

E C O N O M E T R I C A

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

PDF Terms and Conditions

Please note the following Terms and Conditions regarding the use of the attached PDF file of your Article. Use of this PDF file indicates agreement to abide by the following Terms and Conditions.

- The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs.
- You may post the PDF file on your own personal website or your employer's website and/or repository. If you do so, you must provide on any pages linking to the Article a full bibliographic reference and a link to the *Econometrica* website.
- No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society.
- In any use of the Article the following statement must be included:

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

The Econometric Society
Department of Economics, New York University
19 West Fourth Street, 6th Floor, New York, NY 10012, USA
<http://www.econometricsociety.org>

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 80, No. 1 (January, 2012), 1–42

TIMING AND SELF-CONTROL

DREW FUDENBERG

Harvard University, Cambridge, MA 02138, U.S.A.

DAVID K. LEVINE

Washington University in St. Louis, St. Louis, MO 63130-4899, U.S.A.

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

TIMING AND SELF-CONTROL

BY DREW FUDENBERG AND DAVID K. LEVINE¹

The standard dual-self model of self-control, with a shorter-run self who cares only about the current period, is excessively sensitive to the timing of decisions and to the interpolation of additional “no-action” time periods in between the dates when decisions are made. We show that when the shorter-run self is not completely myopic, this excess sensitivity goes away. To accommodate the combination of short time periods and convex costs of self-control, we introduce a cognitive resource variable that tracks how the control cost depends on the self-control that has been used in the recent past. We consider models with both linear and convex control costs, illustrating the theory through a series of examples. We examine when opportunities to consume will be avoided or delayed, and we consider the way in which the marginal interest declines with delay.

KEYWORDS: Self-control, dual selves, present bias, Allais paradox, revealed preference, behavioral economics.

1. INTRODUCTION

MODELS OF LONG-RUN PLANNING and shorter-run² impulsive selves suppose that a single patient self makes decisions in each period to maximize the discounted sum of utility net of a cost of self-control, where this cost depends on the temptations faced by the impatient impulsive self. These models provide a quantitative explanation of a wide variety of “behavioral” paradoxes, including the Rabin paradox (small stakes risk aversion), the Allais paradox, preferences for commitment in menu choice, violations of the weak axiom of revealed preference, nonexponential discounting, and the effect of cognitive load on decision making and reversals due to probabilistic rewards. However, these models, like the quasihyperbolic discounting model,³ have two implications about the role of timing that are at odds with the data. First, because these models have a fixed horizon for the shorter-run self, they cannot explain overwhelming evidence that the length of delay has a continuous impact on decisions, and they make implausible predictions about the value of a commitment that avoids temptation: There is no value for a commitment that must be made the same period that the temptation is faced, while commitment “one period before” the temptation arises can be highly valued, regardless of the length of a time period. Second, past work has identified the importance of allowing the cost of self-control to be a convex as opposed to a linear function

¹This paper was presented as the Fisher–Schultz Lecture at the 2010 World Congress of the Econometric Society in Shanghai. We thank Attila Ambrus, Yuichiro Kamada, Jawaad Noor, Alex Peysakhovich, Tomasz Strzalecki, and Dmitry Taubinsky for helpful comments, and the NSF, Grants SES-0646816, SES-0851315, and SES-0954162, for financial assistance.

²As we are dealing with “short-run selves” who live more than one period, we adopt the suggestions of Rajiv Sarin and refer to “shorter-run selves.”

³Strotz (1955), Phelps and Polak (1968), Laibson (1997), and O’Donoghue and Rabin (1999).

of foregone utility, but it has assumed that the cost function in each period is independent of self-control used in the recent past. As we point out here, that model has the implausible implication that making decisions simultaneously is different from making them with an arbitrarily small delay. The new version of the dual-self model we propose here addresses both of these issues.

To account for the continuous effect of delay and to explore the implications of the timing of decisions, we propose that temptation comes from a shorter-run self or selves who are not completely myopic but who value future utility less than the long-run self does, because there is either a succession of shorter-run selves with random lifetimes or a single shorter-run self whose discount factor is lower than that of the long-run self. This lets us maintain the strength, simplicity, and applicability of the dual-self model. It also lets us model cases where agents are tempted by future consumption, as in Noor (2007, 2011), and explain why this temptation is most significant with respect to payoffs in the near future.

The key modeling question in extending the dual-self model to nonmyopic shorter-run selves is the specification of the cost of self-control. In the one-period model, this cost depends on the amount of utility the shorter-run self foregoes in the current period. When temptation comes from selves who give nonzero weight to future payoffs, expectations about these future payoffs matter; we propose that the control cost of implementing a given action depends on how much that action lowers the highest possible average present value the shorter-run self could obtain from the current period on. This specification is consistent with the interpretation that the shorter-run selves are strategically naive and evaluate foregone utility under the assumption that no self-control will be used in the future; the Supplemental Material (Fudenberg and Levine (2012)) points out some of the complications of an alternative specification.

We begin our analysis with the case where the cost of self-control is linear in the foregone value. This is the simplest version of the model, and the one closest to discounted expected utility, as it is consistent with both the independence axiom for choices over lotteries and the weak axiom of revealed preference. Our first application considers the decision of whether to accept or reject a “simple temptation” that gives an initial positive payoff followed by a negative payoff in future periods. We point out that the agent may prefer to resist a simple temptation when it is a once-and-for-all choice, and yet prefer to give in when that temptation must be faced in every period unless and until it is accepted; we relate this to the effect of “bundling” of decisions noted by Ainslie (2001) and Kirby and Guatsello (2001). We then show that the value of a commitment to avoid temptation can be nonmonotone in the discount factor of the shorter-run selves. In Example 3, we show that the model explains the fact that the ratio between payment in period t and payment in period $t + 1$ that makes subjects indifferent is typically increasing in t and not constant; we focus on data from Myerson and Green (1995), but there have been many similar studies.

Although linear costs are a convenient first cut at self-control problems, there is considerable evidence that the costs are often convex, so that it is more than twice as hard to resist twice the temptation. We therefore extend the model to allow for convex costs. Examples 4 and 5 point out two implications of convex costs for agents faced with simple temptations: First of all, an agent is more likely to resist a temptation that has a low probability of being realized than one whose payoff stream is certain. Second, an agent who is faced with two simultaneous simple temptations may choose to accept one of them, even though he would reject both if they were presented in different periods.

This observation suggests that when costs are convex and time periods are short, we expect that the nonlinearity of control costs should “spill over” from a decision in one period to a subsequent decision soon afterwards, so that making two decisions in rapid succession is similar to making the two decisions simultaneously. Moreover, since the length of the time period in decision problems is an artificial construct, we want the model to apply to cases where the time periods are very short, with decisions made in only a few of them.⁴ Adding such “intermediate” no-action periods makes no difference in classic rational models, but can have counterintuitive implications in models of self-control.

To capture the effect of changing the period length when there are convex costs, we suppose that self-control uses cognitive resources and that these resources are a stock that can be depleted over short time intervals, as argued by Muraven, Tice, and Baumeister (1998) and modeled by Ozdenoren, Salant, and Silverman (2009). From this perspective, the simpler model of the previous sections corresponds to cases where the stock of cognitive resources completely replenishes from one period to the next.

The reason for introducing a stock of cognitive resources is to track variations in the marginal cost of self-control and account for the way that using self-control in one period can alter the self-control cost and decision in a subsequent period. To check that this is all it is doing, we first show in Theorem 3 that when there is a single decision, the stock of cognitive resources is irrelevant, as the agent’s decision will be the same as in a “state-free” model with the appropriately defined cost function. We then show in Theorem 4 that when the marginal cost of self-control is constant, the decisions made by an agent with a stock of cognitive resources that partially replenishes over time are again the same as those made in an associated model without cognitive resources. However, when the agent makes multiple decisions and the marginal cost of self-control varies, the equivalence with the state-free model fails, precisely because of the link between current decisions and the marginal cost of self-control in future periods.

In general, there are three possible sources of nonlinearity in the model, any of which can cause variations in the marginal cost of self-control: the way

⁴In multiplayer games, the period length models how long it takes one player to respond to the play of another.

the stock of cognitive resources is depleted by using self-control; the way the stock is replenished over time; and way the stock provides benefits. Since cognitive resources are not observed directly and have no natural units, there are many equivalent representations of the same preferences, so it is not possible to identify which part of the model is non-linear. Indeed, as Theorem 5 shows, if there is any replenishment at all, it is without loss of generality to specify linear replenishment and lodge all of the nonlinearity in the depletion and benefit functions.

After exploring the general properties of the cognitive resources model, we consider a number of examples with linear depletion and linear (or no) replenishment. Example 7 revisits the example of an agent with convex costs facing two temptations in a row and shows that when resources replenish linearly, the agent makes the same choice whether the two decisions are made exactly simultaneously or in rapid succession. Example 8 shows that when the marginal benefits of resources are concave (so the cost of control is convex), it may be optimal to resist a persistent temptation for a while and then take it, a conclusion that is impossible in the stationary model without a stock variable. Example 9 builds on this by adding the option to pay a fee and permanently avoid the temptation. We show that it may be optimal to resist for a while and then pay the fee, which is consistent with the findings of a suggestive recent experiment of Houser, Schunk, Winter, and Xiao (2010). Examples 7 and 8 simplify by assuming no replenishment of resources at all, which is unrealistic but makes it easier to highlight the logic of the argument. Example 10 reexamines persistent temptation with a general depletion function to highlight how the depletion and benefit functions interact to determine whether the agent will resist for a while before giving in. Finally, Example 11 shows the issues involved in relaxing our assumption that the “willpower technology” is fixed and cannot be changed by the agent.

Some past work has used the device of random *long-run player* lifetimes to explain behavioral anomalies: Dasgupta and Maskin (2005) showed that uncertain long-run player lifetimes can lead to hyperbolic discounting. Halevy (2008) developed a model where a single long-run self faces a stopping (or death) risk that is modified by a convex “transformation function” and so is distinct from the agent’s pure time preference. He used this to explain Keren and Roelofsma’s (1995) data, which show that “present-biased” preference reversals are much less frequent when both the present and future rewards are uncertain.⁵ Epper, Fehr-Duda, and Bruhin (2009) used a similar idea of distorted survival weights to explain present bias as a consequence of prospect theory.

Noor (2007, 2011) developed axioms for infinite-horizon choice problems in which the agent can be tempted by future consumption. His model, like that of

⁵Our (2010) paper (Fudenberg and Levine (2010)) explains the same data as a consequence of a convex cost of self-control.

Gul and Pesendorfer (2001), is more general than ours in relating the temptation values to the objectives of the long-run player, but is less general in imposing the independence axiom; because of this assumption, only the linear-cost model in Section 3 is compatible with his framework. Noor did not investigate how “temptation by the future” depends on the real time between the two periods of his model; his main goal is to show that there can be a self-control problem despite the fact that future temptation results in little demand for commitment, as in our analysis of menu choices in Example 1. In the context of two-period models, Noor and Takeoka (2010a) developed axioms for choices on menus that allow convex costs of self-control, and then (Noor and Takeoka (2010b)) developed a more restrictive representation that they extended to multiple periods and temptation by future payoffs. Brocas and Carrillo (2008) explained the covariance of effort and consumption by assuming the long-run self has incomplete information on the shorter-run self’s cost of effort, and Chatterjee and Krishna (2009) and Dekel, Lipman, and Rustichini (2009) axiomatized cases where second-period preferences are stochastic and can depend on the first-period choice of menu.⁶ For the infinite-horizon problem, Gul and Pesendorfer (2004) developed a recursive extension of their (2001) axioms, including the independence axiom.

2. THE DECISION PROBLEM

In dual-self models, the agent acts to maximize expected discounted utility subject to a cost of self-control that is derived from the preferences of a more impulsive short run self. In this paper, we take this control cost as exogenous; the Supplemental Material explains how the maximization problem we consider corresponds to the equilibrium of a game between the agent’s long-run (LR) self (a “planner”) and a sequence of shorter-run (SR) “temptation selves” (“doers”). To facilitate the exposition and also the comparison of the model with previous work, we use a discrete-time model with periods $n = 1, 2, \dots$. We denote the agent’s discount factor by δ and suppose that the discount factor of the shorter-run self is $\delta\mu$, where $\mu \in [0, 1)$ can also be interpreted as the survival probability of the current shorter-run selves.

We will frequently be interested in how the solution to the model varies with the period length, which we would like to distinguish from the real time between decisions. To do this, we let the period length be τ units of calendar time, and suppose that $\delta = \exp(-\rho\tau)$ and $\mu = \exp(-\eta\tau)$; for small τ , we use the approximations $\mu \approx 1 - \eta\tau$ and $\delta \approx 1 - \rho\tau$.

The space of states, denoted Y , is a measurable subset of a finite-dimensional Euclidean space, as is the space of actions A . For each state, there is a measurable subset of feasible actions $A(y_n) \subseteq A$ and at least one measurable map $a: Y \rightarrow A$ that satisfies $a(y_n) \in A(y_n)$. Dynamics are Markovian: they are given

⁶These two papers differ in how the long-run self views the possible second-period preferences.

by probability distributions $\pi(y_{n-1}, a_{n-1})[dy_n]$ over states next period, conditioned by the current state and action, that are measurable functions of y_{n-1} , a_{n-1} . In applications, the state can take on many different roles: it can correspond to direct and indirect constraints on feasible actions (such as the presence of tempting consumption opportunities and the agent's wealth) and also to the agent's past consumption or other influences on the current period's utility function.⁷

Each period's action is taken after that period's state is known, so the history of play at period n is $h_n = (y_1, a_1, \dots, y_{n-1}, a_{n-1}, y_n)$; the initial history $h_1 = y_1$ is exogenously given. A strategy or plan for the long-run self is then a measurable map \mathbf{a} from histories to actions, so that for each history h_n , the strategy specifies an action in $\mathbf{a}(h_n) \in A(y_n)$.

The shorter-run self (or selves) has utility $u(y_n, a_n)$ in period n if the action a_n is taken in the state y_n . We work with average present values, so that we hold $u(y_n, a_n)$ fixed as we vary the length τ of the time period. The objective of the long-run player is the average present value of these shorter-run self utilities, minus a cost of self-control that is defined with reference to the maximum possible average present value for the shorter-run self. Let $E_{\mathbf{a}, h_n}$ be the conditional expectation generated by the plan \mathbf{a} conditional on the history h_n . The expected average present value of the shorter-run self from period n on under \mathbf{a} is given by

$$U(h_n, \mathbf{a}) \equiv E_{\mathbf{a}, h_n} (1 - \delta\mu) \sum_{\ell=0}^{\infty} (\delta\mu)^\ell u(y_{n+\ell}, a_{n+\ell}),$$

or, equivalently,

$$U(h_n, \mathbf{a}) \equiv E_{\mathbf{a}, h_n} (1 - e^{-(\rho+\eta)\tau}) \sum_{\ell=0}^{\infty} e^{-(\rho+\eta)\ell\tau} u(y_{n+\ell}, a_{n+\ell}).$$

To focus on the application of the model and not standard technical details, we directly impose the following assumption.

ASSUMPTION SR0:

$$\begin{aligned} & E_{\mathbf{a}, h_n} (1 - \delta\mu) \sum_{l=0}^{\infty} (\delta\mu)^l u(y_{n+l}, a_{n+l}) \\ &= (1 - \delta\mu) \sum_{\ell=0}^{\infty} (\delta\mu)^\ell E_{\mathbf{a}, h_n} u(y_{n+\ell}, a_{n+\ell}) \end{aligned}$$

⁷For example, the model is compatible with the "rational addiction" preference structure of Becker and Murphy (1988) and with short-run consumption commitments as in the Gabaix and Laibson (2002) explanation of the equity premium puzzle.

(the expectation and sum operators can be interchanged) and $U(h_n, \mathbf{a})$ has a maximum for each n and h_n .

Because the problem of the shorter-run self is Markov, this maximized value only depends on the state:

THEOREM 1: $\max_{\mathbf{a}} U(h_n, \mathbf{a}) = \bar{U}(y_n)$.

Our earlier work (Fudenberg and Levine (2006, 2010)) assumed that the cost of self-control depends on the amount of utility foregone by the shorter-run self, which is the difference between the maximum possible utility in the current period and the current-period utility the shorter-run self actually receives. When shorter-run selves live more than one period, we must specify how the cost of self-control takes into account the effect of current actions on future payoffs. To do this, we suppose that the level used as a benchmark is the highest value that the shorter-run player could hope to receive.

Specifically, we call $\bar{U}(y_n)$ the *temptation value* in period n starting at state y_n . The *foregone value* is then

$$\begin{aligned} \Delta(y_n, a_n) = & \bar{U}(y_n) - \left((1 - \delta\mu)u(y_n, a_n) \right. \\ & \left. + \delta\mu \int_Y \bar{U}(y_{n+1})\pi(y_{n+1}|y_n, a_n)[dy_{n+1}] \right). \end{aligned}$$

The foregone value is *recursive* in the sense that it depends on the future only through the future temptation value and it attributes control costs to each action as it occurs as opposed to entire contingent plans.⁸ The idea is that the current cost of self-control depends on the foregone value in the current period, which can be nonzero either because the current choice of action a_n lowers $u(y_n, a_n)$ or because it changes the distribution of future states $\pi(y_{n+1}|y_n, a_n)$ and thus the temptation value from the next period on. Future actions that may lower the shorter-run self's value do not incur a current cost; instead they incur costs at the time at which they are taken. Note in particular that there is no foregone value to postponing a decision, provided that the set of utility possibilities in the next period is the same as in this one, although as we will see, the cost of implementing it may change.

⁸This is the same specification of temptation value and foregone value as in Noor (2007). The Supplemental Material discusses an alternative formulation in which the shorter-run selves have rational expectations about future self-control, so each period's control cost depends on the difference between the best expected present value available to a shorter-run self born in that period and the value that will actually be received. We show that when the cost of self-control is linear in the foregone value, as in (1), the two formulations are equivalent, but with nonlinear costs, the recursive formulation has more plausible implications.

One interpretation of the foregone value is that the term

$$\delta\mu \int_Y \bar{U}(y_{n+1})\pi(y_{n+1}|y_n, a_n)[dy_{n+1}]$$

is the shorter-run self's prediction of the expected continuation payoff, and that the shorter-run self predicts that no self-control will be used in the future. Under this interpretation, the shorter-run self is strategically naive and does not anticipate that today's actions can change the amount of self-control that will be used in the future.⁹

Notice that by the principle of optimality, any plan that solves $\max_{\mathbf{a}} U(h_n, \mathbf{a})$ must also solve the dynamic programming problem

$$\max_{a_n} (1 - \delta\mu)u(y_n, a_n) + \delta\mu \int_Y \bar{U}(y_{n+1})\pi(y_{n+1}|y_n, a_n)[dy_{n+1}].$$

Thus we can make the following statement.

THEOREM 2: $\Delta(y_n, a_n) \geq 0$ and if $\hat{\mathbf{a}} \in \arg \max_{\mathbf{a}} U(h_n, \mathbf{a})$, then $\Delta(y_n, \hat{a}_n) = 0$.

The key element of the theory of self-control is the specification of the cost of exerting self-control. We study the linear case first.

3. LINEAR COST OF SELF-CONTROL

We start our analysis with a particularly simple specification of the cost of self-control. We suppose that the cost of self-control is $\Gamma\Delta(y_n, a_n)$, so it is linear in the foregone value Δ , and that the scalar constant $\Gamma \geq 0$ is independent of the state and the period length.

The case of linear self-control costs has been the most widely studied. This type of self-control model satisfies the Gul–Pesendorfer axioms, including the independence axiom. Moreover, while nonlinear costs are important in many applications, many insights still arise in the linear case. We examine the increasing marginal cost of self-control in subsequent sections, along with the idea that willpower is a stock, so that exercising self-control can increase the control cost over the next few periods if periods are short.

The agent's objective function is

$$(1) \quad V(h_n, \mathbf{a}) \equiv E_{\mathbf{a}, h_n} \sum_{\ell=0}^{\infty} \delta^\ell ((1 - \delta)u(y_{n+\ell}, a_{n+\ell}) - \Gamma\Delta(y_{n+\ell}, a_{n+\ell}));$$

⁹This is the same specification of temptation value as in Noor (2007). Depending on the cost function, other interpretations are sometimes possible as well, for example, in the case of linear costs (defined below), the definition of temptation utility is consistent with perfect foresight. But the “naive” interpretation is valid regardless of the cost function.

this is the expected average present value of the per-period utility u minus the discounted sum of control costs.

As in the case of the shorter-run decision problem, we assume the existence of a maximum:

ASSUMPTION 0: $V(h_n, \mathbf{a})$ has a maximum for each n, h_n .

Moreover, because this is a Markov decision problem, there is an optimal plan in which the choice of action depends only on the current state y_n .

The objective function (1) reduces to the linear-costs version of Fudenberg and Levine (2006, 2010) when $\mu = 0$. It is also a special case of the functional forms considered in Noor (2007, 2011); unlike Noor, we assume that the long-run and short-run selves have the same underlying per-period utility function, and we specify how preferences change with the period length.

It is important to note that the term $\Gamma\Delta$ in (1) is not normalized by $1 - \delta$; this is because here we model control cost as an impulse, while we treat utility as a flow.¹⁰ To understand the implications of this, consider the decision to take an action that lowers utility in every period by 1. This action lowers the shorter-run self's value by 1 and so has a foregone value of $\Delta = 1$, independent of period length. Thus its control cost is Γ regardless of period length, which seems like the right conclusion. In contrast, if the $\Gamma\Delta$ terms were multiplied by $(1 - \delta)$, the cost of implementing this action would go to zero as the periods become short. However, an action that lowers utility by 1 for a single period but has no impact on future utilities has foregone value $\Delta = 1 - \delta\mu$ and a cost of $\Gamma(1 - \delta\mu)$, which does become very small when periods are short.

One implication of this model is that long-term commitments will be more attractive than a series of short-term ones: it is cheaper to resist future temptations now than to resist them as they arise. To see the intuition for this, suppose the long-run player undertakes an infinite sequence of actions, each of which lowers utility by 1 in the current period only. The overall cost of this series of actions is $\Gamma(1 - \delta\mu)/(1 - \delta)$, which is greater than the cost Γ of a single action that lowers utility by 1 in every period. Moreover, the cost of committing now to forego 1 util in every period from N on is $\Gamma(\delta\mu)^{N-1}$ and, in particular, is strictly decreasing in N , except in the case $\mu = 0$, where the shorter-run self views all future periods $N > 1$ as equally far away. We illustrate these implications in Example 1 below.

Note that the difference between the long-term and short-term commitments is most extreme in the case $\mu = 0$, where the long-term commitment is no more costly than any of the one-period delays. The difference diminishes as $\mu \rightarrow 1$, holding other parameters constant; although when μ is large, the

¹⁰When we introduce cognitive resources, the control cost will be the reduction in the value of the associated flow.

preferences of the two selves are closely aligned, so that the long-run self has little reason to exert self-control.

Finally, note that sending the time period τ to 0 sends μ to 1, but changes δ as well: lowering utility by 1 forever starting immediately still costs Γ , lowering utility by 1 forever starting at real time $s = N/\tau$ costs $\Gamma \exp(-(\rho + \eta)s)$, and lowering utility by 1 period-by-period immediately costs $\Gamma(\rho + \eta)/\rho$. Thus the difference between the long-term commitment and the series of short-term ones stems not from the period length, but from the greater impatience of the shorter-run self.

Simple Temptations

Several of our examples use as a building block what we call a *simple temptation*, which is a choice between either utility 0 in every period or a flow of $u_g > 0$ that is received for a number of periods N , with $-u_b < 0$ forever after. This choice represents a stereotypical conflict between short-run and long-run preferences that is easy to adapt to varying period lengths and to embed in more complicated decision problems.

The average present values S for the shorter-run self and P for the long-run self of this stream are $S = (1 - (\delta\mu)^N)u_g - (\delta\mu)^N u_b$ and $P = (1 - \delta^N)u_g - \delta^N u_b$. Our interest lies in the case $S > 0$, $P < 0$ so that the shorter-run self would like to take the temptation while the long-run self would prefer to reject it. This conflict arises because the short-run self discounts future periods using discount factor $\delta\mu < \delta$, and will not be present if we send μ to 1 with the other parameters held fixed. However, the reason for interest in the case of μ near 1 is that it corresponds to very short periods. To analyze this case, we fix the calendar length of time T for which the favorable flow lasts, so that $N = T/\tau$. Then when N is an integer, we have $S = (1 - e^{-(\rho+\eta)T})u_g - e^{-(\rho+\eta)T}u_b$, and $P = (1 - e^{-\rho T})u_g - e^{-\rho T}u_b$, independent of τ , even though $\mu = e^{-\eta\tau} \rightarrow 1$.

EXAMPLE 1—Simple and Persistent Temptations With Linear Cost: To begin, consider a choice between accepting and rejecting a simple temptation in the first period, with no other choices to be made. Then the temptation utility is S , so the cost of resisting temptation is $-GS$; thus temptation will be resisted if $P < -GS$. Next note that if the decision can be made at date 1 about whether to accept or reject the temptation in period n , the cost of resisting is $-\Gamma(\delta\mu)^n S$, so the temptation will be resisted if $\delta^n P < -\Gamma(\delta\mu)^n S$ or $P < -\Gamma\mu^n S$; thus as the decision concerns events further in the future, it becomes easier to resist.

Next suppose that the temptation is *persistent*: once the agent consumes the substance, it is gone, but if he does not consume, the substance is still there the next period so that the same choice is faced again in the next period.¹¹ This

¹¹We can formally model this by assuming that there are two states $Y = \{0, 1\}$, where $y_n = 0$ means that the temptation is not available and $y_n = 1$ means that it is. In the state $y_n = 0$, no action

model describes, for example, the temptation to consume a durable good such as a bottle of wine.¹² If the temptation is available, the best possible value for the shorter-run self is S , and the best continuation if the temptation is resisted is $\delta\mu S$. Thus the foregone value from resisting is $\Delta = (1 - \delta\mu)S$, so resisting costs $\Gamma(1 - \delta\mu)S$ each period, and resisting is optimal if $P(1 - \delta) < -\Gamma(1 - \delta\mu)S$. Consequently, the persistent temptation is “harder to resist” than the simple one, and when $(1 - \delta\mu)/(1 - \delta) > |P|/\Gamma S > 1$, the agent would choose to give in to a persistent temptation but resist a simple one. This condition gets increasingly difficult to satisfy as $\mu \rightarrow 1$, holding all other parameters fixed, which corresponds to sending the birth parameter η to 0.

Since the main reason for large μ is that periods are short, it is interesting to study the agent’s choice in the limit of short time periods. Here the agent gives in to the persistent temptation but resists the simple one when $(\rho + \eta)/\rho > |P|/\Gamma S > 1$ or $(\rho + \eta)\Gamma S > \rho|P| > \rho\Gamma S$. This last equation has a simple interpretation: $\rho|P|$ is the value of postponing the negative payoff P for an interval dt , $(\rho + \eta)\Gamma S$ is the flow cost of resisting the persistent temptation, and $\rho\Gamma S$ is the flow or average utility resulting from paying the one-time cost of ΓS to permanently avoid the temptation.

If declining the temptation in period n means that it will not arise again until period $n + \ell$, the situation is intermediate between a persistent temptation ($\ell = 1$) and a simple one ($\ell = \infty$). Then the best continuation value if the option is resisted is $(\delta\mu)^\ell S$, so $\Delta = (1 - (\delta\mu)^\ell)S$, and resisting forever costs

$$\Gamma \sum_{n=0}^{\infty} \delta^{n\ell} (1 - (\delta\mu)^\ell) S = \frac{\Gamma(1 - (\delta\mu)^\ell) S}{1 - \delta^\ell},$$

so resisting is optimal if $|P|/\Gamma S > (1 - (\delta\mu)^\ell)/(1 - \delta^\ell)$. Consequently, resisting is more attractive when the temptation can be avoided for longer and the decision of whether to take at once or resist forever is monotone in ℓ : there is some $\bar{\ell}$ (possibly 0 or infinity) such that the optimum is to take at once if $\ell < \bar{\ell}$ and resist forever if $\ell > \bar{\ell}$.

Intuitively, this is because the shorter-run self is much less concerned about far-off events than the long-run player is, so the gap between the benefit of delay and the cost of buying off the shorter-run player is increasing in the delay length. If the temptation will arrive fairly soon, there is not much point in trying to avoid it by making a commitment now, as the temptation is already

is possible, $A(y_n) = \{0\}$; in the state $y_n = 1$, the space of actions is $A(y_n) = \{0, 1\}$, where 0 means to resist the temptation and 1 means to give in to the temptation. The transition probabilities in state $y_n = 0$ place probability 1 of remaining in that state, $\pi(0|y_n, a_n) = 1$, while in state $y_n = 1$, the transition probability depends on the action taken: $\pi(0|y_n, 1) = 1$ so that if the action is taken, the temptation is off the table, and $\pi(0|y_n, 0) = 0$ so that if the temptation is resisted, it remains for the next period.

¹²We assume here that the consumption option is all or none; perhaps the wine will spoil once opened.

important for the short-run self, so committing to reject it is about as costly as waiting and resisting it when it arrives. On the other hand, it is less costly to commit now to resist temptations that will arrive in the distant future. This is related to Noor's (2007, 2011) point that agents may be tempted by future consumption; the additional structure of our model lets us explain how this effect depends on the real time between the various decisions as opposed to the period length per se.

Next, consider an initial-period choice of whether to accept or reject K simple temptations, the first one at real time 0, the second at real time period T , the third at $2T$, and so forth. If the agent is close to indifferent about whether to take the first temptation, he will strictly prefer to reject the second if that decision could be made in period 1. For this reason, an agent who would accept the simple temptations may choose to reject the "bundle" of them. This is common behavior, as seen, for example, in the experiments of Kirby and Guatsello (2001) on the "bundling" of decisions.¹³ Note that the once-and-for-all decision to decline a simple temptation can be seen as a bundle of all of the "decline today" decisions, and in each case the agent prefers the bundle for the same reason, namely that he is less tempted by future rewards.

EXAMPLE 2—Nonmonotonicity of the Value of Commitment: Here we show how the value of a commitment to avoid temptation can be nonmonotone in η .¹⁴ To do this, we consider a menu choice when some of the choices correspond to future commitments. Specifically, suppose that at time 0 the agent's action is to pick a menu from the list $(\{0\}, \{0, 1\})$, where $\{0\}$ means that no future decision are possible and $\{0, 1\}$ means that at time 1 the agent will face a simple temptation of the form described above. Because the menu $\{0, 1\}$ simply postpones the decision, choosing it will have no foregone value and no control cost, while the foregone value associated with $\{0\}$ depends on the preferences of the shorter-run self. We specialize to the case $u_g = u_b = 1$, $T = 1$, and assume that $\Gamma = 3$ and $\rho = \log(3/2)$. Thus $P = -1/3$ and $S = 1 - (4/3)\exp(-\eta)$; note that S is positive or negative depending on whether η is more or less than $\log(4/3)$. We define $\bar{\eta}$ as the unique solution of $9\exp(-\bar{\eta}) + 6\exp(-2\bar{\eta}) = 1$ and note $\bar{\eta} > \log(4/3)$.¹⁵

¹³See also Chapter 5 of Ainslie (2001). Both Ainslie and Kirby and Guatsello (2001) reported that merely telling subjects they will face the same decision in the future changes choices as well, which the stationary model of this section cannot explain. Note though that Kirby and Guatsello reported a much smaller impact of this "suggested linking" than of actual linking, and that the instructions they used for suggested linking told subjects "the choice you make now is the best indication of how you will choose every time," which may have induced a spurious effect.

¹⁴Section W1 of the Supplemental Material examines a different sort of nonmonotonicity: those who are willing to pay to avoid a temptation can have intermediate utilities from choosing it, as those with very low values may find it easy to resist without commitment, while those with high values will have a correspondingly high control cost for choosing the commitment.

¹⁵This follows from $9\exp(-\log(4/3)) + 6\exp(-2\log(4/3)) = (27/4) + 6\exp(-2\log(3/2)) > 1$.

We will show that for $\eta < \log(4/3)$, the agent is indifferent between $\{0\}$ and $\{0, 1\}$, while for $\log(3/2) < \eta < \bar{\eta}$, the agent strictly prefers $\{0, 1\}$, and for $\eta > \bar{\eta}$, the agent strictly prefers $\{0\}$. The reason for this nonmonotonicity is this: for very small η , the long-run and short-run selves agree on whether to take this temptation, so no self-control is needed to resist it. For larger values of η , resisting requires costly self-control, and when $\eta > \log(3/2)$, the self-control is so costly that when the decision is made at time 1 (that is, “without precommitment”), it is optimal to give in. Moreover, when η is just a little above $\log(3/2)$, it is not optimal to avoid the temptation by choosing the menu $\{0\}$ at time 0, because the time interval between 0 and 1 is short enough that choosing the commitment menu requires substantial self-control. Finally when η is sufficiently large, the short-run self is not tempted by time-1 consumption, so choosing $\{0\}$ is essentially costless.¹⁶

The agent’s choice at time 1 given $\{0, 1\}$ was chosen at time 0:

- If $\eta < \log(4/3)$, then as $S < 0$, there is no self-control problem at time 1; here the optimum is to resist, with payoff 0.
- If $\eta > \log(4/3)$, so that $S > 0$, then S is the temptation value in period 1. Resisting gives payoff $-FS = 4\exp(-\eta) - 3$, so it is optimal to take the temptation if $P > -FS$, that is, $\eta > \log(3/2)$. For $\log(3/2) > \eta \geq \log(4/3)$, it is optimal to resist.

The agent’s choice at time 0:

- If $\eta < \log(4/3)$, then the shorter-run self’s temptation payoff is 0, so the long-run self is indifferent between choosing $\{0, 1\}$ (and subsequently resisting) or choosing $\{0\}$.
- If $\eta > \log(4/3)$, then the shorter-run self’s value is maximized by choosing the menu $\{0, 1\}$ than taking in period 1. The corresponding temptation value in period 0 is $V = \exp(-(\rho + \eta))S = (2/3)(\exp(-\eta) - 4/3 \exp(-2\eta))$. The long-run self’s payoff to $\{0\}$ is thus $-FV = -2(\exp(-\eta) - 4/3 \exp(-2\eta))$. If the menu $\{0, 1\}$ is chosen, it is optimal to take in period 1, resulting in the payoff $\exp(-\rho)P = -2/9$. For $-2(\exp(-\eta) - 4/3 \exp(-2\eta)) > -2/9$ or $\eta > \bar{\eta}$, the long-run self strictly prefers $\{0\}$, and for $\bar{\eta} > \eta > \log(4/3)$, the long-run self strictly prefers $\{0, 1\}$.

EXAMPLE 3—Declining Marginal Interest Rates: Many studies relate agents trade-offs between money today and money at various points in the future; see, for example, the survey in Frederick, Loewenstein, and O’Donoghue (2002) and the early work of Thaler (1981), who asked subjects to state the dollar amount they would require at three different future dates to make them indifferent to receiving \$15 now. Myerson and Green (1995) used a larger set of delays, allowing a better picture of the delay/interest rate trade-off. Subjects

¹⁶If there is a small fee ε for choosing the flexible menu $\{0, 1\}$, then there is a strict failure of nonmonotonicity: for $\eta < \log(4/3)$, the menu $\{0\}$ is strictly preferred, while for $\log(3/2) + \varepsilon < \eta < \bar{\eta} - \varepsilon$, $\{0, 1\}$ is strictly preferred, and for $\eta > \bar{\eta}$, $\{0\}$ is again strictly preferred.

were asked to state how much hypothetical money c_t they would need right now to make them indifferent to receiving a hypothetical \$1,000 after a delay of length t . Using their responses, we can define an incremental interest rate between successive queried times t_i and t_{i-1} as

$$\rho_i = \log(c_{t_i}/c_{t_{i-1}})/(t_i - t_{i-1}).$$

If agents treat cash payments like immediate consumption and the utility of consumption is linear, then in both the quasihyperbolic discounting model and the long-run/shorter-run self model with a fixed period length these incremental interest rates should be time invariant after the first period. From the Myerson and Green data, the rates are¹⁷

Months (t_i)	Incremental Interest Rate	Months (t_i)	Incremental Interest Rate
0.23	132%	36	26.0%
1	82.1%	60	8.0%
6	40.9%	120	9.4%
12	42.7%	300	6.6%

As can be seen, the decline is gradual rather than a one-time decline after the first period.

We now show how nonmyopic short-run selves generate this sort of declining interest rate. We take utility to be linear in consumption and without further loss of generality, we set $u(c) = c$.¹⁸ We compute the amount of consumption c_n that makes the long-run self indifferent between a unit of consumption at

¹⁷Andersen, Harrison, Lau, and Rutström (2008) found evidence of a smaller but still monotone gradual decline of interest rate with delay when real financial incentives are provided in the form of lotteries and adjustments are made for curvature; Benhabib, Bisin, and Schotter (2010) also found evidence of a gradual monotone decline using (small) cash rewards. Since Keren and Roelofsma (1995) already found (in hypothetical experiments) that agents' intertemporal choices are closer to geometric discounting when rewards are stochastic, it is not clear how much of the difference between Andersen et al.'s findings and those of Myerson and Green are due to the fact that subjects were paid.

¹⁸Myerson and Green asked subjects about cash payoffs as opposed to consumption. Both dual-self and quasihyperbolic models need additional structure to explain why subjects (who presumably save) also view cash payoffs as tempting and why the small-stakes risk aversion seen in lab experiments is so much higher than risk aversion for nontrivial fractions of lifetime wealth. Our earlier papers explained this with endogenously determined "mental accounts and showed how this leads to the high level of small-stakes risk aversion seen in the lab. We believe that a similar explanation should be possible with nonmyopic short-run selves and that the mental accounts model can also explain what is known as the magnitude effect: the implicit interest rate seems to shrink with the stakes (Thaler (1981), Green, Myerson, and McFadden (1997), Benhabib, Bisin, and Schotter (2010)). This remains a project for future work. Note also that if utility is a nonlinear function of consumption, then the marginal rate of substitution between two dates depends on the consumption levels at both dates in addition to the time-preference parameters.

time 1 and c_n units in period n ; we then use this to compute effective marginal interest rates on consumption. (In the Supplemental Material, we extend this to the interest rate used at time 1 to discount between any two periods n and ℓ , which is closer to the long-run player's rate of time preference ρ because future consumption is less tempting.) Observe that if the long-run self is indifferent between 1 unit now and c_n units later, then since $\mu < 1$, the initial shorter-run self strictly prefers 1 unit now. Hence the temptation is to consume now, which incurs no control cost, and provides utility 1 for an average present value of $1 - \delta$. The initial short-run self gets an average present value of $1 - \delta\mu$ from consuming at time 1 and gets $(1 - \delta\mu)(\delta\mu)^{n-1}c_n$ from the delayed option, so the control cost of the delayed option is $\Gamma(1 - \delta\mu)(1 - (\delta\mu)^{n-1}c_n)$. Thus the utility of the delayed option is $(1 - \delta)\delta^{n-1}c_n - \Gamma(1 - \delta\mu)(1 - (\delta\mu)^{n-1}c_n)$. Equating the values of the two options determines the consumption level leading to indifference:

$$1 - \delta = (1 - \delta)\delta^{n-1}c_n - \Gamma(1 - \delta\mu)(1 - (\delta\mu)^{n-1}c_n).$$

We can then solve for c_n . This is done in the [Appendix](#); from that computation it follows that as $\mu \rightarrow 1$, we have $c_n \rightarrow 1/\delta^{n-1}$, which is the solution for a single agent without self-control costs. To relate this back to Myerson and Green, the [Appendix](#) also computes the instantaneous interest rate for consumption decisions at real time $t = n\tau$ rate by letting the period length τ go to 0; we show that

$$\lim_{\tau \rightarrow 0} \log\left(\frac{c_{t/\tau+1}}{c_{t/\tau}}\right) / \tau = \rho + \Gamma \frac{\rho + \eta}{\rho} \exp(-\eta t) \eta.$$

In this case, the marginal interest rate, to a good approximation, is equal to the subjective interest rate of the long-run self, plus a term that declines exponentially at rate η . In the case of a shorter-run self who lives exactly one period, that is, $\mu = 0$ or $\eta = \infty$, the marginal interest rate declines after a single period to a constant equal to the subjective interest rate of the long-run self. However, for larger values of μ , we get a more gradual decline, as we see in the data.¹⁹

4. CONVEX COSTS OF SELF-CONTROL

We now consider a simple extension of the model of linear cost of self-control by allowing the cost of self-control to be convex. Specifically, we assume

¹⁹This gradual decline is consistent with hyperbolic (as opposed to quasihyperbolic) discounting. The hyperbolic model is not widely used in economics, perhaps because it is hard to apply in a dynamic setting.

that the objective function is defined by the expected average present value of shorter-run utility net of the control cost,

$$(2) \quad V(h_n, \mathbf{a}) \equiv E_{\mathbf{a}, h_n} \sum_{\ell=0}^{\infty} \delta^\ell ((1 - \delta)u(y_{n+\ell}, a_{n+\ell}) - g(\Delta(y_{n+\ell}, a_{n+\ell}))),$$

where g is a convex function. Allowing g to be convex is important both in light of evidence from the psychology literature and because in the standard dual-self model, convex costs can explain preference “reversals” that arise from failure of the independence axiom, such as the Allais paradox, while linear control costs cannot.²⁰ In addition, convex costs can generate “compromise effect” violations of the axiom of revealed preference, as shown in Fudenberg and Levine (2006).

EXAMPLE 4—Stochastic Temptations: Another implication of convex control costs is that the agent is more likely to resist “stochastic temptations” than certain ones. This is the basis of the explanation of the Allais paradox in Fudenberg and Levine (2010); we give an illustration of the idea here using simple temptations. When faced with a single, and certain, simple temptation, with $S > 0 > P$, it is optimal to choose the temptation if $P > -g(S)$. Now suppose that the agent is faced with the choice between an action that gives probability q of the same simple temptation and complimentary probability $1 - q$ of 0, or resisting, with utility flow 0. Then resisting the temptation has foregone value qS , so resisting is optimal if $-qP < g(qS)$, so when g is convex, it may be optimal to give in to the certain temptation but resist the smaller one.

The same qualitative conclusion extends to the case where the agent learns in period 1 that she will need to make the choice in some future period n : now the temptation value is $(\delta\mu)^{n-1}S$, so the agent resists a lottery that gives probability q of the temptation if $-q\delta^{n-1}P < g(q(\delta\mu)^{n-1}S)$, and it is possible for this inequality to hold for small q but not larger ones. At the same time, though, since there is less of a self-control problem about future decisions, increasing n makes it more likely that the agent resists for all values of q . That is, the model predicts an interaction of the effects of risk and delay: costly self-control in the presence of risk can lead to Allais-type behavior, but since future payoffs are less tempting, individual agents will switch from “paradoxical” to expected-value-maximizing choices when the payoffs are sufficiently far in the

²⁰Linear control costs are consistent with the independence axiom of expected utility theory and so rule out the Allais-type preference reversals. Noor and Takeoka (2010a) weakened the independence axiom in Gul–Pesendorfer axioms to allow nonlinear control costs and then developed axioms that correspond to control costs being convex. Since they work in a two-period model with a single choice of a menu, they do not address the modeling issues we discuss here. Noor and Takeoka (2010b) developed a more restrictive set of axioms that correspond to the cost having the form $\phi(\tilde{U})\Delta$, so that marginal cost is increasing in the size of the largest temptation as opposed to the amount of temptation resisted or value foregone.

future. This is consistent with the data of Baucells and Heukamp (2010): they found that 36% of subjects exhibited preference reversal in a common-ratio Allais paradox, changing from the safer to the expected-value-maximizing choice when the decision is less likely to matter, while only 22% of subjects exhibited this preference reversal when all payoffs were delayed by 3 months.²¹ Note finally that the dependence of the decision on q holds in the case where the agent is initially uncertain whether she will face the certain temptation or the stochastic one: all that matters is that she knows which temptation she is facing at the moment she decides.²²

EXAMPLE 5—Two Tempting Choices: We now consider a variation on Example 4, where instead of a probability of a more or less tempting choice, there is a certainty that two simple temptations will be faced: at both $n_1 = 1$ and $n_2 \geq 1$, the agent has to decide whether to accept or reject a simple temptation with $S > 0, P < 0$.

Our goal is to investigate the sensitivity of the decisions to the timing. Suppose first that $n_2 > 1$, so there is at least some brief delay between the two decisions. Because of the recursive nature of the formulation and the additivity of the utilities, the two decisions are identical. If the option is not taken, utility is 0 and the control cost is $g(S)$. If the option is taken, utility is P and there is no self-control cost, so it is optimal to take at both n_1 and n_2 if $-P < g(S)$ and not to take if $-P > g(S)$. Notice that the solution is the same for any value of $n_2 > n_1$ and for any period length, so it holds in particular if the periods are arbitrarily short.

However, the solution changes if $n_2 = n_1$. In this case the possible actions are not to take, $a = 0$, to take exactly one of the options, $a = 1$, or to take both options, $a = 2$. The temptation is to take both options, so utility is $V(a) = aP - g(2S - aS)$. Then $V(2) - V(1) = 2P - P + g(S) = P + g(S)$ and $V(1) - V(0) = P - g(S) + g(2S)$. When g is strictly convex, $g(2S) - g(S) > g(S)$. If $g(2S) - g(S) > -P > g(S)$, it is optimal to resist each temptation when the options are sequential, but it is not optimal to resist both when they are presented simultaneously. This shows that this model of nonlinear costs is not suited for analyzing decisions that occur in rapid succession. Intuitively, the problem is

²¹Explaining these data with a one-period horizon for the short-run player requires that the period length of the subjects who switched is at least 3 months; this is consistent with the theory but not with data from other experiments. Note also that Weber and Chapman (2005) did not find this effect, although their sample was very small.

²²Convex costs also imply that it is easier to resist smaller deterministic temptations when (as in Example 2) we suppose that utility is a linear function of consumption. As noted in footnote 13, there is evidence that the cost of resisting money temptations does not follow this prediction; we attribute this to the fact that money is not consumed immediately and that its use is governed in part by mental accounts. One advantage of varying the probability of winning and not the prize itself is that the mental accounting for the scale of prize used in lab experiments may well be independent of the probability with which that prize is won.

that the nonlinearity of control costs should “spill over” from one period to the next when time periods are short. The next section extends the model to allow this.

5. WILLPOWER AS A STOCK AND INCREASING MARGINAL COST OF SELF-CONTROL

The reason that control costs are often convex is that self-control can require the use of costly cognitive resources, as argued by Baumeister and various collaborators (for example, Baumeister, Bratslavsky, Muraven, and Tice (1998) and Muraven, Tice, and Baumeister (1998)). This implies that soon after one tempting choice, the marginal cost of another tempting choice will be high; for example, two consecutive decisions a microsecond apart should be about the same as two simultaneous decisions. Thus, to develop a model that is consistent with convex control costs and also robust to the timing of decisions and the granularity of the periods, we need to incorporate the way the willpower stock induces a spillover from one period’s self-control to self-control in the near future.²³ To do this, we develop a generalization of the willpower model of Ozdenoren, Salant, and Silverman (2009).

Specifically, we assume that at the beginning of period n , there is a stock w_n of cognitive resources or willpower available. Note that this is part of the vector y . Foregone value Δ has the same definition as before and, in particular, is not affected by w_n ; the change in the model is that the cost of self-control comes from the fact that it depletes the stock of cognitive resources. Specifically, when $\Delta(y_n, a_n)$ is the foregone value, the end of period stock is $\tilde{w}_n = f(w_n, \Delta(y_n, a_n))$, where $f(w, \Delta)$ is nondecreasing in w and non-increasing in Δ , continuously differentiable in both arguments, and satisfies $f(w, 0) = w$ and $f(w, \Delta) \leq w$. Note that the stock of cognitive resources depends on the action taken only through the foregone value, so actions that maximize the shorter-run self’s value also maximize the end-of-period stock \tilde{w}_n . We do not require that cognitive resources are bounded below; indeed, it will be more convenient to allow them to become arbitrarily negative.

In Ozdenoren, Salant, and Silverman (2009) the stock is depleted, but never replenished. This is a reasonable approximation for the short-duration problem they analyzed, but to adapt the model to longer horizons, we add the possibility that willpower can be replenished. Specifically, we set $w_{n+1} = r(\tilde{w}_n) \geq \tilde{w}_n$, where r is nondecreasing in \tilde{w}_n ; thus for a given w_n , the highest that w_{n+1} can be is $r(w_n)$, and this maximum is attained by actions that set $\Delta = 0$. We assume also that $r(\tilde{w}_n) \leq \bar{w}$, so that there is an upper bound on the stock

²³In the longer term, it is possible that willpower can be built up, that is, that the “willpower technology” can be improved. This introduces a range of issues that our model does not handle well, and we abstract away from it for most of the paper. Section 6 explains some of the complications that arise when willpower can be increased through training.

of cognitive resources. If $r(\tilde{w}_n) = \bar{w}$, then resources are replenished immediately, which is the usual assumption when shorter-run selves live a single period. If $r(\tilde{w}_n) = \tilde{w}_n$, resources are never replenished, as in Ozdenoren, Salant, and Silverman. Self-control costs arise because cognitive resources have alternative uses. Following Ozdenoren, Salant, and Silverman, we assume that an (end-of-period) stock of cognitive resources \tilde{w}_n yields a utility in other uses of $m(y_n, \tilde{w}_n)$, nondecreasing in \tilde{w}_n , and that this is added to the utility from consumption.²⁴ Ozdenoren, Salant, and Silverman view \tilde{w}_n as representing only the stock of willpower, and motivate its assumed value as arising from self-control problems that are not directly modeled. In our earlier work, we provided evidence that cognitive resources matter and that these resources have alternative uses, so we take a broader view of what the uses of these resources might be.

The objective function of the long-run self is then to maximize

$$(3) \quad V(h_n, \mathbf{a}) \equiv E_{\mathbf{a}, h_n} (1 - \delta) \sum_{\ell=0}^{\infty} \delta^\ell (u(y_{n+\ell}, a_{n+\ell}) + m(y_{n+\ell}, f(w_n, \Delta(y_n, a_n))))).$$

Note that the contribution m of the stock of cognitive resources is measured in the same units as utility. Thus if there is a fixed stock \bar{w} of cognitive resources, the stock produces an amount $m(y, \bar{w})$ of utility each period. The function m is assumed to be concave, differentiable, and strictly increasing in w . Recall that with full replenishment, the cost was not normalized by $1 - \delta$ as the benefit is here. We discuss the reason for this difference below.

As in the linear case, we assume LR0, which we repeat here for completeness:

ASSUMPTION LR0: $V(h_n, \mathbf{a})$ has a maximum for each n, h_n .

Note also that once again the long-run self's maximization problem has a solution that depends only on the current state y_n .

In the "cake-eating" problem of Ozdenoren, Salant, and Silverman (2009), there is a cake of fixed size, the only choice is a consumption level a_n that reduces the size of the cake, and the objective of the long-run self is to maximize the discounted value of consumption until the exogenous terminal date, with any remaining cake being useless. They assumed a purely myopic short-run self ($\mu = 0$) and that there is no replenishment ($r(\tilde{w}_n, a_n) = \tilde{w}_n$). Ozdenoren,

²⁴It does not, however, enter into the computation of the temptation utility or the foregone value, as these are a cause of self-control cost, not a consequence. Note also that we do not impose the restriction $\tilde{w}_n \geq 0$ as do Ozdenoren, Salant, and Silverman: we can set $m(y_n, \tilde{w}_n) = -\infty$ if $\tilde{w}_n < 0$ to incorporate that constraint, in which case we modify the differentiability assumptions that we impose below.

Salant, and Silverman used a continuous-time model to specify that the temptation is a fixed upper bound \bar{a} on the consumption rate as long as some cake remains and that the rate of willpower depletion is $\tilde{f}(a_n, w_n)$ for $a_n < \bar{a}$, with \tilde{f} decreasing and strictly convex in a_n . In addition, $m(y_n, \tilde{w}_n) = 0$ until the stock of cake runs out or the time horizon is reached, at which time $m(y_n, \tilde{w}_n) = \bar{m}(\tilde{w}_n)$. In our formulation, $\Delta(y_n, a_n) = u(\bar{a}) - u(a_n)$, so that their model is a special case of ours with $f(w_n, \Delta(y_n, a_n)) = \tilde{f}(u(\bar{a}) - \Delta(y_n, a_n), w_n)$.²⁵ However, their formulation requires that cognitive resource utility is state-dependent. This possibility leads to complications, because it implies that the plan most preferred by the short-run self, which is the plan that has the least temptation, need not minimize the resource cost of self-control. We examine this issue in Section 6, along with the possibility that actions have a direct impact on the evolution of cognitive resources.

ASSUMPTION 4—State-Independence: *m depends on the state only through the stock of willpower; that is, for all $y_n, y'_n, \tilde{w}_n, m(y_n, \tilde{w}_n) = m(y'_n, \tilde{w}_n)$.*

We maintain this assumption until Section 6 and, to lighten notation, we write $m(\tilde{w}_n)$ in place of $m(y_n, \tilde{w}_n)$.

The dual-self model of Fudenberg and Levine (2006, 2010) corresponds to the assumptions (i) $r(\tilde{w}_n) = \bar{w}$, so that replenishment is immediate, and (ii) $\mu = 0$, so the short-run self is completely myopic. Here the temptation value is $\max_a u(y_n, a)$, so $\Delta(y_n, a_n) = \max_a u(y_n, a) - u(y_n, a_n)$, and the benefit derived from cognitive resources in period n is $m(f(\bar{w}, \Delta(y, a_n)))$. We can then define $c(\Delta) = (m(\bar{w}) - m(f(\bar{w}, \Delta)))/(1 - \delta) \geq 0$; since $f(w, 0) = w$, this implies $c(0) = 0$.

Substituting this definition of cost into the objective function, we have

$$V(h_n, \mathbf{a}) \equiv E_{\mathbf{a}, h_n} (1 - \delta) \sum_{\ell=0}^{\infty} \delta^\ell \left(u(y_{n+\ell}, a_{n+\ell}) + m(\bar{w}) - c(\Delta(y_{n+\ell}, a_{n+\ell})) \right),$$

which is equivalent to the objective function in our earlier papers. Note here that neither the function f nor the function m matters on its own: what matters is the composition $m \circ f$, for this is what determines the cost function c .²⁶ There is a similar but more complicated interaction between m and f under partial

²⁵The function f is not constrained to be $w_n - \Delta$. This allows for lower depletion near $w_n = 0$, as in the multiplicative functional form $f(w_n, \Delta(y_n, a_n)) = \Delta(y_n, a_n)w_n$.

²⁶To model the effect of cognitive load (e.g., using short-term memory) on self-control, Fudenberg and Levine (2006) assumed that the control cost depends on the sum of the foregone value and cognitive load; this corresponds to assuming that the benefit derived from cognitive resources in period n is $m(f(\bar{w}, \Delta(y, a_n) + d_n))$, where d_n is the cognitive load in period n .

replenishment, as we see in the next subsection and again in the analysis of Example 9.²⁷

Single Decision Problems

Cognitive resources serve to link the decisions in one period to future control costs and thus subsequent decisions. Several of the examples we have considered so far involve a single decision; in those cases, the resource variable is superfluous. To make this precise, we define what we mean by a *single decision*. Let Y^* be the set of states in which a decision is possible, that is, $y \notin Y^*$ implies $\#A(y) \leq 1$. Then the probability of hitting Y^* from a state $y \in Y^*$ must be zero: once a decision is offered, no further decisions are possible. Notice, though, that if y can occur in different periods, the amount of cognitive resources available for decision making may be different. Suppose that $w_1 = \bar{w}$, so that initially cognitive resources are “topped up.” In this case, we say that resources *start full*.²⁸

For any period n , define the end-of-period resource stocks corresponding to an initial shock of Δ and no subsequent shock by $\tilde{w}_n^n(\Delta) = f(\bar{w}, \Delta)$ and $\tilde{w}_{n+\ell}^n(\Delta) = r(\tilde{w}_{n+\ell-1}^n(\Delta))$ for all $\ell \geq 1$. Then the cost of self-control corresponding to a single shock is

$$g(\Delta_n) = \sum_{\ell=0}^{\infty} \delta^\ell [m(\bar{w}) - m(\tilde{w}_{n+\ell}^n)].$$

The following result is immediate.

THEOREM 3: *If there is a single decision and cognitive resources start full, the maximization problems*

$$E_{a,h_n}(1 - \delta) \sum_{\ell=0}^{\infty} \delta^\ell (u(y_{n+\ell}, a_{n+\ell}) + m(y_{n+\ell}, f(w_n, \Delta(y_n, a_n))))$$

and

$$E_{a,h_n} \sum_{\ell=0}^{\infty} \delta^\ell ((1 - \delta)u(y_{n+\ell}, a_{n+\ell}) - g(\Delta(y_{n+\ell}, a_{n+\ell})))$$

²⁷Note also that cost is normalized differently than m . When there are cognitive resources, the loss due to a single shock is spread out over many periods, so the per-period amount should be in the same units as period utility u . In contrast, when cognitive resources replenish immediately, as implicit in the “cost” formulation, the loss due to a single shock is concentrated in a single period and should be in the same units as the present value of utility.

²⁸The theorem also holds if there is a fixed time at which the decision is possible. That is, if $y_n \in Y^*$ implies $n = n^*$, then we may replace \bar{w} with the fixed amount of cognitive resources w^* available when a decision is possible, as this is invariant to $y \in Y^*$.

have the same set of solutions.

The Linear Case

Next suppose that in addition to conditions (i) and (ii) above, the benefit of cognitive resources is linear in Δ , so that $m(y_n, \tilde{w}_n) = \gamma \tilde{w}_n$, and suppose that resource depletion is linear as well, so that $f(w, \Delta) = w - \Delta$.²⁹ Then the cost function $c(\Delta) = (m(\bar{w}) - m(f(\bar{w}, \Delta)))/(1 - \delta)$ defined above is $c(\Delta) = (1 - \delta)\gamma\bar{w} - (1 - \delta)\gamma(\bar{w} - \Delta) = \gamma(1 - \delta)\Delta$, so the linear model of the previous section, where the cost of self-control is independent of δ , corresponds to scaling the cost by $1/(1 - \delta)$. Intuitively, full replenishment means that all of the cost of self-control is borne in the current period, so if foregone utility reduces the flow benefits of cognitive resources by a proportionate amount, the cost of self-control goes to zero with the period length. Conversely, if the control cost is invariant to the period length and there is full replenishment, the flow cost in a period must become large as the periods get small. This is also true when there are convex costs: the convex cost model of Section 4 can be viewed as a model with full replenishment and linear depletion, where the benefits at \bar{w} are independent of τ , while for smaller stocks, we have $m_\tau(\bar{w} - \Delta) = m(\bar{w}) - g(\Delta)/(1 - \delta)$.

When benefits are linear as well, we have a stronger result: the linear model with partial replenishment is equivalent to the linear model with full replenishment, so that partial replenishment has observable consequences only if at least one of f , r , or m is nonlinear. Specifically, we say the model has *linear replenishment* of resources if $r(\tilde{w}_n) = \tilde{w}_n + \lambda(\bar{w} - \tilde{w}_n)$, where $0 \leq \lambda \leq 1$.

THEOREM 4: *For any model with linear benefit γ , depletion, and replenishment λ , define*

$$\Gamma = \frac{(1 - \delta)\gamma}{1 - \delta(1 - \lambda)}.$$

Then if \mathbf{a} is a solution to the linear model with parameter Γ , then it is a solution to the (λ, γ) model in which actions are independent of w_n , and all such solutions to the (λ, γ) model are solutions to the Γ model.

Theorem 4, which is proven in the [Appendix](#), shows that if depletion, replenishment, and self-control cost are all linear and state-independent, the stock of self-control is irrelevant. Intuitively, with linear costs, all that matters is the average present value of the costs, and not their timing, which is why the stock of

²⁹Note that we define linear depletion to mean that foregone utility is subtracted one-for-one from the stock of resources. In principle, depletion might be linear with a coefficient other than 1, but we can normalize the coefficient to 1 by choosing appropriate units for w .

willpower resources does not matter. Because the theorem maps many (λ, γ) models to the same linear model Γ , it also implies that these models are equivalent in the sense of generating the same decisions. That is, if we change λ to λ' while holding the time period (and thus δ and μ) fixed, and set

$$\gamma' = \gamma \left(\frac{1 - \delta(1 - \lambda')}{1 - \delta(1 - \lambda)} \right),$$

the resulting system (λ', γ') will have the same cost for every self-control decision, even though the time path of the willpower stock in the two models will be different: the agent has the same preferences with resources that replenish quickly and matter a lot or with resources that replenish very slowly and matter little.

Note that it is important for this result that resources are unbounded below; if there is a lower boundary, the model is not linear there and the equivalence with the full replenishment linear model breaks down. Note also that the equivalent linear parameter Γ depends on the replenishment rate: when resources are replenished very quickly ($\lambda = 1$), the cost of self-control is on the order of $(1 - \delta)$ of a single period's utility, while when replenishment is slow, self-control has a long-term cost on the order of 1.

When we vary τ , the length of the period, we want to hold fixed the amount of calendar time required for a given amount of replenishment, so we set $\lambda(\tau) = 1 - \exp(-\kappa\tau)$. This corresponds to assuming that self-control in a given period reduces the stock of willpower at the start of that period. That is, we suppose that when Δ_n is spent in some period n , the stock immediately jumps from w_n to \tilde{w}_n and is then replenished according to the continuous-time differential equation $\dot{w}_t = \kappa(\bar{w} - w_t)$.³⁰ Thus when the period length is τ , we have $w_{n+1} = \bar{w} - \exp(-\kappa\tau)(\bar{w} - \tilde{w}_n)$, so $\lambda(\tau) = 1 - \exp(-\kappa\tau)$. Note that when the period is long, the stock almost completely replenishes. Note also that as $\tau \rightarrow 0$, we have $\lambda(\tau) \rightarrow 0$, $\delta(\tau) \rightarrow 1$.

To characterize the behavior of the linear model as $\tau \rightarrow 0$, we compute the limit of the equivalent marginal cost, which is

$$(4) \quad \Gamma^* \equiv \lim_{\tau \rightarrow 0} \frac{(1 - \exp(-\rho\tau))\gamma}{(1 - \exp(-(\rho + \kappa)\tau))} = \frac{\rho\gamma}{\rho + \kappa}.$$

If κ is very small, so that any reduction in the stock is almost permanent, then $\Gamma^* \approx \gamma$. At the other extreme, when κ is very large, the equivalent cost is near 0. Intuitively, in this case reductions in resources are replaced so quickly

³⁰In the Supplemental Material, we consider alternate specifications of the timing of the reduction in the resources stock caused by self-control, for example, the reduction might happen at the end of the period instead of the start. We show that with linear replenishment, the timing does not matter when periods are small, although, of course, it does matter when periods are longer.

that they are virtually costless, even though the amount of replenishment in a given period goes to 0 with τ .

EXAMPLE 6—Resisting Temptation Within the Linear Case: To illustrate Theorem 4, we reconsider the persistent and delayed temptations of Example 1 in the model with linear benefit, linear depletion, and linear replenishment. This lets us illustrate the continuous-time limit and set the stage for our subsequent analysis of these temptations when the benefit function is concave.

We begin with the case of a persistent temptation, where temptation is present each period unless and until it is accepted. If the benefits are linear, so that $m(w) = \gamma w$, then by Theorem 4, the solution is the same as in the linear case: The agent will resist the temptation if $P(1 - \delta) < -\Gamma(1 - \delta\mu)S$ and accept it when the reverse inequality is satisfied, where $\Gamma = (1 - \delta)\gamma/(1 - \delta(1 - \lambda))$. Using equation (4), we see that when the time period is short, the agent resists the temptation when

$$(5) \quad \rho P < -\gamma\rho S(\rho + \eta)/(\rho + \eta).$$

Here the left-hand side is approximately the gain of postponing P by τ and the right-hand side is the cost of postponement; this is the foregone value of $(\rho + \eta)S$ multiplied by the continuous-time cost parameter Γ^* . Note that when κ is very large compared to the other parameters, the right-hand side is near 0, so it is always better to resist.

To extend this analysis to the case where declining the temptation delays it for a real time s , recall that when declining puts off the temptation for ℓ periods, it is optimal to resist if $P < -\Gamma(1 - (\delta\mu)^\ell)/(1 - \delta^\ell)S$; substituting for Γ yields $P(1 - \delta(1 - \lambda))/(1 - \delta) < -\gamma S(1 - (\delta\mu)^\ell)/(1 - \delta^\ell)$. If we suppose that the delay is s units of real time, assume $\ell = s/\tau$, and send τ to zero while holding s fixed, then it is optimal to resist if $P(\rho + \kappa)/\rho < -\gamma S[(1 - \exp(-(\rho + \eta)s))/(1 - \exp(-\rho s))]$.

In the case $s = \tau$, in which the delay is only a single period, this inequality reduces to $\rho P < -\gamma\rho S[(\rho + \eta)/(\rho + \kappa)]$, which is what we had before. As $s \rightarrow \infty$, the condition reduces to $P < -\gamma\rho S/(\rho + \kappa)$, which is easier to satisfy. More generally, as in the discrete-time model, the decision of whether to take at once or resist forever is monotone in s : there is some \bar{s} (possibly 0 or infinity) such that the optimum is to take at once if $s < \bar{s}$ and resist forever if $s > \bar{s}$. When we reexamine this problem with a concave benefit function, we see that it can be optimal to resist for a finite length of time and then take, and that the optimal time to give in is monotone in the length s of the delay.

6. COGNITIVE RESOURCES, NONLINEARITIES, AND REPLENISHMENT

In the linear model, the stock of willpower does not play a significant role. When there are nonlinearities in f , r , and/or m , the stock of willpower determines the way the cost of foregone utility is allocated between different periods. The main reason for introducing the cognitive resources variable is to

allow for the possibility that (i) the cost of self-control depends on the stock, (ii) the stock does not completely replenish from one period to the next, so that exerting willpower in one period can have a carry-over effect on decisions made soon afterward, and (iii) the agent faces more than one decision, so using self-control in an earlier decision can alter the control cost in a subsequent one. The simplest way to capture this is to suppose that there is no replenishment at all, so that the stock evolves according to $w_{n+1} = f(w_n, \Delta(y_n, a_n))$. This stark assumption is sufficient for demonstrating many of the implications of willpower as a resource that is limited in the short run, but it is not necessary, and many of the same results are obtained provided that replenishment is incomplete.

In general, the dynamic value of cognitive resources given the foregone utility process Δ_n is governed by the depletion equation $\tilde{w}_n = f(w_n, \Delta_n)$, the replenishment equation $w_{n+1} = r(\tilde{w}_n)$, and of course the benefit function $m(y_n, \tilde{w}_n)$. Putting together the depletion and replenishment equations gives the dynamics of cognitive resources $w_{n+1} = r(f(w_n, \Delta_n))$, where $\tilde{w} \geq r(\tilde{w}_n) \geq \tilde{w}_n$ is nondecreasing and $f(w_n, \Delta_n) \leq w_n$, $f(w_n, 0) = w_n$ is nondecreasing in w_n and nonincreasing in Δ_n .

Linear Replenishment

The units in which w_n are measured are arbitrary; by changing them, we change f , r , and m . As we shall see, there is redundancy in these three functions, meaning that we can choose one of them to normalize.

Specifically, let $0 < \lambda < 1$ be a fixed number. We construct a change of units $w' = h^{-1}(w)$ and $\tilde{w}' = h^{-1}(\tilde{w})$ so that the replenishment function, r' that corresponds to w' is linear, with $r'(\tilde{w}') = \tilde{w} - (1 - \lambda)(\tilde{w} - \tilde{w}')$. Given such a function, we may define $f'(w', \Delta) = h^{-1}(f(h(w'), \Delta))$ and $m'(w') = h^{-1}(m(h(w')))$, and the model with the new units and new depletion and benefit functions is the same as the original one.

There are a variety of ways to construct an h function. One simple method is to start with an interval $I_\lambda = (0, \lambda\tilde{w}]$ and a mapping $T(w') \equiv \tilde{w} - \lambda(\tilde{w} - w')$, and then consider the images of the iterated map $T^n(I_\lambda)$ (since T is invertible, we allow negative values of n). Notice that these intervals form a partition of $(-\infty, \tilde{w})$.

Hence for any $w' < \tilde{w}$, there exists a unique integer $n(w')$ (possibly negative) such that $w' \in T^{n(w')}(I_\lambda)$. The interval I_λ for the units w' corresponds to $(0, r(0)]$ in the original units w . Define $h(w') = r^{n(w')}(T^{-n(w')}(w')r(0)/\lambda\tilde{w})$. If $r(0) > 0$ ³¹ and r is strictly increasing for $\tilde{w} < \tilde{w}$, $(0, r(0)]$ is a nonempty interval, so h maps onto $(-\infty, \tilde{w})$. In this case, h is strictly increasing and so is invertible, and by construction $h(\tilde{w} - \lambda(\tilde{w} - w')) = r(h(w'))$. If r is continuous, then

³¹The same construction would work with a different origin, provided there is some w with $r(w) > w$.

h is continuous and $h(0) = 0$; h extends uniquely to a continuous function³² on $(-\infty, \bar{w}]$ by defining $h(\bar{w}) = \bar{w}$.

We summarize this as a theorem.

THEOREM 5: *Suppose that r is continuous and strictly increasing, and that $r(w) > w$ for at least one w . Then the system with replenishment $r'(\tilde{w}') = \bar{w} - (1 - \lambda)(\bar{w} - \tilde{w}')$, depletion $f'(w', \Delta) = h^{-1}(f(h(w'), \Delta))$, and benefit function $m'(w') = h^{-1}(m(h(w')))$ maps strategies to values of the agent's objective function exactly as does the system with replenishment $r(\tilde{w})$, depletion $f(w, \Delta)$, and benefits $m(w)$.*

Notice that the rescaling of units to linearize r is possible only when there is some replenishment ($r(w) > w$) and less than full replenishment (r strictly increasing). When there is full replenishment, we cannot change the units to spread the foregone utility shock over time: as soon as there is nonlinearity, partial replenishment spreads the marginal cost of self-control over time.

Note also that if we start with a system where benefits and depletion are linear, and replenishment is linear with some λ' , then the equivalent system in units of w' for a different value of λ is not linear. This may seem puzzling in light of our observation that Theorem 3 implies an equivalence between linear models with different replenishment rates. However, this is only for average present values, not the stronger sort of equivalence established here, which tracks the moment-by-moment movement of the flow benefit of cognitive resources. The weaker form of equivalence is sufficient when benefits and depletion are linear, but once these functions are allowed to be nonlinear, the stronger sort of equivalence is needed, and this equivalence requires a nonlinear change of units.

Nonlinear Costs and Linear Replenishment

Now we investigate the implications of nonlinear costs when the agent faces multiple decisions, so that self-control in one period can increase the marginal cost of self-control in the next one. To make the computations easier, we pick units so there is linear replenishment, $r(\tilde{w}_n) = \tilde{w}_n + \lambda(\bar{w} - \tilde{w}_n)$. In the examples that follow we will frequently need to compute the average present value of cognitive resources when the stock at the start of period n is some arbitrary w_n and no self-control is used from period n . With linear replenishment and no foregone utility, $w_{n+\ell} = (1 - \lambda^\ell)\bar{w} + \lambda^\ell w_n$. Along this path, the average present value of cognitive resources is

$$(6) \quad M(w_n) = (1 - \delta) \sum_{\ell=0}^{\infty} \delta^\ell m(w_{n+\ell}).$$

³²Notice that h need not be differentiable everywhere: it may have kinks at the boundaries of the intervals.

Recall that to study the effect of varying the period length, we take $\lambda(\tau) = 1 - \exp(-\kappa\tau)$, so for small τ , we have $\lambda(\tau) \approx \kappa\tau$.

EXAMPLE 7—Two Tempting Choices With Linear Depletion and Replenishment: Now we reanalyze the two temptations of Example 5 assuming partial linear replenishment and linear depletion. We show that the agent makes the same decision whether the decisions on the two temptations are made simultaneously or in very rapid succession.

Suppose first that the decisions are made in consecutive periods, so that $n_1 = 1$ and $n_2 = 2$, and to simplify notation, suppose that the initial stock of resources is \bar{w} . The agent has four possible plans:

- *Take both options:* The first option provides direct utility of P , the second option provides δP , and no self-control is used, so overall utility is $(1 + \delta)P + M(\bar{w})$. As $\tau \rightarrow 0$, this converges to $2P + M(\bar{w})$.

- *Take only the second option:* Self-control of S is used in the first period and none thereafter, so $\bar{w}_1 = \bar{w} - S$ and the overall value is

$$(1 - \delta)m(\bar{w} - S) + \delta P + \delta M(\bar{w} - (1 - \lambda)S) = \delta P + M(\bar{w} - S).$$

As $\tau \rightarrow 0$, this converges to $P + M(\bar{w} - S)$.

- *Take only the first option:* No self-control is used in the first period; in the second period, the foregone SR utility (and thus the expenditure of cognitive resources) is S and no self-control is used thereafter, so cognitive resources at the end of the second period are $\bar{w} - S$ and the overall value is

$$\begin{aligned} P + (1 - \delta)m(\bar{w}) + \delta(1 - \delta)m(\bar{w} - S) + \delta^2 M(\bar{w} - (1 - \lambda)S) \\ = P + (1 - \delta)m(\bar{w}) + \delta M(\bar{w} - S). \end{aligned}$$

As $\tau \rightarrow 0$, this converges to $P + M(\bar{w} - S)$.

- *Reject both options:* Self-control is used in both periods, so the value is

$$\begin{aligned} (1 - \delta)m(w_1 - S) + (1 - \delta)\delta m(\bar{w} - (2 - \lambda)S) \\ + \delta^2 M(\bar{w} - (1 - \lambda)(2 - \lambda)S) \\ = (1 - \delta)m(w_1 - S) + \delta M(\bar{w} - (2 - \lambda)S). \end{aligned}$$

As $\tau \rightarrow 0$, this converges to $M(\bar{w} - (2 - \lambda)S)$.

Because the two projects are identical, when the decisions are made simultaneously, there are only three plans to consider:

- *Take both options:* Here the overall value is $2P + M(\bar{w})$. As $\tau \rightarrow 0$, this does not change.

- *Take one option:* Here overall value is $P + (1 - \delta)m(\bar{w} - S) + \delta M(\bar{w} - (1 - \lambda)S)$. As $\tau \rightarrow 0$, this converges to $P + M(\bar{w} - S)$.

- *Reject both options*: Now the overall value is

$$(1 - \delta)m(\bar{w} - 2S) + \delta M(\bar{w} - (1 - \lambda)2S) = M(\bar{w} - 2S).$$

As $\tau \rightarrow 0$, this converges to $M(\bar{w} - 2S)$.

Thus as $\tau \rightarrow 0$, the value for taking both options converges to $2P + M(\bar{w})$, the value for taking either one of the two options converges to $P + M(\bar{w} - S)$, and the value for rejecting both options converges to $M(\bar{w} - 2S)$, regardless of whether the options are sequential or simultaneous, and regardless of which option is taken when only one is taken. Hence simultaneity does not matter in the limit and the optimum is governed by whichever of these three numbers is largest. In the sequential case, if it is optimal to take just one option, the advantage of taking the second option over the first is given by $(1 - \delta)[-P + M(\bar{w} - S) - m(\bar{w})]$. Notice that $m(\bar{w}) = M(\bar{w})$ and that the first option is strictly better if and only if $-P < M(\bar{w}) - M(\bar{w} - S)$. On the other hand, for the first option to be a strict optimum, it must also be strictly better than taking both options, which implies $M(\bar{w}) - M(\bar{w} - S) < -P$. Thus if it is strictly best to take one option, it must be the second option. Intuitively, when the second option occurs, cognitive resources are depleted, so it makes more sense to give in rather than to first give in and then resist.

EXAMPLE 8—Persistent Temptation With Nonlinear Benefits and No Replenishment: To further explore the implications of willpower being a stock that can be depleted over time, we now revisit the persistent temptation of Example 1 in a setting with no replenishment of cognitive resources, linear depletion, and nonlinear benefits. One of the main differences is that with willpower, it may now be optimal to resist a while and then to take the temptation once the marginal benefit of resources becomes sufficiently high. We emphasize that this sort of nonstationary behavior is consistent with perfect foresight and thus need not be interpreted as a sign that agents misperceive their own future intentions.

Because there is no replenishment at all, the stock decreases by $\Delta = (1 - \delta\mu)S$ each time the agent resists; if the agent resists ℓ times before giving in, his value is

$$(1 - \delta) \left[\sum_{n=1}^{\ell-1} \delta^{n-1} m(\bar{w} - n(1 - \delta\mu)S) + \sum_{n=\ell}^{\infty} \delta^{n-1} m(\bar{w} - \ell(1 - \delta\mu)S) \right] + \delta^\ell P.$$

The larger is ℓ , the smaller is the first term and the larger is the second. This implies that a necessary condition for an optimal ℓ is that the value for resisting

$\ell + 1$ times is not larger and that a sufficient condition for $\ell = 0$ optimal is that the value for $\ell = 1$ is lower.

The value for ℓ minus the value for $\ell + 1$ is³³

$$D(\ell) = \delta^\ell [m(\bar{w} - \ell(1 - \delta\mu)S) - m(\bar{w} - (\ell + 1)(1 - \delta\mu)S) + (1 - \delta)P].$$

Observe that since m is concave, $D(\ell)$ is strictly increasing, so it is optimal to take at the first time this expression is positive and never to take if it is always negative.

To characterize the solution for small τ , define

$$\begin{aligned} d(s, \tau) &= D(s/\tau) / (\delta^{(s/\tau)}(1 - \delta)) \\ &= [m(\bar{w} - (s/\tau) \exp(-(\rho + \eta)\tau)S) - m(\bar{w} - ((s/\tau) + 1)(\exp(-(\rho + \eta)\tau)S + \exp(-\rho\tau)P)] \\ &\quad / \exp(-\rho\tau), \end{aligned}$$

where s/τ is the number of periods until real time s . Note that d has the same sign as D , so it can be used to characterize the optimum. Observe also that

$$d(s) = \lim_{\tau \rightarrow 0} d(s, \tau) = m'(\bar{w} - s(\rho + \eta)S)(\rho + \eta)S/\rho + P.$$

Thus if $\lim_{w \rightarrow -\infty} m'(w) < -\rho P/S(\rho + \eta)$, it is optimal never to give in; if $m'(\bar{w}) > -\rho P/(\rho + \eta)S$, it is optimal to give in right away; and if neither corner solution applies, the optimal time to give in is characterized by $m'(\bar{w} - s(\rho + \eta)S) = -\rho P/(\rho + \eta)S$.³⁴

This case of no replenishment is extreme and we will soon revisit this example to allow not only partial linear replenishment, but also a general depletion function. First, though, we want to make a simpler point about the possibility that it is optimal to “wait to commit.”

The case where declining the temptation postpones the decision for a number of periods T is also of interest. Using an argument analogous to that above, it can be shown that the stopping time is increasing in T ; we omit the details.

EXAMPLE 9—Waiting to Commit: Without cognitive resources (or with full replenishment), it is always cheaper to commit now to avoid a future temptation than to do so later when the temptation is more imminent. However,

³³Here we make use of the fact that the two decisions lead to the same time path of resources up to time ℓ .

³⁴A naïve interpretation of these inequalities is that if ρ is large compared to the other terms, the agent resists forever, while if ρ is small, the agent accepts at once, but as both P and S depend on ρ , the comparative statics require more care.

because there is no foregone value associated with delaying the commitment, it is easy to construct examples where it is optimal to postpone a decision if the current stock of resources is quite low. The next example expands on this theme of optimal delay: we now add to Example 8 the possibility of taking the temptation off the table for a cost $F < |P|$, while maintaining the simplifying assumptions of no replenishment and linear depletion. We show that it can be optimal to wait for a while and then pay to commit; once again, a nonstationarity in behavior is consistent with perfect foresight.

Using the argument from Example 8, we see that for small enough τ , it is not optimal to commit immediately if $m'(\bar{w} - F)(\rho + \eta)S/\rho < F$. A sufficient condition is

$$(7) \quad m'(\bar{w} + P) < \rho F/(\rho + \eta)S.$$

Suppose that $m'(-\infty)(\rho + \eta)S/\rho + P < 0$, so that in the absence of the possibility of commitment, it is optimal never to give in, resulting in at least the value $m(\bar{w}) + P$. Taking the temptation off the table in the first period gives value

$$\lim_{\tau \rightarrow 0} (1 - \delta) \left[\sum_{n=1}^{\infty} \delta^{n-1} m(\bar{w} - (1 - \delta\mu)S) - F \right] - F = m(\bar{w} - F) - F.$$

Hence a sufficient condition for committing is $m(\bar{w}) - m(\bar{w} - F) + P + F < 0$. Observe that

$$\begin{aligned} m(\bar{w}) - m(\bar{w} - F) + F + P &\leq m'(\bar{w} - F)F + F + P \\ &\leq m'(\bar{w} + P)F + F + P. \end{aligned}$$

Thus a sufficient condition for committing is $m'(\bar{w} + P) < |P|/F - 1$. This inequality together with (7) are sufficient for it to be optimal to wait a while then to commit. To see that both conditions can be satisfied simultaneously, take $F = |P|/2$. Then the sufficient condition for committing is $m'(\bar{w} + P) < 1$, while (7) becomes $m'(\bar{w} + P) < \rho|P|/2(\rho + \eta)S$. In other words, if the marginal benefit of cognitive resources is low when resources are higher than $\bar{w} + P$ and if F is small but not too small, then it pays to use cognitive resources for self-control until the marginal benefit of cognitive resources is sufficiently high and then to commit to taking the temptation off the table.

Houser et al. (2010) provided a very suggestive experiment that indicates that delay in commitment may occur in practice: subjects were paid to complete certain tasks while they had access to a web browser, but they could pay a cost to remove web access. The experimental instructions did not specify when and whether opportunities for commitment might occur in the future, so it is not clear what subjects believed and whether the perfect foresight analysis applies. Thus although delay was observed, we cannot be certain from these

experiments whether it is the type of delay predicted by this model. We hope that future experiments will shed more light on the type of delay that can occur with cognitive resource depletion.

EXAMPLE 10—A Persistent Temptation With Nonlinear Benefits, Partial Linear Replenishment, and General Depletion: Our final variation on the persistent temptation problem drops the option to commit, and examines the interplay between the benefit function m and the depletion function f . To begin, note that regardless of the form of the benefit and depletion functions, if there is full replenishment, the problem is stationary, so it is never optimal to wait for a while and then take. Defining the cost of self-control to be $g(\Delta) = (1 - \delta)(m(\bar{w}) - m_\tau(f(\bar{w}, \Delta)))$ (as in Section 4), we see that the policy of taking at once gives payoff $P + m(\bar{w})$ and resisting forever gives payoff $m(f(\bar{w}, (1 - \delta\mu)S))$, so resisting forever is optimal if

$$P < m(f(\bar{w}, (1 - \delta\mu)S)) - m(\bar{w}) = -g((1 - \delta\mu)S)/(1 - \delta).$$

Next suppose that $m(w) = \gamma w$ for $w \geq 0$ and $m(w) = -\infty$ for $w < 0$. If $P < -\gamma(1 - \delta\mu)S$, then as in our earlier analysis, the optimum is to take the temptation immediately. If $(1 - \delta(1 - \lambda))P > -\gamma(1 - \delta\mu)S$ the solution is to resist forever if permanent resistance is consistent with resources remaining nonnegative; this is the case if $1 - (1 - \delta\mu)S/\bar{w} > \lambda$. Otherwise, if $(1 - \delta(1 - \lambda))P > -\gamma(1 - \delta\mu)S$ and $(1 - (1 - \delta\mu)S)/\bar{w} < \lambda$, the solution is to resist until one more period of resistance would “exhaust the stock” (that is, make $\tilde{w} < 0$) and then to give in to the temptation.

Now consider the general case, normalizing to have linear replenishment. We assume that m and f are twice continuously differentiable, and that

$$-\frac{dm}{dw}((1 - A)\bar{w} + Aw)\frac{\partial f}{\partial \Delta}(w, 0)$$

(which is positive) is decreasing in w ; we call this *increasing marginal cost of self-control*. Note that the marginal cost of control is increasing if m is strictly concave and f is linear. Equivalently, we can write this condition in terms of second derivatives as

$$\begin{aligned} &-\frac{d^2m}{dw^2}((1 - A)\bar{w} + Aw)\frac{\partial f}{\partial \Delta}(w, 0) \\ &-\frac{dm}{dw}((1 - A)\bar{w} + Aw)\frac{\partial^2 f}{\partial \Delta \partial w}(w, 0) < 0. \end{aligned}$$

If m is concave, this says that the cross-partial derivative of f should not be “too negative.” If the cross-partial is strongly positive, then m need not be concave.

PROPOSITION 6: *Suppose there is increasing marginal cost of self-control and there is strictly partial linear replenishment, $0 < \kappa$.*

(a) There is $\underline{\tau} > 0$ such that if $\tau < \underline{\tau}$, there are $|\bar{P}_\tau| > |\underline{P}_\tau| > 0$ such that it is optimal to resist forever if $|P| > \bar{P}_\tau$, it is optimal to resist until period $\infty > \hat{\ell} > 1$ and then take if $|\bar{P}_\tau| > |P| > |\underline{P}_\tau|$, and it is optimal to take immediately if $|\underline{P}_\tau| > P$.

(b) Let $\bar{P}_0, \underline{P}_0$ denote the limits of $\bar{P}_\tau, \underline{P}_\tau$ as $\tau \rightarrow 0$. Let W_s be the solution to the differential equation

$$\dot{W}_t = \kappa(\bar{w} - W_t) + \frac{\partial f}{\partial \Delta}(W_t, 0)(\rho + \eta)S$$

and let W_∞ be the solution to

$$0 = \kappa(\bar{w} - W_\infty) + \frac{\partial f}{\partial \Delta}(W_\infty, 0)(\rho + \eta)S.$$

Then

$$|\bar{P}_0| = -(\rho + \eta)S \int_0^\infty e^{-(\rho+\kappa)t} m'(W_\infty) \frac{\partial f}{\partial \Delta}(W_\infty, 0) dt,$$

$$|\underline{P}_0| = -(\rho + \eta)S \int_0^\infty e^{-(\rho+\kappa)t} m'(\bar{w}) \frac{\partial f}{\partial \Delta}(\bar{w}, 0) dt,$$

and if $|\bar{P}_0| > |P| > |\underline{P}_0|$, then $\hat{s} = \lim_{\tau \rightarrow 0} \tau \hat{\ell}$ is finite and strictly positive, and is determined by

$$|P| = (\rho + \eta)S \int_0^\infty e^{-(\rho+\kappa)t} m'(\bar{w} - e^{-\kappa t} W_s) \frac{\partial f}{\partial \Delta}(W_s, 0) dt.$$

REMARK 1: One way to read this result is that the agent's choice depends on the magnitude of P , but recall that $P = (1 - \delta^T)u_g - \delta^T u_b$ and $S = (1 - (\delta\mu)^T)u_g - (\delta\mu)^T u_b$, so changing P implies changes in S and/or in δ and μ (or ρ and η in the continuous-time formulation), and any of these other changes will also matter for the decision.

REMARK 2: To better understand the formulas given above, note that when depletion and benefits are both linear,

$$|\bar{P}_0| = |\underline{P}_0| = (\rho + \eta)\gamma S \int_0^\infty e^{-(\rho+\kappa)t} dt = \frac{(\rho + \eta)}{(\rho + \kappa)} \gamma S,$$

which is the same as the condition for the critical value of P given in equation (4).

REMARK 3: To illustrate the fact that concavity of the optimization can come from any of the three functions f , m , and r , consider the case where f and m are linear, and r is piecewise linear:

$$\begin{cases} r(\tilde{w}_n) = \tilde{w}_n + \lambda_1(\bar{w} - \tilde{w}_n) & \text{for } \tilde{w}_n \in [w^*, \bar{w}], \\ \tilde{w}_n + \lambda_2(\bar{w} - \tilde{w}_n) & \text{for } \tilde{w}_n < w^*, \end{cases}$$

where $\lambda_1 = 1 - \exp(-\kappa_1\tau)$, $\lambda_2 = 1 - \exp(-\kappa_2\tau)$, and

$$-\gamma S(\rho + \eta)/(\rho + \kappa_2) < P < -\gamma S(\rho + \eta)/(\rho + \kappa_1).$$

Then if the replenishment rate was fixed at κ_1 , the agent would always resist, while if it was fixed at κ_2 , the agent would take the temptation at once. We claim that the short-time-period solution with the piecewise linear replenishment function is to resist until resources fall to w^* . To see why, first consider the agent's problem when the resource level is w^* . Resisting forever gives exactly the same payoff as when the replenishment rate is fixed at κ_2 , and taking gives a higher payoff than with replenishment fixed at κ_2 , so since taking gives a higher payoff here than with $\kappa = \kappa_2$, the agent takes. When $w > w^*$, the gain from resisting for a short interval and then taking the temptation, instead of taking it now, is exactly as in the case $\kappa = \kappa_1$, so the agent resists; since this causes resources to decrease, the agent will resist until resources fall to w^* .

The proof of Proposition 6 is provided in the [Appendix](#), but the intuition is simple. We first show that because of the increasing marginal cost of self-control and because resisting temptation lowers the stock next period, the gain to waiting one more period is monotone in the number of periods ℓ that the temptation has been resisted. Thus, if P is small enough (sufficiently bad) relative to all the other parameters, it is optimal to wait forever; if P is close enough to 0, it is optimal to take at once; and for intermediate P , it is optimal to wait a while and then take. For an arbitrary length τ of the time period, this intermediate region may be empty, but when τ is very small, the concavity assumption ensures that it is nonempty.

EXAMPLE 11—State-Dependent Marginal Cost: We assume full replenishment of willpower each period, so the stock of willpower is constant and thus irrelevant. As in the case of constant marginal cost, we assume both linear resource depletion and linear value of cognitive resources. However, we drop the assumption that the marginal benefit of cognitive resources is constant and instead let them depend on the state. Specifically, we assume $m(y_n, \tilde{w}_n) = \Gamma(y_n)\tilde{w}_n$ and, more specifically, that the marginal benefit $\Gamma(y_n)$ in period 1 is $\bar{\Gamma} > 0$ while from period 2, it is either $\bar{\Gamma}$ or 0, depending on the first period choice.

In period 1, there is a choice of whether or not to pay a cost F ; think of it as spending time learning self-control, perhaps with the aid of a counselor or

religious or spiritual advisor. If the cost is paid, then there is no problem of self-control at all in future periods, that is, $F(y_n) = 0$; if the cost is not paid, the marginal benefit remains equal to \bar{F} .

In period 2, the agent can decide whether to take or to resist a simple temptation, with shorter-run player value S and direct value P for the long-run player, with $S > 0 > P$ and $P < -\bar{F}S$, so that if the agent does not pay in period 1, it will be optimal to take in period 2.

Now we examine the decision in period 1. The future best value for the short-run self is $\delta\mu S$, regardless of whether F is paid today or not. Thus the temptation utility is $\delta\mu S$ and the utility the SR associates with “pay” is $\delta\mu S - F(1 - \delta\mu)$, so the foregone value of pay is $F(1 - \delta\mu)$ and the self-control cost for this action is $\bar{F}F(1 - \delta\mu)$. Hence it is optimal in the reduced form problem to pay whenever $(\bar{F}(1 - \delta\mu) + (1 - \delta))F < -\delta P$.

In contrast, if paying F today makes taking tomorrow impossible, the foregone value of pay is $\delta\mu S + F(1 - \delta\mu)$, so for some parameters (such as μ close to 1), the commitment will not be optimal even though the arguably equivalent “training” action would be. The difference between commitment and lowering control costs is a consequence of our assumption that the shorter-run selves are strategically naive, so that the shorter-run player is unconcerned by any action that leaves the feasible set unchanged. Models with nonnaive shorter-run players may also be of interest, but they are much more complicated.³⁵

To make this example simple, we kept the stock of willpower constant and assumed that the first-period action had a direct effect on the cost of self-control in the second period. Similar effects could be obtained if we allowed the replenishment function r to depend on the action as well as on the end-of-period willpower, and let the benefits of cognitive resources be slightly concave (so that the cost is slightly convex). Specifically, suppose that acting in the first period increases the willpower stock from 1 to $1 + w^*$ and that the benefit of cognitive resources w is w^α for some $\alpha \in (0, 1)$. Then if the agent does not act in the first period, the cost of resisting second-period temptation is $1^\alpha - (1 - \delta\mu S)^\alpha \geq \alpha\delta\mu S$, while the cost if the agent acts goes to 0 with w^* .

In the Supplemental Material, we discuss some of the complications that arise when the evolution of resources can be state-dependent and can depend directly on actions, not only on the foregone value. The Supplemental Material also shows that when we rule out the kind of endogenous changes in willpower explored in Example 11, we can show that the actions that maximize the long-run self’s objective function correspond to the equilibrium of a game in which

³⁵This example suggests that nonnaivete is necessary to capture St. Augustine’s request “give me chastity and continence, but not yet.”

a benevolent but patient long-run self faces a sequence of shorter-run selves who live for a random length of time. In this game, decisions are made by the shorter-run selves, but the long-run self can alter the preferences of the shorter-run selves by undertaking “self-control” actions that, in general, lower the utility of the shorter-run selves.

7. CONCLUSION

Allowing shorter-run selves to live more than a single period provides a natural way to capture the way preferences change as the “period” becomes shorter. This lets us explain why commitments to avoid far-off temptations are less costly and more attractive than commitments to avoid more imminent temptations, and lets us explain the subjective interest rates decline with delay; it also provides a natural parameterization of the effect of the length of the time interval between potential decision nodes. This is important because the concept of a discrete-time period in these decision problems is simply a convenient construction; the real-time length of delay is what should matter for decision making.

When the marginal cost of self-control is constant, the agent’s decision problem is not affected by the timing of when self-control costs are incurred and there is no need for the model to track the stock of the agent’s cognitive resources. As we saw, the model with linear replenishment, benefits, and depletion is equivalent to the “state-free” model. However, when nonlinearities matter, so does the timing of self-control decisions and costs, and the willpower stock provides a way to model the “spillover” from one period’s self-control to future control costs. Tracking the stock of willpower allows simultaneous decisions to be about the same as almost simultaneous ones, and lets us explain why agents may choose to “resist and then give in” and “wait to commit.”

We explored some, but far from all, of the many possible ways to model these nonlinearities, and there is ample scope for future work on this. In particular, we have looked for plausible properties, such as insensitivity to minor changes in timing; it would be useful to compile these properties in a set of clear and readily interpretable axioms to better understand the universe of models that satisfy them. Also, it would be good to extend the qualitative analysis here by exploring the extent to which we can find, for each individual agent, a stable constellation of preference parameters that fits that agent’s quantitative behavior across a range of problems. This was done to a limited extent in Fudenberg and Levine (2010) for the model where shorter-run selves live a single period, although the calibration was for the median subject across a number of different experiments as opposed to individual subjects observed in many different treatments. However, several of the experiments studied there are better fit by allowing shorter-run selves to have random lifetimes. Finally, while our analysis here has presumed that the cost function c (in the first part of the paper) or the functions governing the evolution of cognitive resources (in the second

part) are fixed and constant, it is straightforward to allow them to have random variation, as long as this variation is exogenous. For example, each period t the cost of foregoing Δ units of value could be Γ_t , where the Γ_t are independent and identically distributed according to some known distribution.

APPENDIX

EXAMPLE 3—Continued Derivation of Marginal Interest Rates: From the indifference condition in the text, we may compute

$$\frac{1 - \delta + \Gamma(1 - \delta\mu)}{(1 - \delta)\delta^{n-1} + \Gamma(1 - \delta\mu)(\delta\mu)^{n-1}} = c_n,$$

so setting $n = t/\tau$ and taking the limit $\tau \rightarrow 0$ shows that

$$\begin{aligned} \text{MR}_t &= \lim_{\tau \rightarrow 0} \log \left(\frac{c_{(t/\tau)+1}}{c_{(t/\tau)}} \right) / \tau \\ &= \lim_{\tau \rightarrow 0} \log \left(\frac{\frac{1 - \delta + \Gamma(1 - \delta\mu)}{(1 - \delta)\delta^{t/\tau} + \Gamma(1 - \delta\mu)(\delta\mu)^{t/\tau}}}{\frac{1 - \delta + \Gamma(1 - \delta\mu)}{(1 - \delta)\delta^{(t/\tau)-1} + \Gamma(1 - \delta\mu)(\delta\mu)^{(t/\tau)-1}}} \right) / \tau \\ &= \lim_{\tau \rightarrow 0} \log \left(\frac{1}{\delta} \frac{(1 - \delta) + \Gamma(1 - \delta\mu)\mu^{(t/\tau)-1}}{(1 - \delta) + \Gamma(1 - \delta\mu)\mu^{t/\tau}} \right) / \tau \\ &= \rho + \Gamma \lim_{\tau \rightarrow 0} (1 - \delta\mu)(\exp(-\eta t)(\mu^{-1} - 1)) / (1 - \delta)\tau \\ &= \rho + \Gamma(\rho + \eta) \exp(-\eta t)\eta / \rho, \end{aligned}$$

where MR denotes the marginal interest rate.

PROOF OF THEOREM 4: Recall that $\tilde{w}_n = w_n - \Delta_n$. With linear replenishment, $w_{n+1} = \tilde{w}_n + \lambda(\tilde{w} - \tilde{w}_n) = w_n - \Delta_n + \lambda(\tilde{w} - w_n + \Delta_n)$. Define the *willpower deficit* $D_n = \tilde{w} - w_n$. Then

$$D_{n+1} = (1 - \lambda)D_n + (1 - \lambda)\Delta_n = (1 - \lambda)^{n+1}D_0 + \sum_{n'=0}^n (1 - \lambda)^{n'+1}\Delta_{n'}.$$

Recall that the average value of cognitive resources in the linear case is

$$M = (1 - \delta) \sum_{n=0}^{\infty} \delta^n \gamma w_n.$$

It follows that the total value of cognitive resources is

$$\begin{aligned}
M/(1-\delta) &= \gamma \left[\bar{w} - \sum_{n=0}^{\infty} \delta^n D_n \right] \\
&= \gamma \left[\bar{w} - \sum_{n=0}^{\infty} \delta^n \left((1-\lambda)^n D_0 \right. \right. \\
&\quad \left. \left. + \sum_{n'=1}^n (1-\lambda)^{n-n'} \Delta_{n'} + \Delta_{n+1} \right) \right] \\
&= \gamma \left[\bar{w} - \sum_{n=0}^{\infty} \delta^n \left((1-\lambda)^n D_0 + \sum_{n'=0}^n (1-\lambda)^{n-n'} \Delta_{n'} \right) \right] \\
&= \gamma \left[\bar{w} - \sum_{n=0}^{\infty} \delta^n (1-\lambda)^n D_0 - \sum_{n=0}^{\infty} \delta^n \sum_{n'=0}^n (1-\lambda)^{n-n'} \Delta_{n'} \right] \\
&= \gamma \left[\bar{w} - \sum_{n=0}^{\infty} \delta^n (1-\lambda)^n D_0 \right. \\
&\quad \left. - \sum_{n'=0}^{\infty} \delta^{n'} \Delta_{n'} \sum_{n=n'}^{\infty} (\delta(1-\lambda))^{n-n'} \right] \\
&= \gamma \left[\bar{w} - \sum_{n=0}^{\infty} \delta^n (1-\lambda)^n D_0 - \frac{1}{1-\delta(1-\lambda)} \sum_{n'=0}^{\infty} \delta^{n'} \Delta_{n'} \right].
\end{aligned}$$

Hence if we define $\Gamma = (1-\delta)\gamma/(1-\delta(1-\lambda))$, we see the equivalence to the simple linear model without replenishment. *Q.E.D.*

PROOF OF PROPOSITION 6: Suppose the agent resists for ℓ periods and then gives in. Let w_n^ℓ be the corresponding time path of cognitive resources. Note that this is a weakly decreasing function of ℓ , strictly decreasing for $n > \ell$. The resulting average value is

$$(1-\delta) \left(\sum_{n=0}^{\ell-1} \delta^n m(f(w_{n+1}^\ell, (1-\delta\mu))S) + \left(\sum_{n=\ell}^{\infty} \delta^n m(w_{n+1}^\ell) \right) \right) + \delta^\ell P,$$

and the larger is ℓ , the smaller is the first term and the larger is the second term (recall that P is negative). This implies that a necessary condition for an optimal ℓ is that the value for $\ell + 1$ is no larger and that a sufficient condition for $\ell = 1$ optimal is that the value for $\ell + 1$ is lower. Let us look at the value at

ℓ minus the value at $\ell + 1$:

$$\begin{aligned}
D(\ell) &= (1 - \delta) \left(\sum_{n=0}^{\ell-1} \delta^n m(f(w_{n+1}^\ell, (1 - \delta\mu)S)) \right. \\
&\quad \left. + \left(\sum_{n=\ell}^{\infty} \delta^n m(w_{\ell+1}^\ell) \right) \right) + \delta^\ell P \\
&\quad - (1 - \delta) \left(\sum_{n=0}^{\ell} \delta^n m(f(w_{n+1}^{\ell+1}, (1 - \delta\mu)S)) \right. \\
&\quad \left. + \left(\sum_{n=\ell+1}^{\infty} \delta^n m(w_{n+1}^{\ell+1}) \right) \right) - \delta^{\ell+1} P.
\end{aligned}$$

Note that for $n \leq \ell$, we have $w_n^\ell = w_n^{\ell+1}$, so we can write this difference as

$$\begin{aligned}
D(\ell) &= (1 - \delta) \delta^\ell \left[(m(w_{\ell+1}^\ell) - m(f(w_{\ell+1}^\ell, (1 - \delta\mu)S))) \right. \\
&\quad \left. + \sum_{n=\ell+1}^{\infty} \delta^{n-\ell} (m(w_{n+1}^\ell) - m(w_{n+1}^{\ell+1})) + P \right].
\end{aligned}$$

Observe that because there is partial replenishment, $w_{\ell+n}^\ell$ strictly decreases in ℓ .

We now use the assumption of increasing cost of self-control to conclude that there is a $\bar{\tau}$ such that for $\tau < \bar{\tau}$, each individual term in $D(\ell)$ is strictly increasing in ℓ and hence that $D(\ell)$ is strictly increasing. The first term, $m(w_{\ell+1}^\ell) - m(f(w_{\ell+1}^\ell, (1 - \delta\mu)S)) = m(f(w_{\ell+1}^\ell, 0)) - m(f(w_{\ell+1}^\ell, (1 - \delta\mu)S))$, strictly decreases in $w_{\ell+1}^\ell$ from increasing marginal cost of self-control at $A = 1$ and the fact that $(1 - \delta\mu)S \rightarrow 0$ as $\tau \rightarrow 0$. Since $w_{\ell+1}^\ell$ decreases in ℓ , these differences increase.

For the terms in the sum, since n runs from $\ell + 1$ to ∞ , the arguments w_{n+1}^ℓ and $w_{n+1}^{\ell+1}$ have the form $w_{\ell+\ell'}^\ell$ and $w_{\ell+\ell'}^{\ell+1}$, and the former decrease with ℓ . The individual terms have the form $m(w_{\ell+\ell'}^\ell) - m(w_{\ell+\ell'}^{\ell+1})$ with $w_{\ell+1}^{\ell+1} = f(w_{\ell+1}^\ell, (1 - \delta\mu)S)$ and $w_{\ell+\ell'}^{\ell+1} = \bar{w} - (1 - \lambda)^{\ell'-1}(\bar{w} - w_{\ell+1}^{\ell+1})$. Putting this together, we have

$$\begin{aligned}
&m(w_{\ell+\ell'}^\ell) - m(w_{\ell+\ell'}^{\ell+1}) \\
&= m(\bar{w} - (1 - \lambda)^{\ell'-1}(\bar{w} - w_{\ell+1}^\ell)) \\
&\quad - m(\bar{w} - (1 - \lambda)^{\ell'-1}(\bar{w} - f(w_{\ell+1}^\ell, (1 - \delta\mu)S))) \\
&= m(A\bar{w} + (1 - A)w_{\ell+1}^\ell) - m(A\bar{w} + (1 - A)f(w_{\ell+1}^\ell, (1 - \delta\mu)S))
\end{aligned}$$

$$= m(A\bar{w} + (1 - A)f(w_{\ell+1}^\ell, 0)) \\ - m(A\bar{w} + (1 - A)f(w_{\ell+1}^\ell, (1 - \delta\mu)S)),$$

where $A = (1 - \lambda)^{\ell'-1}$. When τ is small enough, the increasing marginal cost of self-control implies that this expression is decreasing in $w_{\ell+1}^\ell$ and so increasing in ℓ when τ is small enough.

Because w_n is bounded below by the steady state, D is bounded above as a function of ℓ . If P is large enough in absolute value (it is negative) given all the other parameters, then this expression is negative for all ℓ and it is optimal to wait forever, let \bar{P}_τ be the smallest such P in absolute value. If P is small enough in absolute value, this expression is positive for all S and it is optimal to take immediately, let \underline{P}_τ be the largest such P in absolute value.

Next we assume that τ is small, and show that $|\bar{P}_\tau| > |\underline{P}_\tau|$. Observe that

$$w_{\ell+\ell'}^\ell = \bar{w} - (1 - \lambda)^{\ell'-1}(\bar{w} - w_{\ell+1}^\ell) = \bar{w} - (1 - \lambda)^{\ell'-1}(\bar{w} - w_{\ell+1}^\ell), \\ w_{\ell+\ell'}^{\ell+1} = \bar{w} - (1 - \lambda)^{\ell'-1}(\bar{w} - w_{\ell+1}^{\ell+1}) \\ = \bar{w} - (1 - \lambda)^{\ell'-1}(\bar{w} - f(w_{\ell+1}^\ell, (1 - \delta\mu)S)).$$

Let $d(s, \tau) = D(s/\tau)/\delta^{s/\tau}(1 - \delta)$ with $W_s^t \equiv w_{s/\tau}^t$. Then

$$d(s, \tau) = (m(W_{s+\tau}^s) - m(f(W_{s+\tau}^s, (1 - \delta\mu)S))) \\ + \sum_{n=2}^{\infty} e^{-\rho n} (m(W_{s+n\tau}^s) - m(W_{s+n\tau}^{s+\tau})) + P,$$

where

$$W_{s+n\tau}^s - W_{s+n\tau}^{s+\tau} = (1 - \lambda)^{n-1} [W_{s+\tau}^s - f(W_{s+\tau}^s, (1 - \delta\mu)S)]$$

and

$$W_{s+n\tau}^s = \bar{w} - (1 - \lambda)^{n-1}(\bar{w} - W_{s+\tau}^s).$$

The first term of d converges to zero as $\tau \rightarrow 0$, and since m and f are differentiable, the sum converges to

$$d(s) \equiv (\rho + \eta)S \int_0^\infty e^{-(\rho+\kappa)t} m'(\bar{w} - e^{-\kappa t}(\bar{w} - W_s)) \frac{\partial f}{\partial \Delta}(W_s, 0) dt,$$

where W_s is the solution to the differential equation

$$\dot{W}_t = \kappa(\bar{w} - W_t) + \frac{\partial f}{\partial \Delta}(W_t, 0)(\rho + \eta)S$$

with initial condition $W_0 = \bar{w}$. Thus we have $d(s) \equiv \lim_{\tau \rightarrow 0} d(s, \tau)$. Recall that D is strictly increasing, and that $\hat{\ell} = 1$ is optimal if and only if $D(1) \geq 0$. As $\tau \rightarrow 0$, this is equivalent to

$$\begin{aligned} d(0) &= -(\rho + \eta)S \int_0^\infty e^{-(\rho+\kappa)t} m'(\bar{w}) \frac{\partial f}{\partial \Delta}(\bar{w}, 0) dt + P \\ &= -\frac{(\rho + \eta)m'(\bar{w})}{(\rho + \kappa)} \frac{\partial f}{\partial \Delta}(\bar{w}, 0) \geq 0. \end{aligned}$$

Similarly, $\hat{\ell} = \infty$ is optimal if and only if $\lim_{\ell \rightarrow \infty} D(\ell) \leq 0$, and so is optimal when $d(\infty) \leq 0$.

Finally, resisting for a while and then taking, that is, $1 < \hat{\ell} < \infty$, is optimal if and only if $D(\hat{\ell} - 1) \leq 0$, $D(\hat{\ell}) \geq 0$, hence $d(\hat{s}) = 0$. This gives the characterization of the optimum in the proposition. Finally, the assumption that the marginal cost of self-control is increasing implies $d(s)$ is strictly increasing, so

$$\begin{aligned} |P_\tau| &= -(\rho + \eta)S \int_0^\infty e^{-(\rho+\kappa)t} m'(\bar{w}) \frac{\partial f}{\partial \Delta}(\bar{w}, 0) dt \\ &< -(\rho + \eta)S \int_0^\infty e^{-(\rho+\kappa)t} m'(W_\infty) \frac{\partial f}{\partial \Delta}(W_\infty, 0) dt = |\bar{P}_0| \end{aligned}$$

and hence $|P_\tau| < |\bar{P}_\tau|$ must hold for all sufficiently small τ .

Q.E.D.

REFERENCES

- AINSLIE, G. (2001): *Breakdown of Will*. Cambridge: Cambridge University Press. [2,12]
- ANDERSEN, S., G. HARRISON, M. LAU, AND E. RUTSTRÖM (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583–618. [14]
- BAUCELLS, M., AND F. H. HEUKAMP (2010): "Common Ratio Using Delay," *Theory and Decision*, 68, 149–158. [17]
- BAUMEISTER, R. F., E. BRATSLAVSKY, M. MURAVEN, AND D. M. TICE (1998): "Ego Depletion: Is the Active Self a Limited Resource?" *Journal of Personality and Social Psychology*, 74, 1252–1265. [18]
- BECKER, G., AND K. MURPHY (1988): "A Theory of Rational Addiction," *Journal of Political Economy*, 96, 675–700. [6]
- BENHABIB, J., A. BISIN, AND A. SCHOTTER (2010): "Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs," *Games and Economic Behavior*, 69, 205–223. [14]
- BROCAS, J., AND I. CARRILO (2008): "The Brain as a Hierarchical Organization," *American Economic Review*, 98, 1312–1346. [5]
- CHATTERJEE, K., AND V. KRISHNA (2009): "A Dual Self Representation for Stochastic Temptation," *American Economic Journal: Microeconomics*, 1, 148–167. [5]
- DASGUPTA, P., AND E. MASKIN (2005): "Uncertainty and Hyperbolic Discounting," *American Economic Review*, 95, 1290–1299. [4]
- DEKEL, E., B. LIPMAN, AND A. RUSTICHINI (2009): "Temptation-Driven Preferences," *Review of Economic Studies*, 76, 937–971. [5]

- EPPER, T., H. FEHR-DUDA, AND A. BRUHIN (2009): "Uncertainty Breeds Decreasing Impatience: The Role of Risk Preferences in Time Discounting," Mimeo. [4]
- FREDERICK, S., G. LOEWENSTEIN, AND T. O'DONOGHUE (2002): "Time Discounting and Time Preference: A Critical Review," *Journal of Economic Literature*, 40 (2), 351–401. [13]
- FUDENBERG, D., AND D. K. LEVINE (2006): "A Dual Self Model of Impulse Control," *American Economic Review*, 96, 1449–1476. [7,9,16,20]
- (2010): "Risk, Delay, and Self Control," *AEJ: Microeconomics* (forthcoming). [4,7,9,16,20,35]
- (2012): "Supplement to 'Timing and Self-Control'," *Econometrica Supplemental Material*, 80, http://www.econometricsociety.org/ecta/Supmat/9655_extensions.pdf. [2]
- GABAIX, X., AND D. LAIBSON (2002): "The 6D Bias and the Equity Premium Puzzle," NBER Macroeconomics Annual, NBER. [6]
- GREEN, L., J. MYERSON, AND E. MCFADDEN (1997): "Rate of Temporal Discounting Decreases With Amount of Reward," *Memory and Cognition*, 25, 715–723. [14]
- GUL, F., AND W. PESENDORFER (2001): "Temptation and Self Control," *Econometrica*, 69, 1403–1436. [5]
- (2004): "Self Control and the Theory of Consumption," *Econometrica*, 72, 110–158. [5]
- HALEVY, Y. (2008): "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect," *American Economic Review*, 98, 233–265. [4]
- HOUSER, D., D. SCHUNK, J. WINTER, AND E. XIAO (2010): "Temptation and Commitment in Laboratory," Mimeo. [4,30]
- KEREN, G., AND P. ROELOFSMA (1995): "Immediacy and Certainty in Intertemporal Choice," *Organizational Behavior and Human Decision Making*, 63, 297–297. [4,14]
- KIRBY, K., AND B. GUATSELLO (2001): "Making Choices in Anticipation of Similar Future Choices Can Increase Self Control," *Journal of Experimental Psychology: Applied*, 7, 154–164. [2,12]
- LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 112, 443–477. [1]
- MURAVEN, M., M. D. TICE, AND R. BAUMEISTER (1998): "Self-Control as a Limited Resource," *Journal of Personality and Social Psychology*, 74, 774–789. [3,18]
- MYERSON, J., AND L. GREEN (1995): "Discounting of Delayed Rewards: Models of Individual Choice," *Journal of the Experimental Analysis of Behavior*, 64, 263–276. [2,13]
- NOOR, J. (2007): "Commitment and Self Control," *Journal of Economic Theory*, 135, 1–34. [2,4,7-9,12]
- (2011): "Temptation and Revealed Preference," *Econometrica*, 79, 601–644. [2,4,9,12]
- NOOR, J., AND N. TAKEOKA (2010a): "Uphill Self Control," *Theoretical Economics*, 5, 127–158. [5,16]
- (2010b): "Temptation by Menu," Mimeo. [5,16]
- O'DONOGHUE, T., AND M. RABIN (1999): "Doing It Now or Later," *American Economic Review*, 89, 103–124. [1]
- OZDENOREN, E., S. SALANT, AND D. SILVERMAN (2009): "Willpower and the Optimal Control of Visceral Urges," *Journal of the European Economic Association* (forthcoming). [3,18,19]
- PHELPS, E. S., AND R. A. POLLAK (1968): "On Second-Best National Saving and Game-Equilibrium Growth," *Review of Economic Studies*, 35, 185–199. [1]
- STROTZ, R. H. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 23, 165–180. [1]
- THALER, R. (1981): "Some Empirical Evidence on Dynamic Inconsistency," *Economics Letters*, 8, 201–207. [13,14]
- WEBER, B., AND G. CHAPMAN (2005): "The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect?" *Organizational Behavior and Human Decision Processes*, 96, 104–118. [17]

Dept. of Economics, Harvard University, 310 Littauer Center, Cambridge, MA 02138, U.S.A.; dfudenberg@harvard.edu

and

Dept. of Economics, Washington University in St. Louis, St. Louis, MO 63130-4899, U.S.A.; david@dklevine.com.

Manuscript received November, 2010; final revision received May, 2011.

SUPPLEMENT TO “TIMING AND SELF-CONTROL”
(*Econometrica*, Vol. 80, No. 1, January 2012, 1–42)

BY DREW FUDENBERG AND DAVID K. LEVINE

W1. AN EXAMPLE MOTIVATED BY DELLAVIGNA ET AL.

FIRST WE CONSIDER an example motivated by DellaVigna et al.’s (2010) experiment on door-to-door charitable fund-raising that illustrates a form of non-monotonicity. DellaVigna et al. (2010) provided homeowners with the option to avoid a fundraiser by using a “Do Not Disturb” check box and found that those who chose avoidance were concentrated among people who donate less when avoidance is not possible. As we explain here, this example is consistent with our model: whether the agent is willing to pay to avoid a temptation can be a nonmonotone function of the temptation’s short-run utility. Intuitively, those who are willing to pay to avoid a temptation can have intermediate utilities from choosing it, as those with very low values may find it easy to resist without commitment, while those with high values will have a correspondingly high control cost for choosing the commitment.

EXAMPLE W1—Door-to-Door Sales: This example has two decisions: First, whether to avoid a tempting opportunity and, second, whether to give in to temptation if it was not avoided. (For concreteness, think of the avoidance activity as avoiding a door-to-door salesman.) As we will see, costly self-control leads to a nonmonotonicity: if the temptation is very high or very low, then the opportunity will not be avoided, but it may be avoided for intermediate levels of temptation. The intuition is that when the opportunity is very good, there is little conflict between the long-run self and shorter-run self, so the opportunity should be taken advantage of and not avoided. When the opportunity is very bad, the shorter-run self will not indulge much, so it is not worth paying a fixed cost for avoidance. However, in the intermediate case, there may be a more severe conflict between long-run self and shorter-run self, so the long-run self may choose to commit so as to avoid the temptation.

The example is very simple and stylized. In period 1, a cost $F \geq 0$ may be paid or not; think of this as not being at home when the salesman calls. If the cost is paid, the utility in all subsequent periods is 0. If the cost is not paid, then in period 2, a decision must be made on whether to purchase from the salesman. If the purchase is made, the utility in period 2 is B ; otherwise it is zero. In period 3, if the purchase was made, it must be paid for, resulting in a disutility of -1 .

To solve the model, we first compute temptation values in each period and state, and then compute the agent’s objective function. We then solve the various inequalities to see when each action is best.

We begin by computing the temptation value in the last period in which action is possible, namely in period 2 when the avoidance cost has not been paid. Here the shorter-run self's average present value from doing nothing is 0 and that of purchasing is $(1 - \delta\mu)(B - \delta\mu)$, so $\bar{U}_2 = (1 - \delta\mu) \max\{0, B - \delta\mu\}$. In the initial state, if F is chosen, the shorter-run self's value is $-F(1 - \delta\mu)$, while if it is not chosen, the shorter-run player value is $\delta\mu\bar{U}_2$. If $B - \delta\mu < 0$, then also $B - \delta < 0$, so in period 2 the optimum is not to purchase, which incurs no cost of self-control.

Now we suppose $B - \delta\mu > 0$ and compute the agent's decision. Resisting temptation in period 2 will cost $\Gamma(1 - \delta\mu)(B - \delta\mu)$, so the purchase will be made when $(1 - \delta)(B - \delta) \geq -\Gamma(1 - \delta\mu)(B - \delta\mu)$. If avoidance is chosen, the shorter-run self pays the avoidance cost of $-(1 - \delta\mu)F$ in the first period. Since the temptation value is $\delta\mu\bar{U}_2$, the average present value of avoidance is

$$\begin{aligned} & -(1 - \delta)F - \Gamma(F(1 - \delta\mu) + \delta\mu\bar{U}_2) \\ & = -(1 - \delta)F - \Gamma(F(1 - \delta\mu) + \delta\mu(1 - \delta\mu) \max\{0, B - \delta\mu\}) \end{aligned}$$

and avoidance is optimal if this is higher than the discounted average value of long-run player utility in period 2, which is $\delta \max\{(1 - \delta)(B - \delta), -\Gamma(1 - \delta\mu)(B - \delta\mu)\}$.

Denote the strategy of not paying the avoidance cost and not purchasing as \mathbf{a}^0 , of not paying the avoidance cost and purchasing as \mathbf{a}^1 , and of paying the avoidance cost as \mathbf{a}^F . We have the following characterization of the optimal decision rule:

PROPOSITION W1: *Set*

$$F^* = \frac{\Gamma\delta^2(1 - \delta\mu)(1 - \delta)(1 - \mu)^2}{(1 - \delta + \Gamma(1 - \delta\mu))^2}.$$

If $F \geq F^*$, then \mathbf{a}^0 is optimal for

$$B \leq \delta \frac{1 - \delta + \Gamma(1 - \delta\mu)\mu}{1 - \delta + \Gamma(1 - \delta\mu)} \equiv B^*$$

and \mathbf{a}^1 is optimal if $B \geq B^*$. If $F \leq F^*$, then

$$\begin{aligned} \bar{B} & \equiv \delta \frac{1 - \delta + \Gamma(1 - \delta\mu)\mu^2}{1 - \delta + \Gamma(1 - \delta\mu)\mu} - \frac{1 - \delta + \Gamma(1 - \delta\mu)}{\delta(1 - \delta + \Gamma(1 - \delta\mu)\mu)} F \\ & \geq B^* \geq \delta\mu + \frac{1 - \delta + \Gamma(1 - \delta\mu)}{\Gamma(1 - \delta\mu)\delta(1 - \mu)} F \equiv \underline{B} \end{aligned}$$

and \mathbf{a}^0 is optimal for $B \leq \underline{B}$, \mathbf{a}^F is optimal for $\underline{B} \leq B \leq \bar{B}$, and \mathbf{a}^1 is optimal for $B \geq \bar{B}$.

PROOF: (i) As worked out above, the payoff to F is

$$\begin{aligned} & -(1 - \delta)F - \gamma(F(1 - \delta\mu) + \delta\mu\bar{U}_2) \\ & = -(1 - \delta)F - \gamma(F(1 - \delta\mu) + \delta\mu(1 - \delta\mu)\max\{0, B - \delta\mu\}). \end{aligned}$$

(ii) If \mathbf{a}^0 (do not avoid, do not purchase) is chosen, the direct utility is 0 and the reduced form utility is the temptation cost incurred in the second period:

$$V(\mathbf{a}^0) = -\Gamma\delta(1 - \delta\mu)\max\{0, B - \delta\mu\}.$$

(iii) If \mathbf{a}^1 is chosen, the direct utility is $(1 - \delta)(\delta B - \delta^2)$, while the cost of self-control is in period 2 and is $-\Gamma(1 - \delta\mu)\min\{0, B - \delta\mu\}$, as self-control is needed only when the shorter-run player does not want to purchase. Thus

$$V(\mathbf{a}^1) = (1 - \delta)\delta(B - \delta) + \Gamma\delta(1 - \delta\mu)\min\{0, B - \delta\mu\}.$$

If $B \leq \delta\mu$, then the optimum is not to purchase and there is no temptation cost; here it is also not optimal to avoid in the first period, and the optimum is \mathbf{a}^0 . Next suppose that $B > \delta\mu$ and consider the period-2 choice, assuming the avoidance cost was not paid. If the purchase is not made, the average value from period 2 on is $-\Gamma(1 - \delta\mu)(B - \delta\mu)$, while if it is made, the average value is $(1 - \delta)(B - \delta)$. So the optimum is not to purchase when

$$B \leq \delta \frac{1 - \delta + \Gamma(1 - \delta\mu)\mu}{1 - \delta + \Gamma(1 - \delta\mu)} \equiv B^*.$$

Next observe that since $B > \delta\mu$, the present value of utility from avoiding is given by

$$-(1 - \delta)F - \Gamma(F(1 - \delta\mu) + \delta\mu(1 - \delta\mu)(B - \delta\mu)).$$

Then $V(\mathbf{a}^F) \geq V(\mathbf{a}^0)$ if and only

$$B \geq \delta\mu + \frac{1 - \delta + \Gamma(1 - \delta\mu)}{\Gamma(1 - \delta\mu)\delta(1 - \mu)}F \equiv \underline{B}.$$

Since $B - \delta\mu > 0$, this implies there is a range of sufficiently small F where \mathbf{a}^F is better and a range of F so large that \mathbf{a}^0 is better.

Finally, $V(\mathbf{a}^F) \geq V(\mathbf{a}^1)$ if

$$\begin{aligned} & -(1 - \delta)F - \Gamma(F(1 - \delta\mu) + \delta\mu(1 - \delta\mu)(B - \delta\mu)) \\ & \geq \delta(1 - \delta)(B - \delta) \end{aligned}$$

or

$$B \leq \delta \frac{1 - \delta + \Gamma(1 - \delta\mu)\mu^2}{1 - \delta + \Gamma(1 - \delta\mu)\mu} - \frac{1 - \delta + \Gamma(1 - \delta\mu)}{\delta(1 - \delta + \Gamma(1 - \delta\mu)\mu)}F \equiv \bar{B}.$$

We conclude that \mathbf{a}^F is best when

$$\begin{aligned} & \delta \frac{1 - \delta + \Gamma(1 - \delta\mu)\mu^2}{1 - \delta + \Gamma(1 - \delta\mu)\mu} - \frac{1 - \delta + \Gamma(1 - \delta\mu)}{\delta(1 - \delta + \Gamma(1 - \delta\mu)\mu)} F \\ & \geq B \geq \delta\mu + \frac{1 - \delta + \Gamma(1 - \delta\mu)}{\Gamma(1 - \delta\mu)\delta(1 - \mu)} F. \end{aligned}$$

Straightforward algebra shows that there is a nonempty interval of B where \mathbf{a}^F is best when

$$F \leq \frac{\Gamma(1 - \delta\mu)(1 - \delta)\delta^2(1 - \mu)^2}{(1 - \delta + \Gamma(1 - \delta\mu))^2} = F^*.$$

If $F > F^*$, it is not optimal to use \mathbf{a}^F ; in this case, the optimum is determined from the condition for $V(\mathbf{a}^0) \geq V(\mathbf{a}^1)$ above. If $F \leq F^*$ and if $B \leq \underline{B}$, then $V(\mathbf{a}^F) \leq V(\mathbf{a}^0)$ and $V(\mathbf{a}^1) \leq V(\mathbf{a}^0)$, so \mathbf{a}^0 is optimal; if $\underline{B} \leq B \leq \bar{B}$, then $V(\mathbf{a}^F) \geq V(\mathbf{a}^0)$ and $V(\mathbf{a}^F) \geq V(\mathbf{a}^1)$, so \mathbf{a}^F is optimal; if $B \geq \bar{B}$, then $V(\mathbf{a}^F) \leq V(\mathbf{a}^1)$ and $V(\mathbf{a}^1) \geq V(\mathbf{a}^0)$, so \mathbf{a}^1 is optimal. Finally note that $\delta\mu \leq \underline{B}$, B^* , so that the case $B \leq \delta\mu$ where \mathbf{a}^0 is optimal is included in this result. *Q.E.D.*

Note that the right-hand side inequality in $\underline{B} \leq B \leq \bar{B}$ gets harder to satisfy as $\mu \rightarrow 1$ or as $\Gamma \rightarrow 0$. In the former case, the interests of the shorter-run self are nearly aligned with those of the long-run self, while in the second case, the shorter-run self defers to the wishes of the long-run self. In either case, paying F is just an expensive way to not buy. Paying a small F is attractive as $\mu \rightarrow 0$, as here the first SR self is not very tempted by the second-period outcome, so it is cheap to get him to agree to a commitment that will probably bind on the next self.

DellaVigna et al. (2010) found that if an option to avoid the fundraiser is available, about a quarter of people make use of it, and that if the option is made cheaper by providing a “Do Not Disturb” check box, nearly a third of people choose to avoid the salesperson. If we imagine that without checking the box there is a cost of avoiding, then this is as our model predicts: the lower the cost of avoidance, the more people will choose it. As noted above, DellaVigna et al. found that those who chose avoidance were concentrated among people who donate less when avoidance is not possible. Whether this is the case in our model depends on the distribution of B . If the lowest value of B/\bar{B} in the population is greater than or equal to 1 and the highest value of B/\bar{B} also exceeds 1, then all those who would not contribute when avoidance is not possible ($F = \infty$) will choose avoidance, while only some of those who would contribute choose avoidance; this is what DellaVigna et al. found. On the other hand, if the highest value of B/\bar{B} in the population is less than or equal to 1 while the lowest value of B/\bar{B} is below 1, our model predicts the opposite result. A more elaborate experiment could vary the value of B and

the cutoffs more systematically—for example, in the flier describing the visit, indicating that a level of matching funds are available (\$3 to the charity for every \$1 you donate, for example). This would make it possible to test for the nonmonotonicity in B that the model predicts.

W2. STATE DEPENDENCE

In the text, we defined the value of cognitive resources to be *state-independent* if it depends on the state only through the stock of willpower; in a slight abuse of notation, we write this as $m(y_n, \tilde{w}_n) = m(\tilde{w}_n)$. State-independent resource valuation implies that the action most favored by the shorter-run self maximizes the utility of cognitive resources. To see this, define

$$M(h_n, \mathbf{a}) \equiv E_{\mathbf{a}, h_n}^k \sum_{\ell=0}^{\infty} (\delta\mu)^\ell (1 - \delta\mu) m(\tilde{w}_{n+\ell})$$

and note that the value on the right-hand side is independent of k . If each period's action is chosen to maximize the value $U^k(h_n, \mathbf{a})$ of the current shorter-run self, the foregone value each period is 0. This implies that the level of resources at each period is as high as possible given the initial value; with state-independent resource valuation, any action plan \mathbf{a} that leads to this highest possible path for \tilde{w} also maximizes the flow of benefits $m(\tilde{w}_n)$ in the strong sense that no other action plan leads to a higher value of m in any period along any history. As a consequence, any action plan that maximizes shorter-run utility in each period on each history also maximizes M .

THEOREM W2: *With state-independent resource valuation, $\arg \max_{\mathbf{a}} U(h_n, \mathbf{a}) = \arg \max_{\mathbf{a}} M(h_n, \mathbf{a})$.*

W3. THE GAME BETWEEN LONG-RUN AND SHORTER-RUN SELVES

Here we show that the optimization problem in the text can be identified with the outcome of a game between the long-run self and a sequence of shorter-run selves. To do this, we introduce an augmented state variable Y_k that is defined in any period n in which a new shorter-run self is born and includes along with y_n the value of n as well as available cognitive resources; that is, $Y_k = (y_n, w_n, n)$. Notice that any strategy \mathbf{a} that maps histories to actions induces a well defined stochastic kernel $\Pi(\mathbf{a}, Y_k)[dY_{k+1}]$ based on the original stochastic kernel and the laws of motion for cognitive resources.

In the game formulation, the actions are taken by the shorter-run selves, and the long-run self chooses self-control actions that influence the preferences of the shorter-run self; the control cost we specified in the text will now correspond to a reduction in the utility function as opposed to an additional term.

Each shorter-run self can be thought of as choosing an \mathbf{a} : Although this contains irrelevant information such as how the shorter-run self will behave after he “dies,” we will ignore this in computing the shorter-run self’s payoff. Following Fudenberg and Levine (2006), we assume that before the shorter-run self moves, the long-run self chooses a *self-control action* $e \in E$. It is convenient to take $E = 0 \cup \mathbf{A}$; we then interpret $e \in \mathbf{A}$ as the “suggested action” and $e = 0$ as “no recommendation” or “no self-control.” In the game formulation, cognitive resources follow the exact same equations of motion as in the reduced form model and depend only on the action actually taken by the short-run self and not on the self-control action e . We consider a sequence of stage games between the long-run self and the k th shorter-run self. The k th stage game consists of a choice of self-control action e by the long-run self and a response \mathbf{a} by the shorter-run self. The utility of the k th shorter-run self has the form $u(Y_k, e, \mathbf{a})$, which we specify below.

Histories in this game are sequences of augmented states Y_k along with the chosen actions e_k, \mathbf{a}_k . A strategy from the long-run self is a map \mathbf{e} from the previous history to a self-control action, and a strategy for the k th shorter-run self is a map \mathbf{a}_k from the previous history and choice of the long-run self to an action. The vector of strategies for all shorter-run players is denoted $\bar{\mathbf{a}}$. We define the conditional expectation operator $E_{\mathbf{e}, \mathbf{a}, Y_k}$ given the strategies \mathbf{e}, \mathbf{a} and state Y_k .

Fudenberg and Levine (2006) specified the procedure for deriving a utility function from an underlying objective function and a “cost of self-control” function. We mimic that procedure here to show that the equilibria of the game are equivalent to those of a particular optimization problem that we define below; we then show that the solution to this optimization problem is the same as the solution to the optimization problem in the text.

To do so, we first define $E_{\mathbf{e}, \mathbf{a}, Y_n}^k$ to be the conditional expectation when k is alive. Write¹

$$\tilde{U}(Y_k, \mathbf{a}) \equiv (1 - \delta) E_{0, \mathbf{a}, Y_k}^k \sum_{n=0}^{\infty} (\delta \mu)^{n-k} u(a_{k+n}^S, y_{k+n}).$$

Parallel to the definition of M in Section W2, define

$$\tilde{M}(Y_k, \mathbf{a}) \equiv (1 - \delta) E_{0, \mathbf{a}, Y_k}^k \sum_{n=0}^{\infty} (\delta \mu)^{n-k} m(y_{k+n}, \tilde{w}_{k+n}),$$

¹Since the right-hand side of this equality does not depend on \mathbf{e} , we write the expectation conditional on $\mathbf{e} = 0$ to facilitate later steps.

where again the right-hand side does not depend on \mathbf{e} . Following Fudenberg and Levine (2006), we now define the SR objective function in the game:

$$\tilde{u}(Y_k, e, \mathbf{a}) = \begin{cases} \tilde{U}(Y_k, \mathbf{a}) + \max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a}'), & e = 0, \\ \tilde{U}(Y_k, e) + \tilde{M}(Y_k, e) - \|e - \mathbf{a}\|, & \tilde{U}(Y_k, \mathbf{a}) + \tilde{M}(Y_k, \mathbf{a}) \geq \tilde{U}(Y_k, e) + \tilde{M}(Y_k, e), \\ \tilde{U}(Y_k, \mathbf{a}) + \tilde{M}(Y_k, \mathbf{a}) - \|e - \mathbf{a}\|, & \tilde{U}(Y_k, \mathbf{a}) + \tilde{M}(Y_k, \mathbf{a}) < \tilde{U}(Y_k, e) + \tilde{M}(Y_k, e). \end{cases}$$

Notice that with this objective function the shorter-run self cares about resources, but except when the utility $\tilde{U}(Y_k, \mathbf{a}) + \tilde{M}(Y_k, \mathbf{a})$ from the chosen action is smaller than the utility $\tilde{U}(Y_k, \mathbf{e}) + \tilde{M}(Y_k, \mathbf{e})$ from the suggested action, the shorter-run self views those resources as being outside of his control.

The long-run self is completely benevolent: her payoff in the game is the discounted sum of shorter-run self utilities

$$\tilde{V}(Y_1, \mathbf{e}, \bar{\mathbf{a}}) \equiv (1 - \delta) E_{\mathbf{e}, \bar{\mathbf{a}}, Y_1} \sum_{k=0}^{\infty} \delta^k \tilde{u}(Y_k, e_k, \mathbf{a}_k).$$

We assume that $\tilde{U}(Y_k, \mathbf{a})$, $\tilde{M}(Y_k, \mathbf{a})$ are continuous in \mathbf{a} and define the cost of self-control to be

$$\begin{aligned} C(Y_k, \mathbf{a}) &\equiv \tilde{u}(Y_k, 0, \mathbf{a}) - \max_{e | \tilde{u}(Y_k, e, \mathbf{a}) \geq \tilde{u}(Y_k, e, \cdot)} \tilde{u}(Y_k, e, \mathbf{a}) \\ &= \begin{cases} 0, & \mathbf{a} \in \arg \max_{\mathbf{a}'} \tilde{u}(Y_k, 0, \mathbf{a}), \\ \max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a}') - \tilde{M}(Y_k, \mathbf{a}), & \mathbf{a} \notin \arg \max_{\mathbf{a}'} \tilde{u}(Y_k, 0, \mathbf{a}), \end{cases} \end{aligned}$$

which has the property that $C(Y_k, \mathbf{a}) \geq 0$ and $C(Y_k, \mathbf{a}) = 0$ if and only if $\mathbf{a} \in \arg \max_{\mathbf{a}'} \tilde{u}(Y_k, 0, \mathbf{a}')$.

As in Fudenberg and Levine (2006), we now consider equilibria in which the shorter-run selves optimize following every history, and the long-run player anticipates this reaction and plays like a Stackelberg leader. This is designed to capture what we imagine is the strategic naivete of the shorter-run self: with one-period lifetimes for the shorter-run players, this Stackelberg equilibrium is equivalent to subgame-perfect equilibrium in which the long-run player moves first against each shorter-run player and is equivalent to the weaker concept of SR-perfect Nash equilibrium defined in Fudenberg and Levine (2006). If we assume that the long-run player can choose a self-control action e_k that is observed by shorter-run self k before choosing plan \mathbf{a}_k , SR-perfect Nash equilibrium has the same implication here. However, the assumption that e_k is chosen once and for all at the beginning of the life of shorter-run self k is stronger

when the shorter-run self lives multiple periods. First, the self-control action changes the preferences of the shorter-run self over many periods. Second, the self-control action cannot be “changed” as long as the particular shorter-run self is alive. As we note below, if the long-run self were unable to commit for the life of the shorter-run self, then there would be a nontrivial strategic interaction between the two.

As is the case in which the shorter-run self lives only for a single period, the expectations of the shorter-run self about play by the long-run self do not matter, because the long-run self has already moved. For this reason, the situation does not correspond to a repeated game (which it would in the absence of the commitment assumption). Moreover, the case for subgame perfection may be stronger here than it is in general, as when the long-run self can commit, the predictions of subgame perfections are less sensitive to changes in the information structure.

Fudenberg and Levine (2006) defined a SR-perfect Nash equilibrium profile to be *equivalent* to a solution to the reduced form optimization problem of maximizing

$$E_{\mathbf{a}, Y_1} \sum_{k=0}^{\infty} \delta^k (u(Y_k, 0, \mathbf{a}) - C(Y_k, \mathbf{a}))$$

if the reduced strategy induced from the shorter-run players’ strategy profile is a solution to the optimization problem. Conversely, if there exists a SR-perfect Nash equilibrium profile with this property for a particular solution to the optimization problem, we say that this solution of the reduced form optimization problem is equivalent to the SR-perfect Nash equilibrium profile. Provided that $\tilde{U}(Y_k, \mathbf{a})$ and $\tilde{M}(Y_k, \mathbf{a})$ are continuous in \mathbf{a} , the conditions of Fudenberg and Levine’s (2006) Theorem 1 are satisfied, so this equivalence does indeed hold.²

THEOREM W3: *If $\tilde{U}(Y_k, \mathbf{a})$ and $\tilde{M}(Y_k, \mathbf{a})$ are continuous in \mathbf{a} , then SR-perfect Nash equilibria are equivalent to solutions to the reduced form optimization problem.*

We now wish to relate solutions to the optimization problem equivalent to SR-perfect Nash equilibria

$$(*) \quad E_{\mathbf{a}, Y_1} \sum_{k=0}^{\infty} \delta^k (u(Y_k, 0, \mathbf{a}) - C(Y_k, \mathbf{a}))$$

²These conditions are costly and unlimited self-control, limited indifference, and continuity.

to those of

$$(**) \quad E_{\mathbf{a}, Y_1} \sum_{n=0}^{\infty} \delta^n ((1 - \delta)u(a_n, y_n) + m(y_n, \tilde{w}_n)),$$

the agent's objective function that we used as the starting point in this paper.

THEOREM W4: *(*) and (**) have the same solutions.*

PROOF: Observe that since we have assumed state independent resource valuation, by Theorem W2, $\tilde{U}(Y_k, \mathbf{a})$ and $\tilde{M}(Y_k, \mathbf{a})$ have the same arg max, while $\tilde{u}(Y_k, 0, \mathbf{a}) = \tilde{U}(Y_k, \mathbf{a}) + \max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a}')$ trivially has the same arg max. It follows that if $\mathbf{a} \in \arg \max_{\mathbf{a}'} \tilde{u}(Y_k, 0, \mathbf{a})$, then $\mathbf{a} \in \arg \max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a})$ and, by definition, $C(Y_k, \mathbf{a}) = 0$. The former implies $\max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a}') - \tilde{M}(Y_k, \mathbf{a}) = 0$, so $C(Y_k, \mathbf{a}) = \max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a}') - \tilde{M}(Y_k, \mathbf{a})$. Since this also holds by definition for $\mathbf{a} \notin \arg \max_{\mathbf{a}'} \tilde{u}(Y_k, 0, \mathbf{a})$, it holds for all \mathbf{a} . Hence

$$\begin{aligned} & E_{\mathbf{a}, Y_1} \sum_{k=0}^{\infty} \delta^k \left(\left(\tilde{U}(Y_k, \mathbf{a}) + \max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a}') \right) \right. \\ & \quad \left. - \left(\max_{\mathbf{a}'} \tilde{M}(Y_k, \mathbf{a}') - \tilde{M}(Y_k, \mathbf{a}) \right) \right) \\ & = E_{\mathbf{a}, Y_1} \sum_{t=0}^{\infty} \delta^t (\tilde{U}(Y_k, \mathbf{a}) + \tilde{M}(Y_k, \mathbf{a})). \end{aligned}$$

Let A_{kt} be the probability k is alive at t . Then we may write

$$\begin{aligned} & E_{\mathbf{a}, Y_1} \sum_{t=0}^{\infty} \delta^t (\tilde{U}(Y_k, \mathbf{a}) + \tilde{M}(Y_k, \mathbf{a})) \\ & = E_{\mathbf{a}, Y_1} \sum_{k=0}^{\infty} \delta^k (1 - \delta) \left(\sum_{n=k}^{\infty} (\delta \mu)^{n-k} A_{k, k+n} u(a_{k+n}^S, y_{k+n}) \right. \\ & \quad \left. + \sum_{n=k}^{\infty} (\delta \mu)^{n-k} A_{k, k+n} m(y_{k+n}, \tilde{w}_{k+n}) \right) \\ & = E_{\mathbf{a}, Y_1} \sum_{n=0}^{\infty} \delta^n \sum_{k=1}^n (\mu)^{n-k} A_{k, k+n} (1 - \delta) \\ & \quad \times (u(a_{k+n}^S, y_{k+n}) + m(y_{k+n}, \tilde{w}_{k+n})) \\ & = E_{\mathbf{a}, Y_1} \sum_{n=0}^{\infty} \delta^n (1 - \delta) (u(a_{k+n}^S, y_{k+n}) + m(y_{k+n}, \tilde{w}_{k+n})). \quad Q.E.D. \end{aligned}$$

Thus the reduced form of the game is the same agent's objective function that we used in our analysis, hence our study of the solutions of the agent's objective function can be interpreted as an equilibrium of this game.

Notice that we have assumed that the long-run self can commit for the lifetime of the shorter-run self. This is intended to capture the strategic naivete of the shorter-run self as a passive actor. If the long-run self simply moves first each period but cannot commit to contingent plans for future periods, the equilibrium here is still a SR-perfect equilibrium, since we have shown that the solution to the reduced form optimization problem is Markov, so that the long-run self has no wish to renege on his commitment. However, without commitment there can be other equilibria in which the shorter-run self chooses a plan different from that suggested by the long-run self as part of a repeated game equilibrium. However, we regard such equilibria as inconsistent with our notion of the nature of the shorter-run self.

W4. RECURSIVE VERSUS OPPORTUNITY COST

The text supposes that the control cost depends on the foregone value, which is defined with respect to the maximum possible SR utility from tomorrow onward given tomorrow's state. This specification does not necessarily satisfy the property of being an opportunity cost. In general, an *opportunity cost* for the short-run self would have the form

$$\tilde{C}(Y_k, \mathbf{a}) = G(\bar{U}(Y_k) - \tilde{U}(Y_k, \mathbf{a})),$$

where, as in Section W3, $\tilde{U}(Y_k, \mathbf{a}) \equiv (1 - \delta)E_{0, \mathbf{a}, Y_k}^k \sum_{n=0}^{\infty} (\delta\mu)^{n-k} u(a_{k+n}^S, y_{k+n})$. With this specification, the control cost is computed each period by the difference between the best expected present value available to a shorter-run self born in that period and the present value actually received, taking into account what will actually happen in future periods.

With this alternative specification of the cost, the alternative objective function for the long-run self in period 1 is

$$\begin{aligned} \tilde{V}(h_1, \mathbf{a}) &\equiv E_{\mathbf{a}, h_1} \sum_{n=0}^{\infty} \delta^n (u(y_{1+n}, a_{1+n}) \\ &\quad - (1 - \mu)\Gamma[\bar{U}(y_{1+n}) - U^{1+n}(h_{1+n}, \mathbf{a})]) \\ &\quad - \mu\Gamma[\bar{U}(y_1) - U^1(h_1, \mathbf{a})]. \end{aligned}$$

Here the control cost is computed each period by the difference between the best expected present value available to a short-run self born in that period and the present value actually received. It has the form of an *opportunity cost* that depends only on the best present value available to the short-run selves and the actual utility they receive. Note that the name of the short-run self born in period $1 + n$ does not matter, so we may compute the self-control cost

without loss of generality for $k = 1 + n$. This expected present value cost is weighted by $1 - \mu$, which is the probability of a new short-run self being born in a given period. In period 1, however, the weight must be taken to be 1, since the optimization problem always begins with the birth of a new short-run self.

In contrast, the formulation in the paper computes the foregone value and thus the cost in each period “as if” no self-control will be used in future periods. However, in the linear case, these two formulations are equivalent.

THEOREM W5: *If $\tilde{C}(Y_k, \mathbf{a}) = \Gamma(\bar{U}(Y_k) - \tilde{U}(Y_k, \mathbf{a}))$, then $\tilde{C}(Y_k, \mathbf{a}) = C(Y_k, \mathbf{a})$.*

PROOF: Since the identity of the SR self does not matter, it suffices to prove this for $k = 1$. So we must show

$$\begin{aligned} & E_{\mathbf{a}, h_1} \sum_{n=0}^{\infty} \delta^n \Delta(y_{1+n}, a_{1+n}) \\ &= E_{\mathbf{a}, h_1} \sum_{n=0}^{\infty} \delta^n ((1 - \mu)[\bar{U}(y_{1+n}) - U^{1+n}(h_{1+n}, \mathbf{a})]) \\ & \quad + \mu[\bar{U}(y_1) - U^1(h_1, \mathbf{a})]. \end{aligned}$$

We do so by showing that we can apply the principle of optimality for the shorter-run self to compute the opportunity cost as a sum of current and future foregone utilities; then we rearrange the resulting sum to get the desired result. As noted in the text, the principle of optimality for the shorter-run self gives the opportunity cost as a sum of weighted increments,

$$\begin{aligned} \bar{U}(y_n) - U^n(h_n, \mathbf{a}) &= \bar{U}(y_n) - E_{\mathbf{a}, h_n}^n \sum_{\ell=0}^{\infty} (\delta\mu)^\ell u(a_{n+\ell}, y_{n+\ell}) \\ &= E_{\mathbf{a}, h_n}^n \left(\sum_{\ell=0}^{\infty} (\delta\mu)^\ell (\Delta(y_{n+\ell}, a_{n+\ell})) \right). \end{aligned}$$

Writing out the full average present value of opportunity costs, we can in turn express that as a weighted sum of foregone utilities,

$$\begin{aligned} & E_{\mathbf{a}, h_1} \sum_{\ell=0}^{\infty} \delta^\ell ((1 - \mu)[\bar{U}(y_{1+\ell}) - U^{1+\ell}(h_{1+\ell}, \mathbf{a})]) + \mu[\bar{U}(y_1) - U^1(h_1, \mathbf{a})] \\ &= E_{\mathbf{a}, h_1} \sum_{\ell=0}^{\infty} \delta^\ell \left((1 - \mu) \sum_{\ell'=0}^{\infty} (\delta\mu)^{\ell'} (\Delta(y_{1+\ell+\ell'}, a_{1+\ell+\ell'})) \right) \\ & \quad + \mu \left[\sum_{\ell'=0}^{\infty} (\delta\mu)^{\ell'} (\Delta(y_{1+\ell'}, a_{1+\ell'})) \right]. \end{aligned}$$

Set $\ell'' = \ell + \ell'$. The final step is to rearrange this sum to get the recursive cost

$$\begin{aligned}
& E_{\mathbf{a}, h_1} \sum_{\ell''=0}^{\infty} \sum_{\ell' \leq \ell''}^{\infty} ((1 - \mu)(\delta^{\ell''} \mu^{\ell'} \Delta(y_{1+\ell''}, a_{1+\ell''}))) \\
& \quad + \mu \left[\sum_{\ell'=0}^{\infty} (\delta \mu)^{\ell'} (\Delta(y_{1+\ell'}, a_{1+\ell'})) \right] \\
& = E_{\mathbf{a}, h_1} \sum_{\ell''=0}^{\infty} (1 - \mu) \delta^{\ell''} \Delta(y_{1+\ell''}, a_{1+\ell''}) \sum_{\ell' \leq \ell''}^{\infty} \mu^{\ell'} \\
& \quad + \mu \left[\sum_{\ell'=0}^{\infty} (\delta \mu)^{\ell'} (\Delta(y_{1+\ell'}, a_{1+\ell'})) \right] \\
& = E_{\mathbf{a}, h_1} \sum_{\ell''=0}^{\infty} \delta^{\ell''} \Delta(y_{1+\ell''}, a_{1+\ell''}) (1 - \mu^{\ell''+1}) \\
& \quad + \mu \left[\sum_{\ell'=0}^{\infty} (\delta \mu)^{\ell'} (\Delta(y_{1+\ell'}, a_{1+\ell'})) \right] \\
& = E_{\mathbf{a}, h_1} \sum_{\ell''=0}^{\infty} \delta^{\ell''} \Delta(y_{1+\ell''}, a_{1+\ell''}),
\end{aligned}$$

which is the desired result.

The key idea is that the principle of optimality for the shorter-run self enables us to write the overall loss to the shorter-run self as a sum of recursively computed losses,

$$\begin{aligned}
& \bar{U}(y_n) - U^n(h_n, \mathbf{a}) \\
& = \bar{U}(y_n) - (1 - \delta \mu) E_{\mathbf{a}, h_n}^n \sum_{\ell=0}^{\infty} (\delta \mu)^{\ell} u(y_{n+\ell}, a_{n+\ell}) \\
& = E_{\mathbf{a}, h_n}^n (1 - \delta \mu) \left(\sum_{\ell=0}^{\infty} (\delta \mu)^{\ell} (\Delta(y_{n+\ell}, a_{n+\ell})) \right).
\end{aligned}$$

Hence the opportunity cost is just a weighted sum of the increments $\Delta(y_{n+\ell}, a_{n+\ell})$, and the proof simply consists of bookkeeping to verify that the weights are the same as in the recursive case. *Q.E.D.*

In the linear case, in other words, it does not matter whether the cost of imposing self-control on the shorter-run self arises from recursive considera-

tions or from an opportunity cost. The reason we adopted the recursive formulation in the text is that in the nonlinear case, the model of opportunity cost leads to implausible predictions about timing, such as changes in behavior when a shorter-run self “dies.”

The recursive formulation of the text defines a cost of self-control for each state and action to be a convex function g of the difference between the most utility obtainable for any shorter-run self starting in that state and the utility that would be obtained if the shorter-run self next period gets the most possible utility,

$$C(y, a) = g\left(\bar{U}(y) - \left(u(y, a) + \delta \int_Y \bar{U}(y') \pi(y'|y, a)[dy']\right)\right).$$

The long-run self then solves the recursive problem directly without reference to which shorter-run self is born at a particular time, resulting in reduced form utility

$$\tilde{V}_{a, h_n}^n \equiv E_{a, h_n}^n \sum_{\ell=0}^{\infty} \delta^\ell [u(y_{n+\ell}, a_{n+\ell}) - C(y_{n+\ell}, a_{n+\ell})].$$

In the opportunity cost formulation, the cost of self-control for shorter-run self n born in period n is a convex function g of the difference between the most utility obtainable for the shorter-run self and the utility actually obtained,

$$C_{a, h_n}^n = g(\bar{U}(y_n) - U_{a, h_n}^n).$$

Note that there will typically be many shorter-run selves and that the long-run player pays the control cost for the n th shorter-run when that self is “born.”

EXAMPLE W2: We now explore the difference between the two formulations through a simple example. In the example, the recursive formulation captures the simple intuition that when the probability of a temptation is reduced, it becomes less tempting. In the opportunity cost formulation, reducing the probability of a temptation has a complicated effect that depends on the exact timing of when the temptation occurs relative to the “birth” of a new short-run self. We argue that this is both unintuitive and inconsistent with experimental results.

In the example it is known that at some point in the future, a simple temptation will arrive, and it will have value P for the long-run self and value S for the shorter-run self. The exact nature of this opportunity is initially unknown; it is equally likely to be highly tempting (H) or less so (L). The agent learns which opportunity he will face at a time $t_1 + 1$, where t_1 follows an exponential waiting time with parameter p_1 , which is the conditional probability of arrival each period. Once t_1 arrives, the agent is informed which state prevails: there

is a second exponential waiting time with parameter p_2 until the time t_2 when the agent can choose whether to take the tempting action or decline it. Notice that t_1 and t_2 take on the value zero with positive probability, meaning, for example, that if $t_1 = t_2 = 0$, the agent learns immediately that he faces an opportunity and he takes an action during the same period. Note also that when p_1 and p_2 are large compared to the birth rate $1 - \mu$ of the short-run players, the initial short-run self will be the one who makes the eventual decision, while in the reverse case, a short-run self will probably be born after t_1 and soon before t_2 .

In the highly tempting state H , the simple temptation if chosen is received for sure. In the less tempting state L , there is only a probability q if the temptation is chosen that it will be received.

EXAMPLE W2a—Recursive Cost of Self-Control: In the recursive case, the only relevant cost of self-control is in the period in which action is taken. The maximization problem can be written conditional on whether state H or L has occurred, so the condition for taking is exactly as in the text: in H the optimum is to take if

$$P < g(S),$$

while in L the condition is

$$qP < g(qS).$$

Note moreover that this solution holds regardless of the values of p_1 and p_2 .

EXAMPLE W2b—Opportunity Cost of Self-Control: Here the long-run self faces one decision problem before t_1 , a different decision problem once the information is revealed at t_1 but before action at t_2 , and yet a third at the decision time t_2 . Each of these problems is stationary and independent of the past history. Each one corresponds to a shorter-run self who faces a different temptation. To compute that temptation, we simply compute the probability distribution over arrivals of the event that the action becomes available, conditional on which type the shorter-run self is.

To calculate the temptation utility for any shorter-run self born at or after t_1 and strictly before t_2 in state H , we compute the expectation of the discounted value of S over different values of t_2 :

$$\begin{aligned} U^{2H} &= E(\mu\delta)^{t_2} S = \sum_{t_2=1}^{\infty} (\mu\delta)^{t_2} S p_2 (1 - p_2)^{t_2-1} \\ &= \frac{\mu\delta p_2 S}{(1 - \mu\delta(1 - p_2))}. \end{aligned}$$

In state L , the temptation utility is just $U^{2L} = qU^{2H}$. For the short-run self born at exactly t_2 , the temptation utilities are $U^{2H} = S$ and $U^{2L} = qS$. In either case, the temptation costs are given by $g(U^{2H})$ and $g(qU^{2H})$. If p_1 is large relative to p_2 , then it is far more likely that the shorter-run self is born after t_1 than before, so most of the weight in the objective function of the long-run self is on these temptation costs. This leads to an analysis that is qualitatively similar to the recursive case. In the laboratory, however, the time between when the subjects learn they will face a particular decision problem and the time they are asked to take an action is usually on the order of minutes, while we expect the mean length of life of a shorter-run self³ to be on the order of a day, so that it is much more likely that the shorter-run self is born before t_1 rather than after.

The temptation utility for a self born before t_1 is the expected value of taking the temptation, averaged over the arrival times t_1 and t_2 . Since these times are independent, the temptation value is

$$\begin{aligned} U^1 &= E[(\mu\delta)^{(t_1+t_2)}0.5(1+q)S] \\ &= \frac{0.5\mu\delta p_1 p_2(1+q)S}{(1-\delta\mu(1-p_1))(1-\delta\mu(1-p_2))}. \end{aligned}$$

Let us focus on the case where most of the weight is on shorter-run selves born before t_1 . Let a^J be an indicator function of whether the long-run self takes at $J \in \{H, L\}$. Let $Q = 0.5(a^H + qa^L)$ be the overall probability of the long-run self taking. Then the objective of the long-run self is

$$\begin{aligned} V(Q) &= \frac{\delta p_1 p_2}{(1-\delta(1-p_1))(1-\delta(1-p_2))}QP \\ &\quad - g\left(\frac{\mu\delta p_1 p_2 S}{(1-\delta\mu(1-p_1))(1-\delta\mu(1-p_2))}(0.5(1+q) - Q)\right). \end{aligned}$$

In particular, the objective function depends only on Q . Notice that there are four relevant values of Q : $0, 0.5q, 0.5, 0.5(1+q)$. Note moreover that since g is convex,

$$\begin{aligned} V'_0 &\equiv \frac{V(0.5q) - V(0)}{0.5Q} \geq V'_{0.5q} \equiv \frac{V(0.5) - V(0.5q)}{0.5 - 0.5q} \\ &\geq V'_{0.5} \equiv \frac{V(0.5(1+q)) - V(0.5)}{0.5(1+q) - 0.5}. \end{aligned}$$

It follows that the optimal choice of Q is 0 if $V'_0 \leq 0$; it is $0.5q$ if $V'_0 \geq 0$, $V'_{0.5q} \leq 0$, it is q if $V'_{0.5q} \geq 0$, $V'_{0.5} \leq 0$, and it is $0.5(1+q)$ if $V'_{0.5} \geq 0$, and since

³See Fudenberg and Levine (2010).

increasing P lowers the slope of V , the optimal Q is a decreasing function of P . If P is close to zero, then Q should be equal to 1, meaning that the long-run self takes regardless of the state. If P is highly negative, then Q should be equal to zero, meaning that the long-run self never takes. As P is reduced from zero, eventually the probability Q must be reduced from 1 to 0.5, meaning that the optimum is to take only in the state H . This is the same behavior exhibited in the case in which most shorter-run selves are born after t_2 or that of a recursive shorter-run self. As we reduce P further, the probability Q must drop again from 0.5 to $0.5q$, meaning that the optimum is to take only in the state L . This is the opposite direction from that of a shorter-run self born after t_2 or a recursive shorter-run self. It also is contrary to indicating that reversals have the form that as the probability of a prize is reduced, behavior favored by the long-run self is more likely to be observed.⁴

Dept. of Economics, Harvard University, 310 Littauer Center, Cambridge, MA 02138, U.S.A.; dfudenberg@harvard.edu

and

Dept. of Economics, Washington University in St. Louis, St. Louis, MO 63130-4899, U.S.A.; david@dklevine.com.

Manuscript received November, 2010; final revision received May, 2011.

⁴See Fudenberg and Levine (2010).