assumptions rules out the deterministic case studied by O'Donoghue and Rabin.

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

(ii) When $\underline{x} < x^* < \bar{x}$,

$$-\gamma x^* + \delta V = x^* + \delta W(x^*, \gamma) \quad (9).$$

Moreover, $dx^*/d\gamma < 0$, so expected waiting time is increasing in the cost of self control.

Proof: See the Appendix.

Let x_n (for “naïve”) be the cutoff the agent would use if forced to make a choice between stopping now and never stopping. Then

$$-\gamma x_n^* + \delta V = x_n^* + \delta \frac{\mu}{(1 - \delta)}$$

When the optimal rule is interior, the agent eventually does choose to stop, and so the continuation value W^* exceeds $\delta\mu/(1 - \delta)$, and equation (9) shows that $x^* < x_n$. This shows that the standard “option value of waiting” consideration carries over to the dual-self model.

In independent work, Miao [2005] applies a dual-self model to a very similar problem. He considers the cases of immediate costs and future benefits, as here; immediate benefits and future rewards, where the temptation is to act too soon, and immediate costs and immediate rewards. His model of future benefits differs from ours in that the long-run benefit is stochastic and the short-run cost is fixed, as in the literature on wage search.³⁹

O’Donoghue and Rabin [2001] analyze the implications of quasi-hyperbolic discounting in a deterministic and stationary infinite-horizon stopping-time problem. They consider a range of assumptions about the agent’s beliefs about his own future preferences; we will focus on the case of “sophisticates” who are correctly perceive their own self-control problems. For these

sophisticates, they show that in any pure-strategy equilibrium there is an “intended delay” d such that in every period where the agent plans to do the task, he must predict that waiting would lead to a delay of exactly d periods. We view these cyclic equilibria as artificial and unappealing; they also complicate the analysis of the equilibrium set.⁴⁰

One obvious difference between their conclusions and ours is that in a deterministic setting, our model predicts that the agent either acts at once or never acts; partial delay is never observed. Some readers have suggested that our model makes the wrong prediction here, but we are not convinced of this. Most real-world stopping time problems have both a stochastic component (either in costs and benefits, or simply in feasibility of acting at certain dates) and some non-stationarity, and we are unaware of relevant experimental data. Moreover, the assumptions of Theorem 3 concern only the support of the distribution of opportunity costs, and not the distribution itself; the assumptions are consistent with there being a very high probability of costs in a small interval. Thus, the appropriate question is not whether observed behavior resembles the prediction of our model in a deterministic stationary environment, but the calibration question of whether the amount of uncertainty and non-stationarity required for observed delays is plausible. In a related vein, note that if O’Donoghue and Rabin sophisticates were constrained to place a small minimum probability on each action, as in a trembling-hand perfect equilibrium, then their model predicts there would be delay cycles; we would be surprised if such cycles were observed.

To further explore the relationship between the dual-self and quasi-hyperbolic models, in the appendix we extend the O’Donoghue and Rabin analysis of sophisticates to the stochastic environment presented above. Our analysis shows that when the support of the random term is sufficiently broad, the quasi-hyperbolic model has a stationary equilibrium that is qualitatively

similar to the unique outcome of the dual-self model. Unlike the dual self model, though, there are typically has multiple equilibria. For example, there can be equilibria in which the agent uses cut-off x^1 in odd-numbered periods and uses cut-off x^2 in even-numbered periods.⁴¹ Of course one can get a unique stationary equilibrium from the quasi-hyperbolic model here simply by fiat, that is by deciding to rule out all other equilibria; we think it is more satisfactory to have this be a result than an assumption.

Theorem 4:

a) *If $v > \mu$ and $\bar{x} > \delta\beta v$, the “sophisticated quasi-hyperbolic model” has a stationary equilibrium with $\underline{x} < x^{**} < \bar{x}$.*

b) *There is an open set of parameters that satisfies the restrictions of part a) and for which there are other equilibria.*

The proof is in the Supplementary Appendix.

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 quasi-hyperbolic preferences and “lack of sophistication,” meaning that consumers misperceive their own quasi-hyperbolic parameter and thus incorrectly forecast their health club usage.⁴² 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” in waiting.⁴³ 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.⁴⁴ 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.⁴⁵

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. The solution in our model depends on γ : If $\gamma < 10/3$, the agent acts at the start, and if prevented from doing so, then act immediately when given the chance; if $\gamma > 10/3$, the agent never acts. Thus (except for the knife-edge case of $\gamma = 10/3$) the agent never plans to act in intermediate periods. The equilibrium of the quasi-hyperbolic model with sophisticates depends on β . For a range of values around $\beta = 1/2$, as the sophisticated agent acts in the second period, but intends to wait until the fourth period if prevented from acting when planned. 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.⁴⁶

5. Cognitive Load and Non-Linear Cost of Self Control

In this section we show how to extend our dual-self model to capture the effect of cognitive load on the cost of self-control. This is of interest both as an illustration of the scope and tractability of the model, and as a motivation for exploring the case of a non-linear cost of self-control report. Our starting point is the following experiment by Shiv and Fedorikhin [1999]:

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.⁴⁷ 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.⁴⁸

The finding that increasing cognitive load increases responsiveness to temptation can be easily captured in our model by the assumption that the marginal cost of self-control is higher when the long-run self, which we identify with cognitive processing, has other demands on its resources. In order to accomplish this, we need to allow the cost of self-control to be dependent on the state y . Let d be the cognitive resources consumed by tasks other than self-control. In this application, it is natural to identify the state y as the level of cognitive load d . This enables us to consider situations in which there are different levels of cognitive load.

To use our framework formally, we need to specify how the state, as measured by the external cognitive load d , impacts on self-control costs. One possibility is to 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 are consuming d resources, the utility of the short-run self can be reduced by at most $D - d$; greater self-control is simply infeasible. Alternatively, we suppose that the cognitive center does not face a fixed resource constraint, but instead has increasing marginal cost as d increases. In the former case, Assumption 2 is violated, since some actions have infinite costs. Moreover, neither version makes sense with the linearity assumption, 5'. Specifically, we want to explain how ranking are changed by self-control costs, so that we want the cognitive load d to change the ranking of alternatives.

To make a formal connection with our model, let the possible actions be h (chocolate) and f (fruit). The short-term utilities (before subtracting any costs of self control or cognitive load) are u^h and u^f , with $u^h > u^f$, so there is no self-control cost in choosing h . To make this into a self-control problem we also assume that the long-term utility of fruit is higher than that of chocolate. Recall that we identify the state y with the external cognitive load d . Then, from assumption 5, the self-control cost of action f with load d is $C(d, f) = c(d, u^h, u^f)$. To explain the data we need this cost to be increasing in d .

One way to explain the Shiv and Fedorikhin data would be to set

$$C(d, f) = d \cdot (u^h - u^f).$$

This linear specification is consistent with the linearity assumption A5', but seems unsatisfactory: If using cognitive resources for memorization increases the marginal cost of self-

control, then we would expect that using these resources for self-control should also change the marginal cost of self control, and hence the ranking of alternatives. That is, we would expect that having a more attractive foregone alternative should have the same effect as a higher value of d .

We therefore prefer the non-linear specification

$$C(d, f) = g(d + u^h - u^f) - g(d),$$

where g is an increasing convex function with $g' > 0$ and $g'' > 0$. While the linear and non-linear formulations are equivalent on the data of the “deserts” experiment, they make different predictions about how behavior changes when a third alternative is added, as we show by example in the next section. In addition, the non-linear formulation, but not the linear one, fits the psychological evidence that self-control is a limited resource, as discussed by Muraven, Tice and Baumeister [1998] and Muraven and Baumeister [2000].

To accommodate this additive cost of load, we weaken Assumption 5 as follows. Let $d(y)$ be the amount of cognitive load in state y .

Assumption 5'' (Cost of Self Control with Cognitive Load) *If $d(y) = d(y')$,*

$$\max_{a'} u(y, 0, a') \geq \max_{a'} u(y', 0, a') \text{ and } u(y, 0, a) \leq u(y', 0, a) \text{ then } C(y, a) \geq C(y', a).$$

To make a precise connection between this example and our general model, let r^* be the minimal amount of self-control needed to choose f , which we assume to be a constant that is not influenced by the state. Set

$$u(d, 0, h) = u^h - g(d); u(d, 0, f) = u^f - g(d), \text{ and } u(d, r^*, f) = u^f - g(d + r^*).$$

By definition, $C(d, f) = u(d, 0, f) - u(d, r^*, f)$, so that the given utility functions imply the non-linear cost specification $C(d, f) = g(d + r^*) - g(d)$.

One could also explain the effect of cognitive load in the quasi-hyperbolic model by assuming that the parameter 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 [1996] theory that vividness influences the effect of temptation.

6. Axiomatizations of Self Control

At this point, we would like to relate our model to the axiomatizations of Gul and Pesendorfer and Dekel, Lipman, and Rustichini.⁴⁹ As we indicated earlier, their axioms need not be satisfied if we make Assumption 5 but not the stronger Assumption 5'. To make this claim precise, we use a dynamic programming approach to reduce our dynamic problem to the two-period problems considered in Gul and Pesendorfer [2001] and Dekel, Lipman and Rustichini [2005].⁵⁰ So we now suppose that the initial period action a has no utility consequences for the short-run self, and serves solely to establish the state y determining utility possibilities starting in period 2, as the banking action does in the model of Section 4. For any initial period choice of action, we can then consider the feasible actions in period 2. These period-2 actions determine both the second-period utility u of the short-run self, and also (by determining the future value of the state) the flow utilities u at all future dates; these correspond to a present value of V for the long-run self starting in period 3. Thus self-control is only an issue in period 2.

From the viewpoint of the long-run agent at the start of the second period, what matters are the feasible utility consequences (u, V) , so we can think of the long-run self in the initial period as choosing a set of feasible (u, V) pairs. In our model, the resulting utility to the initial long-run self choosing the set W is given by

$$\Phi(W) = \max_{(u,V) \in W} [u + \delta V - C(\max(u' \mid (u', V') \in W), u)].$$

This gives rise to a preference ordering over sets W and we can ask when and whether this preference ordering satisfies various axioms.

Gul and Pesendorfer [2001] show that their axioms are equivalent to a representation

$$\max_{(u,V) \in W} [h(u, V) + H(u, V)] - \max_{(u,V) \in W} H(u, V).$$

When

$$C(\max(u' \mid (u', V') \in W), u) = \gamma(\max(u' \mid (u', V') \in W) - u),$$

we take $h(u, V) = u + \delta V$ and $H(u, V) = \gamma u$, so under Assumption 5' our model is consistent with the axioms.⁵¹

On the other hand, when C is not linear, as in our preferred explanation of the cognitive load experiment, our model is consistent with neither their axioms nor the weaker ones of Dekel, Lipman, and Rustichini. To see that a non-linear C is inconsistent with the Gul-Pesendorfer axioms, consider Example 2 of Dekel, Lipman, and Rustichini. There are three possible actions, broccoli (b), frozen yogurt (y), and ice cream (i), with $\{b, y\} \succ \{y\}$ and $\{b, i, y\} \succ \{b, i\}$. Here the frozen yogurt is a ‘‘compromise’’ option that is appealing in the face of strong temptations but

not when faced weaker ones. Dekel, Lipman, and Rustichini show that this preference over choice sets is not consistent with the Gul and Pesendorfer axioms, but it is consistent with their more general axioms. It is also consistent with our Assumptions 1-5: take $(u, V)(b) = (0, 100)$; $(u, V)(y) = (8, 30)$, and $(u, V)(i) = (14, 0)$, and let

$$\Phi(W) = \max_{(u, V) \in W} [u + .9V - .5(\max(u' \mid (u', V') \in W) - u)^2].$$

Then confronted with the choice between b and y the long-run self computes the value of broccoli to be $90 - .5(64) = 58$, which exceeds the value of yogurt, which is 35 in both the set $\{b, y\}$ and the set $\{y\}$. However, in the choice set $\{b, y, i\}$, the value of yogurt is $35 - .5(36) = 17$, the value of ice cream is 14, and the value of broccoli is $90 - .5(196) = -8$.⁵²

Note that the choice function of our agent in this example violates the independence of irrelevant alternatives, as he chooses y from $\{b, i, y\}$ but b from $\{b, y\}$. Dekel, Lipman and Rustichini propose a different explanation for the same preferences, based on the idea of that the long-run self is uncertain of the short-run self's choice function; their explanation satisfies a form of independence of irrelevant alternatives for stochastic choice functions.⁵³ Notice also that when self-control costs are linear, the action that will be chosen from a set W is $\arg \max_{(u, V) \in W} [(1 + \gamma)u + \delta V - \gamma \max(u' \mid (u', V') \in W)]$, so that the temptation $\max(u' \mid (u', V') \in W)$ is a "sunk cost" that does not effect the plan (u, V) that will be chosen from W . In other words, in the case of linear cost, the independence of irrelevant alternatives is satisfied.

In addition to leading to violations of independence of irrelevant alternatives, non-linear self-control costs also lead to a violation of the independence axiom on sets of lotteries used by Gul and Pesendorfer and Dekel, Lipman, and Rustichini.⁵⁴ To see this, consider a lottery that takes place after a is chosen by the short-run self, so that u, V represent random variables rather than constants. In our model, what matters to either the short- or long-run self is the expected present value of this lottery, so our ranking over choice sets is given by

$$\Phi(W) = \max_{(u, V) \in W} [u + \delta V - C(\max(Eu' \mid (u', V') \in W), Eu)].$$

But this ranking can violate the independence axiom, because it is not given by an expected value, but rather by a non-linear function of an expected value.⁵⁵ Our conclusion is that in this setting the independence axiom is less compelling here than in the standard model.⁵⁶ In terms of observed choices, the independence axiom implies that we would see the same distribution of choices if the two dessert options were replaced by lotteries that gave a probability p of the chosen desert and probability $1 - p$ of no desert at all. By way of contrast, our model with convex costs implies that more agents should chose fruit when the probability p of a desert is lower.

To complete the comparison of our model to the current literature on the axiomatization of preference for self-control, we show that our Assumption 5 implies set-betweenness. This shows that the linearity restriction of Assumption 5' is not needed for set-betweenness, and conversely that the set-betweenness axiom on its own is not sufficient for Gul and Pesendorfer's representation theorem. It also shows that Assumption 5 rules out the preferences in Dekel, Lipman, and Rustichini's Example 1, which are $\{b\} \succ \{b, h\}$, $\{b\} \succ \{b, p\}$, $\{b, p\} \succ \{b, h, p\}$, and $\{b, h\} \succ \{b, h, p\}$. These preferences could describe a situation where the long-run self is

uncertain which of h and p will be more tempting, so the example illustrates the fact that our model requires that no uncertainty is realized between the time that the long-run self chooses r and the time that the short-run self picks a . As we remarked in Section 2, this is more plausible when this time interval is very short.

Theorem 5: *Under Assumptions 1-5, the induced preferences Φ over choice sets satisfy set-betweenness. That is, for all choice sets W, Z either $W \succeq W \cup Z \succeq Z$ or $Z \succeq W \cup Z \succeq W$.*

Proof: See the Appendix.

In summary, our Assumptions 1-5 imply the set-betweenness property assumed by Gul and Pesendorfer, but are not consistent with their axioms taken as whole; on the other hand their axioms are more general about the nature of temptations. The extra generality obtained by dropping linear costs is needed to explain the effect of self-control when the short-run self is choosing between lotteries, and to model the idea that self-control is a limited resource. The usual arguments and evidence for independence of irrelevant alternatives and the independence axiom are not compelling when an intermediate decision such as self-control is involved. More generally, once one leaves the domain of the standard model, where the assumptions on behavior have some normative appeal, it becomes more difficult to evaluate axioms on *a priori* grounds. In this situation insights from psychology and neuroscience, such as the presence of multiple systems within the brain that respond differentially to the timing of rewards, can be particularly useful in suggesting useful assumptions. We have chosen what we feel is the most straightforward way to capture these insights, namely to cast them as assumptions on the implicit preferences of the various systems or selves. This is simple and direct, but unlike the axiomatic approach, it does not lead to a complete characterization of the behavioral implications of what

has been assumed. We think such a characterization would be interesting and useful, but it is a topic for further research.

7. Conclusion and Discussion

Sections 2, 3, and 4 show how a tightly parameterized version of the dual-self model can be used in some simple economic settings that have previously been analyzed with quasi-hyperbolic utility. We are confident that the dual-self model would have similar advantages in other settings where quasi-hyperbolic utility has been applied, for example to Laibson's [1997] finding that quasi-hyperbolic consumers overinvest in illiquid assets as a commitment device, so that their consumption is more correlated with income shocks than it would be in the standard model. More speculatively, we conjecture that the dual-self model with a bank can explain the "magnitude effect" in time discounting. In addition, we conjecture that the dual-self model with complete information can explain the co-variance of effort and consumption that Brocas and Carrillo [2005] derive in a dual-self model where the long-run self does not know the preferences of the short-run selves: Suppose that each period the agent simultaneously chooses hours worked and consumption, and that the wage is stochastic. When the wage is high, the long-run self wishes to work more, and if the cost of self-control is non-linear, the increased effort should be accompanied by an increase in consumption (or other forms of short-run gratification) to reduce the cost of self-control.

Our resolution of the Rabin paradox shows how the dual self model can capture some sorts of context effects, as the model makes different predictions about the response to unanticipated payments depending on whether they are received on the floor of a casino or the lobby of a bank. Cues are obviously the key to understanding the effect of context in general and framing in particular. 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, ” which is the most difficult modeling issue in confronting these types of issues. This suggests that one might be able to use experimental and physiological data to determine the relevant contexts and frames. 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.

Most existing work on cues, such as Laibson [2001] and Bernheim and Rangel [2004], abstracts away from self-control costs. Laibson analyzes a “rational addiction” model. This is a model with a single, fully rational, agent, in a setting where there are two cues, namely “green lights ” and “red lights.” The utility from engaging in the addictive utility when a given light is on depends on one’s past behavior under that particular light. Because it is based on rational choice, the model has a unique equilibrium, but that equilibrium has four steady states: never indulge, always indulge, indulge iff green, and indulge iff red. This shows how the agent’s experience can determine the importance of the cues, but does not allow the agent to have a preference for self-control. Bernheim and Rangel [2004] consider an addiction model where the agent sometimes enters a “hot mode” in which he consumes the addictive good whether or not his “cool self” wants him to.⁵⁷ This model, like ours, captures a value for self-control while avoiding the multiple equilibria of a multiple-selves model, but it differs in a number of respects. Most notably, in their model self control is costless in the cool mode and infinitely costly in the hot one, and the actions of the agent in the hot mode can be called “mistakes.”⁵⁸

More broadly, from the viewpoint of observed choices there may not be much difference between regretting a mistaken choice and regretting an action that would not have been taken if one’s self-control cost had not been unusually high. Our Assumption 2 rules out infinite costs, but the fact the self-control is sometimes possible, albeit costly, is what underlies our finding that

the agent responds differently to small versus large gambles in our banking model: a sufficiently large windfall will trigger self control and the long-term perspective.

We focus on the case of “sophisticated” agents who are aware of their own self-control costs. Many papers on self-control problems 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. For example, O'Donoghue and Lowenstein [2004] consider an extension similar to allowing agents to misperceive the future value of γ . O'Donoghue and Lowenstein also point out that that if the current long-run self has correct expectations but does not care about future costs of self control, the decision making process becomes a game between the long-run selves, and the result is equivalent to the usual quasi-hyperbolic discounting.⁵⁹

All models of self-control raise the issue of how the preferences of the different selves should be weighted in thinking about welfare analysis. Because the long-run self in our model has the same short-run preferences as the short-run self, the utility of the long-run self seems to be the obvious welfare criterion, just as it is in the models of Bernheim and Rangel [2004] and of Benhabib and Bisin [2004]. This approach follows traditional welfare economics in preserving the idea of consumer sovereignty in the sense that there is a single set of preferences for the consumer that are taken as primitive for welfare analysis. In other models, consumer sovereignty may be less compelling or fail to make sense, and the appropriate notion of welfare is less obvious, for example if agents are quasi-hyperbolic, or in a dual-self model where the long-run self and short-run self have different short-run preferences.⁶⁰

An important, yet complicated, extension of the model is to allow the preferences of the short-run selves to respond in some way to future consequences. The related work of Benhabib and Bisin [2004] does allow for one form of this responsiveness. They consider a consumption-

savings model in which exercising self-control is a 0-1 and costly decision, made by a “cognitive control center” that corresponds to the long-run self in our model. Temptation is stochastic, and its strength is determined by an exogenous i.i.d. sequence of cues, so that costly self-control is only used when giving in to temptation is sufficiently costly. The equilibrium again corresponds to the solution of an optimization problem by the long-run self, who takes the behavior rule of the affective self as given. In contrast to our own paper and other related work, the behavior of the “affective self” (the mental unit that is susceptible to temptation) is not required to be myopic. Instead, the affective self’s strategy can depend on expectations of future play, but it must be independent of the distribution of temptations and the cost of self control. As an example, they suggest the case where the behavior of the affective self is a Markov-perfect equilibrium of the game where self-control is impossible. If the actual outcome is that the cognitive center does sometimes exercise self-control, this raises the question of how the adaptive selves would come to learn the equilibrium play of the wrong game.⁶¹

Instead of following the expectations-based approach, we would like to model the long-run self having “taught” the short-run self to attach positive affective weight to certain variables that have long-run consequences, as in the learning of cues. Here the use of stimulus-response models of learning may play an important role. The standard forms of these models seem to be a poor fit for many aspects of human cognition.⁶² For example, faced with no observations people will respond differently depending on their prior knowledge of a situation.⁶³ More strikingly, we know that people can learn by “figuring things out” without any external stimulus at all. However, these cognitive activities can be sensibly regarded as aspects of the long-run self, while it makes sense to model the expectations of the short-run self as arising from a process of stimulus-response learning that depends solely on the past history, and does not involve forward-

looking expectations. We would then theorize, based on introspection and casual empiricism, that the long-run self can train the short-run self by manipulating this stimulus-response learning. This leads to the possibility of the deliberate formation of “habits” with a view towards lessening the cost of self control. It also raises the conjecture that the use of “rules of thumb” like “only have sweets at dinner” aids self-control by reducing the vividness of the temptation. One way to investigate this would be to first try to identify the use of rules of thumb by observing choices and perhaps administering questionnaires, and then show agents pictures of various foods while in a scanner, asking them to choose which food to have at the end of the experiment.

Finally, we should point out that the presence of self-control leads to some potentially perverse implications for public policy. For example, as Wertenbroch [1998] discusses, a common self-control strategy is to reduce temptation by avoiding stockpiling of durable “temptation goods.” As an example, he points to the fact that some people buy cigarettes only by the pack even though cartons are much cheaper. The rationale is that this reduces the temptation of having lots of cigarettes around the house, and so mitigates the cost of self-control. Suppose a law was passed that cigarette store were open only on the first day of each month. Then some people who would otherwise have purchased cigarettes by the pack will instead choose to stockpile cartons for the month. In the face of the greater temptation, their smoking will be increased. Of course, other people may reduce or quit smoking as a result of the law, so the overall effect is ambiguous, but there is at least the potential for a public policy achieving the opposite of the desired effect. A more serious application is to that of illegalization of recreational drugs. Making sales illegal, while making possession legal or at most a mild offense raises the fixed cost of conducting a transaction.⁶⁴ The more transactions, the greater the chance the dealer will get caught. Consequently the likelihood of consumers stockpiling drugs is

increased, creating a subsequent increase in consumption due to the problem of self-control. For these types of goods a policy such as legalization together with a high excise tax (as is the case for cigarettes) may prove more effective in reducing consumption than abolition. Notice also that there is a difference for “temptation” goods that are durable and those, such as services, that are not, so criminalizing prostitution should have a different effect than criminalizing marijuana.

References

- Benabou, Roland and Pycia, Marek. "Dynamic Inconsistency and Self-Control." *Economics Letters*, 2002, 77 (3), pp. 419-424.
- Benabou, Roland and Tirole, Jean. "Willpower and Personal Rules." *Journal of Political Economy*, 2004, 112 (4), pp. 848-887.
- Benahib, Jess and Bisin, Alberto. "Modeling Internal Commitment Mechanisms and Self-Control: A Neuroeconomics Approach to Consumption-Saving Decisions." Unpublished Paper, 2004.
- Bernheim, B. Douglas and Rangel, Antonio. "Addiction and Cue-Triggered Decision Processes." *American Economic Review*, 2004, 94 (5), pp. 1558-1590.
- Bodner, Ronit and Prelec, Drazen. "The Diagnostic Value of Actions in a Self-Signaling Model," in Isabelle Brocas and Juan Carrillo, eds., *The Psychology of Economic Decisions, Vol. 1*. Oxford: Oxford University Press, 2003.
- Brocas, Isabelle and Carrillo, Juan. "The Brain as a Hierarchical Organization." Unpublished Paper, 2005.
- Camerer, Colin; Loewenstein, George and Prelec, Drazen. "Neuroeconomics: Why Economics Need Brains." *Scandinavian Journal of Economics*, 2004, 106 (3), pp. 555-579.
- Dekel, Eddie; Fudenberg, Drew and Levine, David K. "Learning to Play Bayesian Games," *Games and Economic Behavior*, 2004, 46 (2), pp. 282-303.

Dekel, Eddie; Lipman, Barton and Rustichini, Aldo. "Temptation-Driven Preferences."

Unpublished Paper, 2005.

Dekel, Eddie; Lipman, Barton and Rustichini, Aldo. "Representing Preferences with a Unique Subjective State Space." *Econometrica*, 2001, 69 (4), pp. 891-934.

DellaVigna, Stefano and Malmendier, Ulrike. "Overestimating Self-Control: Evidence from the Health Club Industry." Stanford GSB Research Paper No. 1800, 2003.

DellaVigna, Stefano and Malmendier, Ulrike. "Contract Design and Self Control: Theory and Evidence." *Quarterly Journal of Economics*, 2004, 119 (2), pp. 353-402.

Frederick, Shane; Loewenstein, George and O'Donoghue, Ted. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 2002, reprinted in Colin Camerer, George Loewenstein and Matthew Rabin, eds., *Advances in Behavioral Economics*. Princeton: Princeton University Press, 2004.

Fudenberg, Drew; Kreps, David and Levine, David K. "On the Robustness of Equilibrium Refinements." *Journal of Economic Theory*, 1988, 44 (2), pp. 354-380.

Fudenberg, Drew and Levine, David K. "Subgame-Perfect Equilibria of Finite- and Infinite-Horizon Games." *Journal of Economic Theory*, 1983, 31 (2), pp. 251-258.

Fudenberg, Drew and Levine, David K. *Theory of Learning in Games*. Cambridge, MA: MIT Press, 1998.

Gul, Faruk and Pesendorfer, Wolfgang. "Temptation and Self Control." *Econometrica*, 2001, 69 (6), pp. 1403-1436.

Gul, Faruk and Pesendorfer, Wolfgang. "Self-Control and the Theory of Consumption." *Econometrica*, 2004, 72 (1), pp. 119-158.

Gul, Faruk and Pesendorfer, Wolfgang. "Self-Control, Revealed Preference and Consumption Choice." *Review of Economic Studies* (forthcoming).

Harris, Christopher and Laibson, David. "Instantaneous Gratification." Unpublished Paper, 2004.

Haruno, Masahiko; Kuroda, Tomoe; Doya, Kenji; Toyama, Keisuke; Kimura, Minoru; Samejima, Kazuyuki; Imamizu, Hiroshi and Kawato, Mitsuo. "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*, 2004, 24 (7), pp. 1660-1665.

Hellwig, Martin. "Changing Preferences and Sequential Decisions." Essay 2 in *Sequential Models in Economic Dynamics*, Ph.D. thesis, MIT, 1973.

Krusell, Per; Kuruscu, Burhanettin and Smith, Anthony T. "Temptation and Taxation." Unpublished Paper, 2005.

Krusell, Per and Smith, Anthony T. "Consumption-Savings Decisions with Quasi-Geometric Discounting." *Econometrica*, 2003, 71 (1), pp. 365-375.

Laibson, David. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*, 1997, 112 (2), pp. 443-477.

Laibson, David. "A Cue-Theory of Consumption." *Quarterly Journal of Economics*, 2001, 116 (1), pp. 81-120.

Loewenstein, George. "Out of Control: Visceral Influences on Behavior." *Organizational Behavior and Human Decision Processes*, 1996, 65 (3), pp. 272-292.

Loewenstein, George and O'Donoghue, Ted. "Animal Spirits: Affective and Deliberative Processes in Economic Behavior." Unpublished Paper, 2004.

- McClure, Samuel M.; Laibson, David; Loewenstein, George and Cohen, Jonathan D. "Separate Neural Systems Value Immediate and Delayed Monetary Rewards." *Science*, 2004, *306* (5695), pp. 503-507.
- McIntosh, D. *The Foundations of Human Society*. Chicago: University of Chicago Press, 1969.
- Metcalfe, Janet and Mischel, Walter. "A Hot/Cool-System Analysis of Delay of Gratification: Dynamics of Willpower." *Psychological Review*, 1999, *106* (1), pp. 3-19.
- Miao, Jianjun. "Option Exercise with Temptation." Unpublished Paper, 2005.
- Muraven, Mark and Baumeister, Roy F. "Self-Regulation and Depletion of Limited Resources: Does Self-Control Resemble a Muscle?" *Psychological Bulletin*, 2000, *126* (2), pp. 247-259.
- Muraven, Mark, Tice, Dianne M. and Baumeister, Roy F. "Self-Control as a Limited Resource." *Journal of Personality and Social Psychology*, 1998, *74* (3), pp. 774-789.
- Noor, Jawwad. "Temptation, Welfare and Revealed Preference." Unpublished Paper, 2005.
- O'Doherty, John. "Reward Representations and Reward-Related Learning in the Human Brain: Insights from Neuroimaging." *Current Opinion in Neurobiology*, 2004, *14* (6), pp. 769-776.
- O'Donoghue, Ted and Rabin, Matthew. "Doing It Now or Later." *American Economic Review*, 1999, *89* (1), pp. 103-124.
- O'Donoghue, Ted and Rabin, Matthew. "Choice and Procrastination." *Quarterly Journal of Economics*, 2001, *116* (1), pp. 121-160.

- Platt, Michael L. and Glimcher, Paul. "Neural Correlates of Decision Variables in Parietal Cortex." *Nature*, 1999, 400 (6741), pp. 233-238.
- Prelec, Drazen. "Decreasing Impatience: A Criterion for Non-stationary Time Preference and 'Hyperbolic' Discounting." Unpublished Paper, 2004.
- Rabin, Matthew. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica*, 2000, 68 (5), pp. 1281-1292.
- Shefrin, Hersh and Thaler, Richard. "The Behavioral Life-Cycle Hypothesis." *Economic Inquiry*, 1988, 26 (4), pp. 609-643.
- Shiv, Baba and Fedorikhin, Alexander. "Heart and Mind in Conflict: The Interplay of Affect and Cognition in Consumer Decision Making." *Journal of Consumer Research*, 1999, 26 (3), pp. 278-292.
- Strotz, Robert H. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies*, 1955, 23 (3), pp. 165-180.
- Thaler, Richard and Shefrin, Hersh. "An Economic Theory of Self-Control." *Journal of Political Economy*, 1981, 89 (2), pp. 392-406.
- Ward, Andrew and Mann, Traci. "Don't Mind If I Do: Disinhibited Eating Under Cognitive Load." *Journal of Personality and Social Psychology*, 2000, 78 (4), pp. 753-763.
- Weber, Roberto A. "Learning and Transfer of Learning with No Feedback: An Experimental Test Across Games." Unpublished Paper, 2003.
- Wertenbroch, Klaus. "Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice." *Marketing Science*, 1998, 17 (4), pp. 317-337.

Wertebroch, Klaus; Soman, Dilip and Nunes, Joe. "Debt Aversion as Self-Control: Consumer Self-Management of Liquidity Constraints." *Journal of Consumer Research* (under revision).

Appendix

Theorem 1 (Equivalence of Subgame Perfection to the Reduced Form): *Under Assumptions 1-4, every SR-perfect Nash equilibrium profile is equivalent to a solution to the reduced form optimization problem and conversely.*

Proof: The proof will make use of the fact that the payoff to the long-run self in any SR-perfect profile is no greater than that for any equivalent reduced form strategy. This follows from the fact that the payoff to the SR-perfect profile, namely

$$\sum_{t=1}^{\infty} \delta^{t-1} \int u(y_t(h), r_t(h), a_t(h)) d\pi_t(h),$$

is less than

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

from the definition of C and the assumption of SR-perfection.

First we show that given a solution to the reduced form optimization problem, we can construct a SR-perfect Nash equilibrium that is equivalent to it. Note that it is sufficient to consider pure solutions to the reduced form optimization problem, since both players will be indifferent between the corresponding SR-perfect Nash equilibria. Suppose that σ_{RF} is a pure solution to the reduced form optimization problem. We extend σ_{RF} from H^{AY} to H by defining it to be constant as a function of the history of r . Because the costs of self-control is continuous and depends only on the current state y , there is a measurable pure strategy $r_{RF} : H \times Y \rightarrow R$ such that $r_{RF}(h, y) \in \arg \max_{\{r' | u(y, r', \sigma_{RF}(h, y)) \geq u(y, r', \cdot)\}} u(y, r', \sigma_{RF}(h, y))$. Let

σ_{SR} be a SR-perfect pure strategy profile for the short-run players such that $\sigma_{SR}(h, y, r) = \sigma_{RF}(h, y)$ whenever $r = r_{RF}(h, y)$.

We claim that that the profile (r_{RF}, σ_{SR}) is a Nash equilibrium. Since by construction the profile is SR-perfect, we need only check that it is also long-run optimal. By construction, (r_{RF}, σ_{SR}) gives the long-run self the same payoff as the solution to the optimization problem, that is $U_{LR}(r_{RF}, \sigma_{SR}) = U_{RF}(\sigma_{RF})$. If the long-run self deviates, the resulting profile $(\sigma'_{LR}, \sigma_{SR})$ will still be SR-perfect, an equivalent reduced-form strategy σ'_{RF} exists, and by the first paragraph of the proof, yields at least as high a payoff as the deviation, $U_{RF}(\sigma'_{RF}) \geq U_{LR}(\sigma'_{LR}, \sigma_{SR})$. Since σ_{RF} was a solution to the optimization problem, we have $U_{LR}(r_{RF}, \sigma_{SR}) = U_{RF}(\sigma_{RF}) \geq U_{RF}(\sigma'_{RF}) \geq U_{LR}(\sigma'_{LR}, \sigma_{SR})$, so that the long-run self did not benefit from this deviation.

Second, we show that every SR-perfect Nash equilibrium profile $(\sigma_{LR}, \sigma_{SR})$ is equivalent to a solution to the optimization problem. Consider any equivalent reduced form strategy σ_{RF} . Suppose that it is not optimal. Note that by the first paragraph it yields at least as much utility as the equilibrium. Then there is a pure reduced form strategy \hat{s}_{RF} and a positive constant $\varepsilon > 0$ such that \hat{s}_{RF} yields a gain of ε in the optimization problem over σ_{RF} . By Assumption 4, following each history h and y , we may find a $r(h, y)$ such that $\hat{s}_{RF}(h, y)$ is the unique best response of the short-run player and by Assumption 3, we may assume that

$$u(y(h), r(h, y), \hat{s}_{RF}(h, y)) \geq u(y(h), 0, s_{RF}(h, y)) - C(y(h), \hat{s}_{RF}(h, y)) - \varepsilon/2.$$

Let $\tilde{\sigma}_{LR}$ be the strategy of playing $r(h, y)$. By SR-perfection, the short-run player strategy σ_{SR} must play the unique best response $\sigma_{SR}(h, y, r(h)) = \hat{s}_{RF}(h, y)$. But this gives utility to the long-

run self following each history within $\varepsilon/2$ of \hat{s}_{RF} , and so the same is true of the present value payoff. This implies that $\tilde{\sigma}_{LR}$ gives a gain of $\varepsilon/2$ over σ_{LR} , contradicting the hypothesis that σ_{LR} was an equilibrium strategy.

☑

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

(ii) When $\underline{x} < x^* < \bar{x}$,

$$-\gamma x^* + \delta V = x^* + \delta W(x^*, \gamma) \quad (9).$$

Proof:

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

$$\delta V - \gamma\underline{x} = \frac{\delta v}{(1-\delta)} - \gamma\underline{x},$$

so acting is better if $(1-\delta)\underline{x} + \delta\mu < \delta v - (1-\delta)\gamma\underline{x}$ or $\delta v > (1-\delta)(1+\gamma)\underline{x} + \delta\mu$.

(ii) If the optimal cutoff x^* is in the interior of $[\underline{x}, \bar{x}]$, optimality implies that the agent is indifferent between waiting and acting when $x = x^*$. Thus $-\gamma x^* + \delta V = x^* + \delta W(x^*, \gamma)$, which establishes (9).

Because x^* maximizes steady state payoff, $W_x|_{x=x^*} = 0$, Thus

$$(1 + \gamma) \frac{dx^*}{d\gamma} = -x^* - W_\gamma(x^*(\gamma), \gamma),$$

so $\frac{dx^*}{d\gamma} < 0$ if $x^* > -W_\gamma(x^*(\gamma), \gamma)$

Now from (8) and the envelope theorem

$$\begin{aligned} -W_\gamma &= \frac{P(x^*)}{1 - \delta + \delta P(x^*)} E(x | x < x^*) \\ &< \frac{P(x^*)}{(1 - \delta)1 + \delta P(x^*)} x^* \\ &< x^*. \end{aligned}$$

So $\frac{dx^*}{d\gamma} < 0$.

□

Quasi-Hyperbolic Preferences in the Stochastic Procrastination Problem

Consider a game between selves, one at each date t , where the date- t self has payoff $\delta\beta V$ if it acts, payoff

$$x_t + \beta\delta(x_{t+1} + \delta x_{t+2} + \dots + \delta^\tau V)$$

if a future self acts at date $t + \tau$, and payoff 0 if no self ever acts.

First we look for a stationary equilibrium, in which each agent acts if its opportunity cost is less than a cut-off x^{**} . Let W^{**} be the expected discounted sum of payoffs from tomorrow on if the agent does not act, so that the payoff to waiting at t with leisure opportunity x_t is $x_t + \beta\delta W^{**}$. Then

$$W^{**} = P(x^{**})(\delta V) + (1 - P(x^{**}))(E(x | x > x^{**}) + \delta W^{**}),^{65}$$

so that

$$W^{**} = \frac{P(x^{**})(\delta V) + (1 - P(x^{**}))E(x | x > x^{**})}{1 - \delta + \delta P(x^{**})}. \quad (10)$$

The agent is indifferent about acting when

$$\delta\beta V = x^{**} + \delta\beta W^{**}, \text{ or}$$

$$x^{**} = \delta\beta(V - W^{**}). \quad (11)$$

In the supplementary on-line appendix we prove the following result

Theorem 4:

a) If $v > \mu$ and $\bar{x} > \delta\beta v$, the “sophisticated quasi-hyperbolic model” has a stationary equilibrium with $\underline{x} < x^{**} < \bar{x}$.

b) There is an open set of parameters that satisfies the restrictions of part a) and for which there are other equilibria.

PREFERENCES ON MENUS

Theorem 5: Under Assumptions 1-5, the induced preferences Φ over choice sets satisfy set-betweenness. That is, for all choice sets W, Z either $W \succeq W \cup Z \succeq Z$ or $Z \succeq W \cup Z \succeq W$.

*Proof*⁶⁶: Let (u_W, V_W) be the act that is chosen from W , (u_Z, V_Z) be the act that is chosen from Z , and (u_Q, V_Q) be the act that is chosen from $Q = W \cup Z$, and let w', z' , and q' be the corresponding temptations, that is, $w' = \max_{w \in W} u$, $z' = \max_{z \in Z} u$ and $q' = \max(w', z')$.

Then

$$\begin{aligned}\Phi(Q) &= u_Q + \delta V_Q - C(q', u_Q) \\ &= \max_{(u, V) \in W \cup Z} \{u + \delta V - C(q', u)\}.\end{aligned}$$

Suppose without loss of generality $(u_Q, V_Q) \in W$. Then

$$\begin{aligned}\Phi(Q) &= \max_{(u, V) \in W} \{u + \delta V - C(q', u)\} \\ &\leq \Phi(W)\end{aligned}$$

so $W \succeq W \cup Z$.

Moreover, if $q' = z'$ then

$$\begin{aligned}\Phi(Q) &\geq \max_{(u,V) \in Z} \{u + \delta V - C(q', u)\} \\ &= \max_{(u,V) \in Z} \{u + \delta V - C(z', u)\} \\ &= \Phi(Z)\end{aligned}$$

so $W \cup Z \succeq Z$, and set-betweenness is satisfied, while if $q' = w'$ then

$$\begin{aligned}\Phi(Q) &= \max_{(u,V) \in W} \{u + \delta V - C(w', u)\} \\ &= \Phi(W)\end{aligned}$$

so $W \cup Z \sim W$.



Supplementary Appendices

(to be posted on the web only)

Appendix S1: The Consumption-Savings Model

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

Consider the problem of maximizing

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

over all feasible plans \vec{a} , i.e. plans that satisfy $a_t \in [0,1]$ and the wealth equation $y_t = Ra_{t-1}y_{t-1}$. 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 $u(0) = \lim_{c \rightarrow 0} u(c) = -\infty$. Let \bar{U} be the supremum in this problem.

Proposition S1.1: *Suppose $R > 1$ and*

$$(S1.1) \quad \sum_{t=1}^{\infty} \delta^{t-1} u(y_0 R^t) < \infty.$$

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

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

Proof: (i) Define $\chi_+(x) = \max\{x, 0\}$ and $\chi_-(x) = -\chi_+(-x)$.

We can write any finite sum as the sum of negative and positive parts, so for any sequence (\vec{a}, \vec{y})

we have

$$\begin{aligned} & \sum_{t=1}^T \delta^{t-1} [(1 + \gamma)u((1 - a_t)y_t) - \gamma u(y_t)] = \\ & \sum_{t=1}^T \delta^{t-1} \chi_+ [(1 + \gamma)u((1 - a_t)y_t) - \gamma u(y_t)] + \\ & \sum_{t=1}^T \delta^{t-1} \chi_- [(1 + \gamma)u((1 - a_t)y_t) - \gamma u(y_t)] \end{aligned}$$

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 $\vec{a}^n \rightarrow \vec{a}^*$. 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 $U(\vec{a}^n) \rightarrow \bar{U}$. Choose any $\varepsilon > 0$ and pick N large enough that $\bar{U} - U(\vec{a}^n) < \varepsilon/2$ for all $n > N$. If we now pick T that

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

we know that

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

for all $n > N$ and $\tau > T$. Since per-period payoffs are continuous at any (a, y) with $a > 0$,

$\vec{a}^n \rightarrow \vec{a}^*$, and $y_t^n \rightarrow y_t^*$, it follows that

$$\bar{U} - \sum_{t=1}^{\tau} \delta^{t-1} [(1 + \gamma)u((1 - a_t^*)y_t^*) - \gamma u(y_t^*)] \leq \varepsilon \text{ for all } \tau > T.$$

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

(iv) Now consider a feasible sequence (\vec{a}^n, \vec{y}^n) with $U(\vec{a}^n) \rightarrow \bar{U}$. 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 S1.1 that an optimum \vec{a}^* exists.

Proposition S1.2: *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 $\vec{a}^2 = (a_1^*, a_1^*, a_2^*, a_3^*, \dots)$ 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 $\vec{a}^n = (a_1^*, a_1^*, \dots, a_1^*, a_2^*, a_3^*, \dots)$ that play a_1^* for the first n periods such that $U(\vec{a}^n) = U(\vec{a}^*) = \bar{U}$. Clearly \vec{a}^n converges in the product topology to the plan of choosing the fixed savings rate a_1^* . From Proposition A.1 (iii), this limiting plan is feasible and gives utility \bar{U} ; that is, it is optimal.

□

III. We have shown that 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

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

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

$$dU/da \propto [(1+\gamma)(1-a) - \gamma(1-a)^\rho] \delta R(Ra)^{-\rho} - (1+\gamma)(1-\delta(Ra)^{1-\rho})$$

which gives the first-order condition for an optimum⁶⁷

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

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

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

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

$$\begin{aligned}
U_2(y_2) &= \frac{y_2^{1-\rho}}{1-\rho} \frac{(1-a^*)^{1-\rho}}{1-\delta(Ra^*)^{1-\rho}} - \frac{1}{(1-\delta)(1-\rho)} \\
&= \frac{y_2^{1-\rho}}{1-\rho} \frac{1}{(1-a^*)^\rho} - \frac{1}{(1-\delta)(1-\rho)} \\
&= \frac{y_2^{1-\rho}}{1-\rho} \frac{1}{(1-\delta^{1/\rho}R^{(1-\rho)/\rho})^\rho} - \frac{1}{(1-\delta)(1-\rho)}
\end{aligned}$$

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

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

The first order condition for optimal consumption is

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

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

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

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

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

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

Appendix S2: Hyperbolic Procrastination and Delay

Theorem 4:

a) If $v > \mu$ and $\bar{x} > \delta\beta v$, the “sophisticated quasi-hyperbolic model” has a stationary equilibrium with $\underline{x} < x^{**} < \bar{x}$.

b) There is an open set of parameters that satisfies the restrictions of part a) and for which there are other equilibria.

proof: Substituting (10) into (11) we see that

$$\begin{aligned} x^{**} &= \delta\beta \left(\frac{V - \delta V + \delta P(x^{**})V - P(x^{**})(\delta V) - (1 - P(x^{**}))E(x | x > x^{**})}{1 - \delta + \delta P(x^{**})} \right) \\ &= \\ &\delta\beta \left(\frac{v - (1 - P(x^{**}))E(x | x > x^{**})}{1 - \delta + \delta P(x^{**})} \right). \end{aligned}$$

$$\text{Let } F(x^{**}) = x^{**} - \delta\beta \left(\frac{v - (1 - P(x^{**}))E(x | x > x^{**})}{1 - \delta + \delta P(x^{**})} \right).$$

To prove the theorem it suffices to show that there is an x^{**} where $F(x^{**}) = 0$. The assumption

that $v > \mu$ implies that $F(0) = 0 - \delta\beta \left(\frac{v - \mu}{1 - \delta} \right) < 0$, and the assumption that $\bar{x} > \delta\beta v$

implies that $F(\bar{x}) = \bar{x} - \delta\beta v > 0$.

b) Suppose that $x^1 = \underline{x}$ and $x^2 = \bar{x}$ so that the odd-numbered agents never act and even ones always act. Then the equilibrium payoff of an even-numbered agent is $\delta\beta V$, and the payoff of an even-numbered agent with cost x_t who chooses to wait is $x_t + \delta\beta\mu + \delta^3\beta V$, so the even agents' strategy is a best response for all x_t if

$\delta\beta V(1 - \delta^2) > \bar{x} + \delta\beta\mu$, or equivalently

$$\delta\beta(v(1 + \delta) - \mu) > \bar{x}. \quad (\text{S2.1})$$

The payoff of the odd-numbered agents who wait is $x_t + \delta^2\beta V$, and the payoff to acting is $\delta\beta V$, so waiting is better for all x_t if

$$\underline{x} > \delta\beta(1 - \delta)V = \delta\beta v. \quad (\text{S2.2})$$

To complete the proof we must show that there is an open set of parameters such that (S2.1) and (S2.2) and the restrictions $v > \mu$ and $\bar{x} > \delta\beta v$. To do this, fix $v > 0$ and $\delta, \beta \in (0, 1)$, and some $\varepsilon > 0$. Set $\underline{x} = (1 + \varepsilon)\delta\beta v$, $\mu = (1 + 2\varepsilon)\delta\beta v$, and $\bar{x} = (1 + 3\varepsilon)\delta\beta v$. (Note that these conditions are consistent with a range of distributions, including the uniform.) By construction this satisfies $\bar{x} > \delta\beta v$ and (S2.2).

If $\varepsilon < \frac{1 - \delta\beta}{2\delta\beta}$ then $\mu = (1 + 2\varepsilon)\delta\beta v < (1 + \frac{1 - \delta\beta}{\delta\beta})\delta\beta v = v$, and for

$$\varepsilon < \frac{\delta(1 - \beta)}{(3 + 2\delta\beta)} \text{ we compute that}$$

$$\begin{aligned}
& \delta\beta(v(1+\delta) - \mu) = \delta\beta(v(1+\delta) - (1+2\varepsilon)\delta\beta v) \\
& = \delta\beta v(1 + \delta(1-\beta) - 2\varepsilon\delta\beta) \\
& > \\
& \delta\beta v(1 + 3\varepsilon) \\
& = \bar{x}.
\end{aligned}$$

This shows that there is a ‘2-cycle equilibrium’ whenever ε is sufficiently small. Since the inequalities in (12) and (13) hold strictly for the specified relationship between the parameters, they hold for an open set of β, δ, V, μ ; the inequalities also hold for a range of distributions with the given mean and endpoints.

☑

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 Bart Lipman, Mihai Minea, Drazen Prelec, Antonio Rangel, and three referees for detailed comments on earlier drafts. 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.

³ The economics literature on time-inconsistent preferences, following Strotz [1955] and Laibson [1997], uses the now-familiar (β, δ) form. This is called ‘quasi-hyperbolic discounting’ 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 ‘decreasing impatience.’

⁴ They say ‘Parts of the limbic system associated with the midbrain dopamine system [...] are preferentially activated by decisions involving immediately available rewards. In contrast, regions of the lateral prefrontal cortex and posterior parietal cortex are engaged uniformly by intertemporal choices irrespective of delay. Furthermore, the relative engagement of the two systems is directly associated with subjects’ choices, with greater relative fronto-parietal activity when subjects choose longer term options.’

⁵ Crude because reward-related information seems to be processed in many different brain regions, and because the links between these regions are more complex than the top-down control assumed in our model. (See for example

O’Doherty [2004] and Platt and Glimcher [1999].) But as economists our primary goal is a model that fits reasonably well, and is simple enough to apply in a range of economic situations. Consistency with what is known about brain function is a plus, as it gives more reason to be confident about the model’s predictions, but this is not our primary objective.

⁶ The only model of quasi-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.

⁷ This is a very stark assumption, but it leads to a much simpler model, and may be a reasonable approximation in some cases of interest. The conclusion discusses the complications introduced by forward-looking “short-run selves.”

⁸ Bernheim and Rangel [2004] favor the “dual systems” terminology.

⁹As we explain below, the cash effect has the same impact in the quasi-hyperbolic discounting model. Note also that this result is in the opposite direction from those of Gul and Pesendorfer [2001, 2004], who do not consider environments with mechanisms such as banks that substitute for self-control.

¹⁰ Benabou and Pycia [2002] analyze a two-period model where the long-run self and the short-run self compete for control by expending resources, with the probability that a given self takes control equal to its share of the total expenditure; it is not clear to us how to extend their model to multiple periods. Brocas and Carrillo [2005] analyze a two-period dual self model of a consumption-leisure choice. They focus on the case where self-control is costless, but the short-run self has private information, so that the long-run self offers a “menu” of consumption-effort pairs to the short-run self, as in agency models.

¹¹ More precisely, Gul and Pesendorfer [2001] and Dekel, Lipman, and Rustichini [2005] give axioms on preferences over static choice sets; Gul and Pesendorfer [2001] complement this with a conditions for a representation of joint preferences over both choice sets and the choice from a given set. Gul and Pesendorfer [2004] extend their earlier representation to an infinite horizon, but only for preferences over choice sets.

¹² This assumption is a modeling convenience, but it can be relaxed to allow for history-dependent action sets.

¹³ Throughout the paper all measure spaces are endowed with the Borel sigma-algebra.

¹⁴ We allow this utility to take on the value $-\infty$ but not $+\infty$.

¹⁵ Because the objective function is linear in probabilities, there will be a deterministic solution to the maximization problem whenever a maximum exists. We allow for mixed strategies here to facilitate comparisons of the solutions of the reduced-form optimization with the equilibria of the game.

¹⁶ Each space of finite histories is a complete separable metric (Polish) space. Such a space is small enough to guarantee that we can find regular conditional probability measures, meaning that the conditional probabilities may be chosen to (almost surely) be countably additive – that is, constitute a probability measure.

¹⁷ Because of sets of measure zero there are generally a great many equivalent reduced form strategies.

¹⁸ Because we take the short-run self's action to be the savings rate and not total savings, the feasible actions A are independent of the long-run self's actions. Note also that since the short-run self makes all consumption/savings decisions, the model satisfies our requirement that the evolution of the state depends only on the actions of the short-run self.

¹⁹ Krusell, Kurusac, and Smith [2005] consider a more general family of preferences; the model of this section is the special case of theirs in which $\beta = 0$. (Larger values of β correspond to a temptation cost that reflects quasi-hyperbolic discounting by the “doer.”) Instead of showing that the solution is stationary, we do, they take as primitive a stationary value function that satisfies their equation (4); they then restrict attention to cases where this stationary solution is the unique limit of their finite-horizon models. As they note, equation (4) can sometime admit multiple solutions, but in our case of $\beta = 0$ equation (4) is a contraction and has a unique solution. However, since infinite-horizon games can have equilibria that are not limits of exact finite horizon equilibria, their argument seems to leave open the possibility that there are other equilibria where the stationarity of (4) is not satisfied.

²⁰ As we have written the problem, with the savings rate as the control, the state evolution equation is not concave. If we change variables so that the control is the absolute level of consumption c_t , the state evolution equation is linear but the per-period payoff becomes $(1 + \gamma)\log(c_t) - \gamma\log(y_t)$, which is not concave in the state if $\gamma > 0$. For this reason our proof technique does not rely on concavity. 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 CRRA case studied in the Supplemental Appendix, the solution given there remains unchanged provided we define $R^{1-\rho} = E(\tilde{R}^{1-\rho})$.

²¹ 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 Krusell, Kurusucu, and Smith obtain the same formula for the savings rate when they specialize their model to our case.

²² 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,

²³ This is true for example of the many studies cited in Frederick, Loewenstein, and O'Donoghue [2004], p. 173.

²⁴ 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 stimuli. 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.”

²⁵ Hellwig’s [1973] Ph.D. thesis discusses a multiple-selves model of changing preferences, and applies it to a model of banking where “in the bank, the agent has a higher preference for savings than when he sits down for dinner.” Hellwig allows the interval between visits to the bank to be endogenous, and studies how it compares to the solution to the single-agent control problem.

²⁶ The Supplementary Appendix provides the parallel computations for the CRRA case.

²⁷ Note that in the case where the probability lottery has a positive probability, the only consequence is that z_1^* changes by a small amount.

²⁸ In the CRRA case with relative risk aversion ρ , the former also dominates as δ approaches $1/R^{1-\rho}$, which is relevant when $\rho \leq 1$.

²⁹ This is because $.5 \log(2100 - 100) + .5 \log(2100 + 105) \approx \log(2100)$.

³⁰ The usual daily limit in the U.S. for ATM withdrawals is \$300.

³¹ The relationship between pocket cash and wealth depends on δ and hence on the period length. The relationship is $\delta = 1 - (x_1 / y_1)$, so for example, if pocket cash is \$300 and wealth \$300,000, then the interest rate corresponding to δ is 1/1000, or if the annual rate is 10%, the period between bank visits is 3-4 days.

³² For example, a subject might be indifferent between \$15 now and \$60 in a year, and also be indifferent between \$250 immediately and \$350 in a year. This is called the “magnitude effect;” see Frederick, Lowenstein, and O’Donoghue [2002].

³³ The mechanism in our model is very simple, and covers only the case where the current account is for one day’s spending. While we do think that cash on hand has some impact on propensity to spend, and hence on risk aversion, even myopic agents may be able to smooth consumption over a few days. To accommodate this, we would need to interpret the period length in our model as the length of time over which the short-run self is willing to smooth consumption. This is easy to do if one takes the mental frames as exogenous. However, since mental frames are heuristics for responding to cognitive limits and self-control costs, it would be more satisfying to derive the nature of the frames from assumptions on those fundamentals.

³⁴ We do not know if it is the only equilibrium.

³⁵ 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.

³⁶ We thank a referee for emphasizing the role of the environment in which these gambles will be resolved.

³⁷ Our later assumptions will imply that $v > \mu$, so that we can think of the agent continuing to enjoy leisure activities in every period after he acts. An alternative interpretation of this set up is that life consists of a series of tasks. Once a task is completed, life continues with a series of additional tasks, which we model by means of a continuation payoff. Procrastination here means delaying life in order to gain some short-term leisure.

³⁸ They suppose that there is a short-run cost $c > 0$ of acting, and no short run benefit of not acting, while we assume that the "cost" of acting is the opportunity cost of foregone leisure, but this is only a normalization

³⁹ Unlike the difference between costs and opportunity costs of acting, this is a substantive difference and not simply a normalization, as in our case the agent knows the current realization of the stochastic term when making a decision.

⁴⁰ O'Donoghue and Rabin conjecture that their model has only one additional mixed-strategy equilibrium. The pure-strategy equilibria correspond to the set of limits of exact equilibria of the finite horizon games as the horizon goes to infinity, but this restriction relies on long chains of backwards induction and is not robust to even small payoff uncertainty, as shown by Fudenberg, Kreps, and Levine [1988].

⁴¹ We suspect that there can be many other sorts of equilibria as well, but given our focus on the dual-self model a full characterization is well beyond the scope of this paper.

⁴² 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 quasi-hyperbolic and dual-self models.

⁴³ 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.

⁴⁴ 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 quasi-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 . The same concern arises in the O’Donoghue and Rabin $(\beta, \hat{\beta})$ model, where the agent in every period has quasi-hyperbolic parameter β but believes that his future play will correspond to the subgame-perfect equilibrium of the multiple-selves game with parameter $\hat{\beta}$. Dekel et al [2004] analyze this issue in games between multiple agents; their results show that when β is fixed over time, assumptions that allow the agent not to learn β are typically too weak to justify restricting attention to equilibrium. Thus, while agents may sometimes have misperceptions about their own future play, it seems hard to justify the use of equilibrium models to analyze these misperceptions.

⁴⁵ If consumers have heterogeneous expectations about whether the health club usage is worthwhile, then the agents who sign up for the health club will have relatively optimistic expectations, while those who get misleadingly negative draws will not sign up and hence not be part of the sample. Thus the explanation based on misperceptions about the attractiveness of the health club does not require that the population as a whole has a bias towards optimistic expectations.

⁴⁶ 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.

⁴⁷ They base this last hypothesis on the work of Lowenstein [1996].

⁴⁸ Ward and Mann [2000] report a similar effect of cognitive load (a combination of a memorization task and a response-time task) on individuals who a pre-experiment survey identified as “restrained eaters.”

⁴⁹ See also Noor [2005] for an alternative approach to axiomization that is based on choice functions instead of preferences. He uses preferences on the timing of rewards to identify the agent’s normative preferences as the limit of preferences on rewards in the far-distant future.

⁵⁰ Gul and Pesendorfer [2004] consider an infinite-horizon dynamic programming problem. We do not compare our model with those axioms because they are more complicated, and because this paper does not provide axioms for the choices of actions to complement its axioms on choices over choice sets.

⁵¹ Benabou and Pycia [2002] show that the equilibrium of their two-period “lobbying” model is consistent with these same axioms. As we remarked earlier, the Gul and Pesendorfer axioms are in some ways less restrictive than Assumption 5, and hence than Assumption 5’. Specifically, we constrain the “temptation” to be the foregone short-run utility, while their axioms are consistent with other “temptation functions.”

⁵² Dekel, Lipman, and Rustichini point out that these preferences are consistent with either the Gul and Pesendorfer set-betweenness axiom or their independence axiom, but not both.

⁵³ Noor [2005] gives a different example of the violation of independence of irrelevant alternatives that is not compatible with our Assumptions 1- 5: he supposes that $\{b, y\} \sim \{b\}$ and $\{b, y, i\} \succ \{b, i\}$. Intuitively, $\{b, y\} \sim \{b\}$ implies that the long run self prefers b to y and that y is not tempting in the menu $\{b, y\}$. To accommodate these preferences we could relax Assumption 5 to allow the “temptingness” of a choice to depend on the menu from which it is chosen, or drop Assumption 1 so that self-control is not always costly.

⁵⁴ The axiom in question is Axiom 3 in both papers; it was first proposed in Dekel, Lipman and Rustichini [1999]. Because it applies to combinations of sets of lotteries, it is not a consequence of the usual independence axiom on lotteries.

⁵⁵ More precisely, no monotone transformation of this function is linear in the expected value of the utilities.

⁵⁶ Gul and Pesendorfer do not specify any particular processes that might underlie their assumptions on preferences, so in particular they do not assume that self-control requires effort; our shows a difficulty with one possible interpretation of their assumptions. The conclusion of Dekel, Lipman and Rustichini mentions another possible reason for the failure of the independence axiom based on “guilt” as opposed to an intervening action.

⁵⁷ The “hot-cool” model originates in Metcalfe and Mischel [1999].

⁵⁸ Lowenstein’s [1996] notion of visceral influences causing people to “act against their self-interest in full knowledge that they are doing so” can also be interpreted this way.

⁵⁹ They also use their model to explain non-linear probability weighting in the assessment of risks.

⁶⁰ In thinking about welfare criteria, it is helpful to distinguish between political economy – how we think the government should make decisions – and moral philosophy – how we would advise others to behave. Even if we believe people do make systematic errors in evaluating how various choices will influence the appropriately-defined measure of their “welfare,” we might not trust that the government or policy analysts would make better evaluations. For this reason, it is consistent to believe both that people make mistakes and that government policy should (with a few exceptions) be based on the assumption that people’s actions and ex-ante predictions are the best guide to what is in their own interests.

⁶¹ It is of course possible that either the short-run or long-run self or both may misperceive the future. However, learning makes it less likely that these misperceptions will involve frequently experienced variables such as γ , or more broadly, one’s own ability to exercise self-control. Moreover, when little learning occurs, it is not clear why one should expect to see equilibrium play.

⁶² 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.

⁶³ For a particularly striking experiment where learning takes place in the absence of feedback, see Weber [2003].

⁶⁴ The non-linearity of penalties with respect to the amount possessed works in the opposite direction; we ignore this effect here.

⁶⁵ Because the reward V here is received two periods from the present, it is multiplied by δ but not by β .

⁶⁶ This proof closely follows the intuition in Gul and Pesendorfer [2001]; we give the proof as their formal argument uses assumptions that we have not imposed.

⁶⁷ We do not know if the first-order condition has a unique solution, except in the logarithmic case.