POVERTY AND SELF-CONTROL

B. Douglas Bernheim
Stanford University and NBER

Debraj Ray
New York University

Şevin Yeltekin
Carnegie Mellon University

January 2013, revised March 2014

"*When you ain't got nothin', you got nothin' to lose.*" Bob Dylan

ABSTRACT

The absence of self-control is viewed as an important correlate of persistent poverty. But the direction of causality — whether poor self-control leads to poverty, or poverty perpetuates itself by undermining self-control — is unclear. According to psychologists, a common method of self-control employs contingent self-reinforcement (both punishment and reward). We interpret such behavior through the lens of game theory, and explore the potential for exercising credible self-reinforcement in a standard intertemporal allocation problem with credit constraints, where the individual has quasi-hyperbolic preferences. Our main result demonstrates that low initial assets can limit the scope for self-control, so that some people become trapped in poverty, while those with high initial assets can accumulate indefinitely. Thus, policies designed to help the poor *initiate* an accumulation plan may be effective, even if they are temporary. We examine various implications: the effect of access to credit on saving, the demand for commitment devices, the design of financial accounts to promote accumulation, and the variation of the marginal propensity to consume across classes of resource claims. We also explore the structure of optimal contingent self-reinforcement, demonstrating that it has a simple and behaviorally plausible structure that is immune to self-renegotiation.

Bernheim's research was supported by National Science Foundation Grants SES-0752854 and SES-1156263. Ray's research was supported by National Science Foundation Grant SES-1261560. We are grateful to Severine Toussaert for her careful reading of the manuscript. We would also like to thank the three anonymous referees and the co-editor of *Econometrica*, participants of the seminars at BYU Computational Public Economics Conference in Park City, Utah, Brown University, Brown Experimental and Economic Theory (BEET) Conference on Temptation and Self Control, Carnegie Mellon University, the ECORE Summer School at Louvain-la-Neuve, Indian Statistical Institute Delhi, Koc University in Istanbul, Minneapolis FED, National Graduate Institute for Policy Studies in Tokyo, Nottingham, Stanford GSB, the Stanford SITE Workshop on Psychology and Economics, ThReD Workshop, Harvard-MIT, UCSD and Yale University for their suggestions and comments.

1. Introduction

Recent research indicates that the poor not only borrow at high rates,[1] but also forego profitable small investments.[2] These behavioral patterns contribute to the persistence of poverty, particularly (but not exclusively) in developing countries. Traditional theory (based on high rates of discount, minimum subsistence needs, and/or aspiration failures) can take us only part of the way to an explanation. Because the poor also exhibit a demand for commitment,[3] it is likely that time inconsistency plays an important contributory role.

It is generally understood that time inconsistency can create self-control problems.[4] Our interest lies in the possibility that difficult economic circumstances may exacerbate these problems. If self-control (or the lack thereof) is a fixed trait independent of economic circumstances, then the outlook for policy interventions that encourage the poor to invest in their futures – particularly short-term interventions – is not good. But if poverty perpetuates itself by impairing the ability of time-inconsistent decision makers to exercise self-control,[5] even temporary policies that help the poor *begin* to accumulate assets may be highly effective.

The absence of self-control is on display when an individual is unable to follow through on a desired plan of action. What then constitutes the *exercise* of self-control? At its core, the term implies an internal mechanism, though a broader interpretation would include the deliberate adoption of externally-enforced commitment devices. We focus on methods of self-control that involve the contingent use of self-reinforcement, including self-punishment and self-reward —

[1] Informal interest rates in developing countries are notoriously high; see, e.g., Aleem (1993). So are formal interest rates; e.g., Bangladesh recently capped microfinance interest rates at 27% per annum, a restriction frowned upon by the *Economist* ("Leave Well Alone," November 18, 2010). Citing other literature, Banerjee and Mullainathan (2010) argue that such loans are taken routinely rather than on an emergency basis.

[2] High returns have been documented for agricultural investments in Ghana, even on small plots (Goldstein and Udry (1999) and Udry and Anagol (2006)), the use of small amounts of fertilizer in Kenya (Duflo, Kremer, and Robinson (2011)), and microenterprise in Sri Lanka (de Mel, McKenzie, and Woodruff (2008)). See Banarjee and Duflo (2011) for additional references.

[3] See, for example, Shipton (1992) on the use of lockboxes in Gambia, Benartzi and Thaler (2004) on employee commitments to save out of future wage increases in the United States, and Ashraf, Karlan, and Yin (2006) on the use of a commitment savings product in the Philippines. Aliber (2001), Gugerty (2007), and Anderson and Baland (2002) view ROSCA participation as a commitment device; see also the theoretical contributions of Ambec and Treich (2007) and Basu (2011). Duflo, Kremer, and Robinson (2011) attribute low fertilizer use in Kenya to a lack of commitment opportunities. Commitment issues also feature prominently in the ongoing debate over whether to replace the public system for distributing food in India with a cash-based program; see Khera (2011).

[4] See, for instance, Strotz (1956), Phelps and Pollak (1968), Ainslie (1975, 1991), Thaler and Shefrin (1981), Akerlof (1991), Laibson (1997), O'Donoghue and Rabin (1999) and Ashraf, Karlan, and Yin (2006). Such inconsistency may be internal to the individual, or have social origins stemming from discordance within the household (e.g., spouses with different discount factors) or from demands made by the wider community (e.g., sharing among kin or extended family).

[5] Aspiration failures can create similar traps. See, e.g., Appadurai (2004), Ray (2006), Genicot and Ray (2006) and the *United Nations Development Program* Regional Report for Latin America, 2010. However, this complementary approach does not generate a demand for commitment devices.

concepts that have played an important role in psychology — to establish incentives. Formally, we model time-inconsistent behavior by examining the subgame-perfect Nash equilibria of a dynamic game played by successive incarnations of a single decision maker with *quasi-hyperbolic preferences* (also known as $\beta\delta$-discounting).[6] Within that setting, we interpret history-dependent equilibrium strategies as methods of deploying credible self-punishment/reward. We discuss the psychological foundations for this interpretation in Section 2. In contrast, the bulk of the existing literature focuses almost exclusively on Markov-perfect equilibria, thereby ruling out contingent self-reinforcement.[7]

We study a standard intertemporal allocation problem in which the individual faces a credit constraint. To avoid trivially building in (by assumption) a relationship between initial assets and the rate of saving, we take preferences to be homothetic and the accumulation technology to be linear. To determine the full scope for self-control, we study the set of *all* subgame-perfect Nash equilibria. This approach allows us to identify the conditions under which wealth accumulation can occur through the exercise of self-control. In particular, we ask whether self-control is more difficult when initial assets are low than when they are high.

It is notoriously difficult to characterize the set of subgame-perfect Nash equilibria for all but the simplest dynamic games. The problem of self-control we study here is, alas, no exception. Obviously, for some parameter values, equilibrium accumulation will occur (or, alternatively, will never occur) regardless of initial assets. Yet there are also parameter values (with intermediate degrees of time inconsistency) for which equilibrium accumulation will depend on the initial asset level. Our main result demonstrates that, in every such case, there is an asset level below which liquid wealth is exhausted in finite time (a poverty trap), as well as an asset level above which the most attractive equilibria give rise to unbounded accumulation.

Figure 1 illustrates our main result computationally.[8] The horizontal axes in each panel measure assets in the current period. The vertical axis in panel (A) similarly measures continuation asset choices for the next period. Thus, points above, on, and below the 45 degree line indicate asset accumulation, maintenance, and decumulation respectively. The example displays a poverty trap – i.e., an asset threshold below which *all* equilibria lead to decumulation because the individual cannot exercise sufficient self-control through contingent self-reinforcement. However, above

[6] Strotz (1956) pioneered this general approach. Models of quasi-hyperbolic discounting were popularized by Laibson (1994, 1996, 1997) and O'Donoghue and Rabin (1999).

[7] We expand further on this point in Section 7. Exceptions to the use of Markov equilibrium include Laibson (1994), Bernheim, Ray, and Yeltekin (1999), and Benhabib and Bisin (2001).

[8] For a complete explanation of our computational methods, and for details concerning all computational examples presented in the text, see the Online Appendix. For this exercise, we set the rate of return equal to 30%, the discount factor equal to 0.8, the hyperbolic parameter (β) equal to 0.4, and the constant elasticity parameter of the utility function equal to 0.5. We chose these values so that the interesting features of the equilibrium set are easily visible; qualitatively similar features arise for more realistic parameter values.

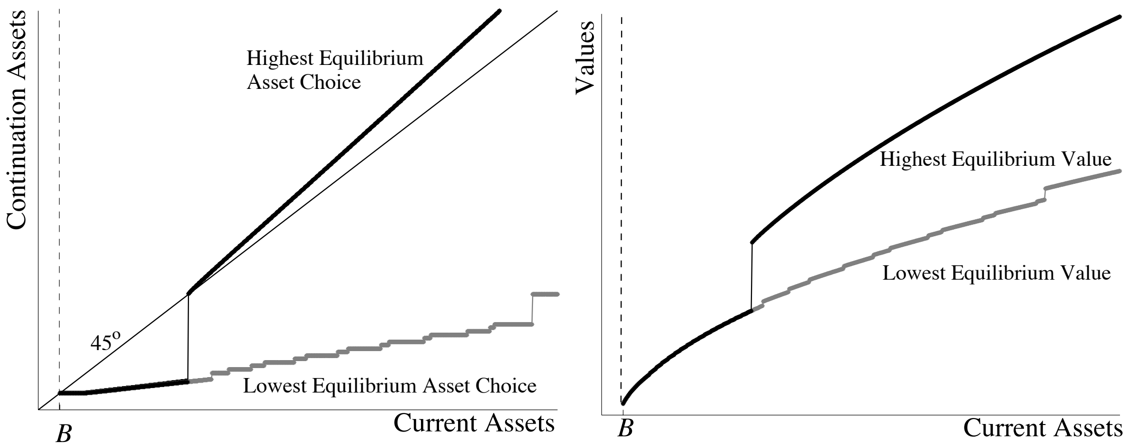

(A) Equilibrium Asset Choices

(B) Equilibrium Values

FIGURE 1. ACCUMULATION AND VALUES AT DIFFERENT ASSET LEVELS.

that threshold, credible patterns of self-reinforcement allow her to accumulate greater assets. Indeed, as we will see later, the most attractive equilibria starting from above the critical threshold lead to unbounded accumulation.[9]

A natural and intuitive explanation for this result is that credible self-punishments become (proportionately) more severe as assets accumulate. Intuitively, if self-discipline collapses, the decision maker does not have much to lose if she is up against her credit constraint, but she has a great deal to lose if she possesses substantial wealth. It turns out, however, that the problem is considerably more complicated than this simple intuition suggests,[10] because the credit constraint infects feasible behavior (and hence worst punishments) at *all* asset levels in subtle ways. The example in Figure 1 illustrates this point dramatically: there are asset levels at which the *lowest* level of continuation assets jumps up discontinuously. As assets cross those thresholds, the worst punishment becomes *less* rather than *more* severe. This is shown in panel (B) of the Figure, which plots equilibrium values: note that the lowest such value (which serves as the worst punishment) jumps upward at several asset levels.[11] To prove our main result, we show that at high asset levels, self-punishments become sufficiently severe to sustain accumulation, but plainly they do not do so monotonically.

[9] This is a more subtle point that cannot be seen directly from Figure 1, though it is indicative. It is more subtle because repeated application of the highest continuation asset need not be an equilibrium, and moreover, even if it were, it need not be the most attractive equilibrium.

[10] The overwhelmingly numerical nature of our earlier working paper, Bernheim, Ray, and Yeltekin (1999), bears witness to this assertion.

[11] The jagged nature of the lowest value in panel (B) is not a numerical artifact; it reflects actual jumps.

One might object to an examination of the entire set of subgame-perfect equilibria on the grounds that many such equilibria involve unreasonably complex patterns of contingent self-reinforcement. Yet we show that the worst credible punishments involve a surprisingly simple, intuitive, and behaviorally plausible pattern of contingent self-reinforcement. In effect, the individual sets a personal standard of behavior, which specifies how much she should save. If she fails to meet this standard, she self-punishes. The punishment involves a temporary binge, which we prove cannot exceed two periods; once it ends, she rededicates herself to her best achievable personal standards. Thus, if she "falls off the wagon," she soon climbs back on; she responds to a lapse by "getting it out of her system" so she can adhere to her standards. As discussed in Section 6.4, these punishments are also renegotiation-proof in the sense of Bernheim and Ray (1989) and Farrell and Maskin (1989).

Our analysis has a number of provocative implications for economic behavior and public policy. First, the relationship between assets and self-control argues for the use of "pump-priming" interventions that encourage the poor to start saving, while relying on self-control to sustain frugality at higher levels of assets. Second, policies that improve access to credit can help people become savers. Intuitively, with greater access to credit, the consequences of a break in discipline become more severe, and hence discipline is easier to sustain. But there is an important qualification: those who still cannot exercise self-control fall further into debt. Third, external commitment devices can undermine the effectiveness of internal self-control mechanisms. Consequently, when the latter are reasonably effective, people may avoid the former, even if they understand their self-control problems. Our theory therefore potentially accounts for the puzzling lack of demand for commitment devices observed in many contexts, particularly among the non-poor.[12] Fourth, it may be possible to increase the effectiveness of incentives to save through special accounts (e.g., IRAs) by requiring the individual to establish a savings target, locking up all funds until the target is achieved, and then removing the lock (thereby rendering all of the funds liquid). Pilot programs with such features have indeed been tested in developing countries.[13] Finally, our analysis provides a potential explanation for the observation that the marginal propensity to consume differs across classes of resource claims, and offers a new perspective on the "excess sensitivity" of consumption to income.

Related literature. As noted above, we build on our unpublished working paper Bernheim, Ray, and Yeltekin (1999), which made its points through simulations, but did not contain analytical results. Banerjee and Mullainathan (2010) also argue that self-control problems give rise to low asset traps, but their analysis has little in common with ours. They examine a novel model of time-inconsistent preferences, in which rates of discount differ from one good to another. "Temptation goods" (those to which greater discount rates are applied) are presumed

[12] See, e.g., Bryan, Karlan, and Nelson (2010) and DellaVigna (2009).
[13] See Ashraf, Karlan, and Yin (2006), as well as Karlan et al. (2010).

to be inferior goods: this *assumed* non-hometheticity of preferences automatically builds in a tendency to dissave when resources are limited. The validity of this central assumption (e.g., whether a poor person spends relatively more of his budget on alcohol than a richer person does on, say, designer drugs or Ipads) is plainly an empirical matter. Our approach requires no such assumption; indeed, we adopt a standard model of time-inconsistency in which preferences are homothetic and the accumulation technology linear. Scale effects arise only from the interplay between credit constraints and equilibrium conditions.[14] Consequently, our analysis is essentially orthogonal (and hence potentially complementary) to that of Banerjee and Mullainathan (2010): theirs is driven by assumed scaling effects in rewards, while ours is driven by scaling effects in punishments arising from assumed credit market imperfections.[15]

Organization of the paper. Section 2 discusses psychological foundations. Section 3 describes the model and defines equilibrium. Section 4 provides a characterization of the equilibrium value set that is useful both conceptually and in developing a numerical algorithm. Section 5 defines self-control, and Section 6 studies its relationship to initial wealth. Sections 7 discusses Markov equilibrium. Section 8 explores additional implications. Section 9 presents conclusions and some directions for future research. Proofs and technical details for numerical computations appear in appendices.

2. Psychological Foundations

The defining feature of the self-control mechanisms that we model in this paper is that people intentionally create and execute plans for self-reinforcement, establishing incentives by punishing themselves for deviations from (or rewarding themselves for conformity with) desired behavior.[16] Psychological foundations for this mechanism are found in the literatures on self-regulation and behavior modification. For example, in one early study of self-regulation, Bandura and Kupers (1964) observed, "[b]y contrast [to rats or chimpanzees], people typically make self-reinforcement contingent on their performing certain classes of responses which they have come to value as an index of personal merit. They often set themselves relatively explicit criteria of achievement, failure to meet which is considered undeserving of self-reward and may elicit

[14] Moav and Neeman (2012) analyze a model in which conspicuous consumption generates poverty traps. They also assume homothetic preferences, and emphasize the same difference from the Banerjee and Mullainathan (2010) approach.

[15] Our analysis is also related to that of Laibson (1994) and Benhabib and Bisin (2001), both of whom examine a model similar to ours, except that there is no credit constraint. Because the resulting model is fully scalable, so is the equilibrium set; consequently, there is no relationship between poverty and self-control.

[16] Certainly, one can frame any form of self-reinforcement as either a reward or a punishment; it is the *difference* in outcomes that creates incentives. Psychologists do not, however, view this framing as neutral, and there is some evidence that people can achieve self-control more effectively by framing self-administered consequences as rewards rather than as punishments; see, e.g., Mahoney, Moura, and Wade (1973). Such framing effects are beyond the scope of our investigation.

self-denial or even self-punitive responses; on the other hand, they tend to reward themselves generously on those occasions when they attain their self-imposed standards."[17] Rehm (1977) notes that "[s]elf-reinforcement has been a major focus of self-control research and many clinical uses of self-administered reward and punishment programs have been described."[18] Likewise, according to Kazdin (2012), "Self-reinforcement and self-punishment techniques have been incorporated into intervention programs and applied to a wide range of problems..."

As explained in the Introduction, we model the decisions of a time-inconsistent individual by studying a dynamic game played by his successive incarnations. In that setting, the individual engages in self-reinforcement by deploying history-dependent strategies, which specify contingent patterns of behavior that serve as rewards and/or punishments, as in the psychological literature. Subgame perfection takes this a step further: it ensures that self-reinforcement is *credibly* implementable. With this interpretation, the scope for exercising self-control through self-punishment/reward is sharply defined by the set of outcomes that can arise in subgame-perfect Nash equilibria.

Not surprisingly, psychologists do not typically employ the language of game theory or the formal logic of subgame perfection. Yet they have long recognized that credibility problems limit the scope for effective self-reinforcement. Ainslie (1975) succinctly summarizes the problem thus: "Self-reward is an intuitively pleasing strategy until one asks how the self-rewarding behavior is itself controlled..."[19]

The logic of using history-dependent strategies to overcome the credibility problem is a recurring theme in Ainslee's work. In particular, he observes that people often successfully adopt personal rules (e.g., "always go to bed early"), which they enforce by construing local deviations to have global significance (e.g., "if I go to bed late today, then I will go to bed late every night"); see Ainslie (1975, 1991). Viewed through our game-theoretic lens, a personal rule is an equilibrium

---

[17] Similarly, Mischel (1973) observes:"The essence of self-regulatory systems is the subject's adoption of contingency rules that guide his behavior in the absence of, and sometimes in spite of, immediate external situational pressures. Such rules specify the kinds of behavior appropriate (expected) under particular conditions, the performance levels (standards, goals) which the behavior must achieve, and the consequences (positive and negative) of attaining or failing to reach those standards." See also Bandura (1971, 1976).

[18] Similarly, Bandura (1976) notes that "[a]mong the various self-regulator phenomena that have been investigated within [the social learning] framework, self-reinforcement occupies a prominent position. In this process, individuals regulate their behavior by making self-reward conditional upon matching self-prescribed standards of performance...[C]ontrol is vested to a large extent in the hands of individuals themselves: they set their own goals, they monitor and evaluate their own performances, and they serve as their own reinforcing agents."

[19] He goes on to argue: "A subject does not actually recruit additional reward by planning to delay a cigarette until he has finished a difficult task. On the contrary, he sets himself a second task: He must both defer smoking and work on his original task on the basis of the same differential reward that has always confronted him." See also Rachlin (1974) and Kazdin (2012).

path for a dynamic intrapersonal game, and the global consequences that support it are the off-equilibrium paths triggered by deviations.[20]

Our interpretation of Ainslie's writings differs from that of both Benabou and Tirole (2004) and Ali (2011). They interpret an individual's contingent beliefs about his own future actions as evidence-based forecasts rather than deliberately contrived arrangements. Some passages in Ainslie's writings are consistent with this interpretation.[21] Yet in other passages, he emphasizes that people intentionally create this conditionality. Ainslie (1991) is particularly instructive on this point. For example, he writes that "…insofar as [the individual] has become aware of this phenomenon, he will be able to induce it where it has not occurred spontaneously, by arbitrarily defining a category of gratification-delaying behaviors that will thereafter prevail or not as a set." Accordingly, he describes personal rules as the mechanism by which "...the person can *arrange* consistent motivation" (emphasis added) for a "prolonged course of action." Indeed, after describing the choices of a time-inconsistent decision maker as involving an "intertemporal prisoner's dilemma," he characterizes personal rules as "a solution to the bargaining problem" between an individual's "successive motivational states." Moreover, citing Klein and Leffler (1981), he notes that "[t]he same logic is the basis for what is called a 'self-enforcing contract' between individuals." To illustrate the use of a personal rule, he examines a simple numerical model of a hyperbolic discounter who, in each of a succession of periods, decides whether to stay up late or go to bed early. Because his model involves no uncertainty concerning preferences, it entails no inference problem. Yet he informally describes a subgame-perfect equilibrium in which the individual exercises self-discipline (going to bed early every night) by contriving a conditional self-punishment (staying up late for ten consecutive nights), and he describes this solution in game theoretic terms: "Insofar as [the individual] sees his current choice as a precedent and not an isolated incident, he will face the incentives of a repeated prisoner's dilemma."

Portions of Ainslie (1991), including the aforementioned numerical example, appear to invoke Nash (or, more generally, Markov) reversion as a solution to the credibility problem. Yet Ainslie (1975) also describes more complex patterns of self-reinforcement. For instance, he discusses the case of an individual who, in order to keep his shoes shined, adopts a personal rule specifying that he must shine them before breakfast (otherwise he will refrain from shining them in the future). Ainslie posits that, upon oversleeping, the subject might be tempted not only to skip this chore, but also to skip the punishment. A secondary punishment is required for that contingency; Ainslie suggests that the subject might skip breakfast (and thereby conform to the letter of his

[20] Laibson (1997), Bernheim, Ray, and Yeltekin (1999), and Benhabib and Bisin (2001) have previously adopted this interpretation.

[21] For example, Ainslie (1991) writes, "If [an individual] makes an impulsive choice, he will have little reason to believe he will not go on doing so, and if he controls his impulse, he has evidence that he may go on doing that." However, in context, one can also read this passage as a reference to the expectations that prevail in a particular intrapersonal equilibrium.

rule). Thus, according to Ainslie (1975), complex patterns of self-reinforcement simply require "skill at private side bets."

3. MODEL

3.1. **Feasible Set and Preferences.** Feasibility constraints link current assets, current consumption and future assets, starting from an initial asset level A_0:

$$c_t = A_t - (A_{t+1}/\alpha) \geq 0, \tag{1}$$

where $\alpha - 1 > 0$ is the rate of return. Assets must also respect a lower bound:

$$A_t \geq B > 0. \tag{2}$$

We interpret B as a credit constraint.[22] For instance, if the individual earns a constant income y, then $A_t = F_t + [\alpha y/(\alpha - 1)]$, where F_t is financial wealth. If she can borrow only some fraction $(1 - \lambda)$ of future income, then $B = \lambda \alpha y/(\alpha - 1)$.

Individuals have quasi-hyperbolic preferences: lifetime utility is given by

$$u(c_0) + \beta \sum_{t=1}^{\infty} \delta^t u(c_t),$$

where $\beta \in (0, 1)$ and $\delta \in (0, 1)$. We assume that u has the constant-elasticity (and homothetic) form

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

for $\sigma > 0$, with the understanding that $\sigma = 1$ refers to the logarithmic case $u(c) = \ln c$. With a linear accumulation technology and homothetic preferences, all scale effects must arise from the interplay between credit constraints and equilibrium conditions.

3.2. **Restrictions on the Model.** The *Ramsey program* from A is the path $\{A_t\}$ that maximizes

$$\sum_{t=0}^{\infty} \delta^t \frac{c_t^{1-\sigma}}{1-\sigma},$$

with initial stock $A_0 = A$. It is constructed without reference to the hyperbolic factor β. This program is well-defined provided utilities do not diverge, which we ensure by assuming that

$$\gamma \equiv \delta^{1/\sigma} \alpha^{(1-\sigma)/\sigma} < 1. \tag{3}$$

We presume throughout that the Ramsey program exhibits growth, which requires

$$\delta \alpha > 1. \tag{4}$$

[22] Another interpretation of B is that it is an investment in an illiquid asset. We return to this interpretation when we discuss policy implications.

Under (3) and (4), the value of the Ramsey program is finite, and $c_t = (1-\gamma)A_t$, while assets grow exponentially: $A_{t+1} = A_0 \left(\delta^{1/\sigma}\alpha^{1/\sigma}\right)^t = A_0 \left(\gamma\alpha\right)^t$.

For $\sigma \geq 1$, utility is unbounded below, and so it is possible to sustain virtually any outcome using punishments that impose zero consumption, or a progressively more punitive sequence of vanishingly small consumptions; see Laibson (1994). Such devices are contrived and unrealistic; we rule them out by assuming that consumption is bounded below at every asset level:

$$c_t \geq \upsilon A_t, \tag{5}$$

where υ is some small but positive number. Formally, it is enough to take $\upsilon < 1-\gamma$, so that Ramsey accumulation is feasible, but we think of υ as tiny. We assume the lower bound on consumption is proportional to assets so as to avoid introducing artificial scale effects.

3.3. **Equilibrium.** Continuation asset A' is *feasible* given A if $A' \in [B, \alpha(1-\upsilon)A]$. A *history* h_t at date t is a feasible sequence of assets $(A_0, \ldots, A_t)$ through date t. A *policy* ϕ specifies a feasible continuation asset $\phi(h_t)$ following every history. If the history h_t is followed by the asset choice x, we write the resulting history as $h_t.x$. A policy ϕ yields a *value* V_ϕ, as follows:

$$V_\phi(h_t) \equiv \sum_{s=t}^{\infty} \delta^{s-t} u\left(A(h_s) - \frac{\phi(h_s)}{\alpha}\right),$$

where $A(h_t)$ denotes the last element of h_t and, recursively, $h_{s+1} = h_s.\phi(h_s)$ for $s \geq t$. Similarly, ϕ also yields a *payoff* P_ϕ:

$$P_\phi(h_t) \equiv u\left(A(h_t) - \frac{\phi(h_t)}{\alpha}\right) + \beta\delta V_\phi(h_t.\phi(h_t)).$$

Note that payoffs include the hyperbolic factor β, while values do not.

An *equilibrium* is a policy such at every history h_t,

$$P_\phi(h_t) \geq u\left(A(h_t) - \frac{x}{\alpha}\right) + \beta\delta V_\phi(h_t.x). \tag{6}$$

for each x feasible given $A(h_t)$.[23]

For some of our results, it will be useful to assume that the set of equilibrium continuation values is always convex. We therefore suppose that, following any asset choice, the continuation plan can be chosen (if needed) using a public randomization device.[24] Generalizing our notation to encompass public randomization is routine; we skip the details for the sake of brevity.

[23] Given the many-self perspective that we maintain in the paper, the "one-shot deviation principle" embedded in this definition is natural. Coordinated deviations among multiple selves belong more properly to the subject of renegotiation; see Section 6.4 and Footnote 39.

[24] Here, "public" randomization involves conditioning the continuation equilibrium on the realization of a random variable that the individual does not subsequently forget.

4. EXISTENCE AND CHARACTERIZATION OF EQUILIBRIUM

For each $A \geq B$, let $\mathcal{V}(A)$ be the set of all equilibrium values available at A. If $\mathcal{V}(A)$ is nonempty, let $H(A)$ and $L(A)$ be its supremum and infimum values. Under (3) and (5), it is obvious that

$$-\infty < L(A) \leq H(A) \leq R(A) < \infty,$$

where $R(A)$ is the Ramsey value. In fact, a tighter bound is available for worst values:

OBSERVATION **1.** *Suppose that $\mathcal{V}(A)$ is nonempty for every $A \geq B$. Then*

$$L(A) \geq u\left(A - \frac{B}{\alpha}\right) + \delta L(B) \geq u\left(A - \frac{B}{\alpha}\right) + \frac{\delta}{1-\delta} u\left(\frac{\alpha - 1}{\alpha} B\right) \tag{7}$$

Observation 1 establishes a baseline for iterating a self-generation map. To this end, consider a nonempty-valued correspondence $\mathcal{W}$ on $[B, \infty)$ such that for all $A \geq B$,

$$\mathcal{W}(A) \subseteq \left[u\left(A - \frac{B}{\alpha}\right) + \frac{\delta}{1-\delta} u\left(\frac{\alpha - 1}{\alpha} B\right), R(A)\right]. \tag{8}$$

Say that $\mathcal{W}$ *supports the value* w *at asset level* A if there is a feasible asset choice x and $V \in \mathcal{W}(x)$ — a *continuation* $\{x, V\}$, in short — with

$$w = u\left(A - \frac{x}{\alpha}\right) + \delta V, \tag{9}$$

while for every feasible x',

$$u\left(A - \frac{x}{\alpha}\right) + \beta\delta V \geq u\left(A - \frac{x'}{\alpha}\right) + \beta\delta V'. \tag{10}$$

for some $V' \in \mathcal{W}(x')$. That is, the value w at A is "incentive-compatible" given suitable choices of continuation values from $\mathcal{W}$. Now say that $\mathcal{W}$ *generates* the correspondence $\mathcal{W}'$ if for every $A \geq B$, $\mathcal{W}'(A)$ is the convex hull of all values supported at A by $\mathcal{W}$. We use the convex hull to capture public randomization (i.e., an asset choice can yield a lottery over continuation values).

Standard arguments tell us that the equilibrium correspondence $\mathcal{V}$ generates itself, and indeed it contains any other correspondence that does so. Accordingly, define a sequence of correspondences on $[B, \infty)$, $\{\mathcal{V}_k\}$, by

$$\mathcal{V}_0(A) = \left[u\left(A - \frac{B}{\alpha}\right) + \frac{\delta}{1-\delta} u\left(\frac{\alpha - 1}{\alpha} B\right), R(A)\right]$$

for every $A \geq B$, and recursively, $\mathcal{V}_k$ generates $\mathcal{V}_{k+1}$ for all $k \geq 0$. It is obvious that the graph of $\mathcal{V}_k$ contains the graph of $\mathcal{V}_{k+1}$. We assert

PROPOSITION **1.** *An equilibrium exists from any initial asset level, so that the equilibrium correspondence $\mathcal{V}$ is nonempty-valued. Moreover, for every $A \geq B$,*

$$\mathcal{V}(A) = \bigcap_{k=0}^{\infty} \mathcal{V}_k(A). \tag{11}$$

Also, $\mathcal{V}$ is convex-valued and has closed graph.

Despite the fact that assets are unbounded (so that the graph of $\mathcal{V}_k$ is not a compact object), the proof of this result involves fairly standard arguments. Note that it both establishes existence and suggests a computational algorithm, which we employ for all numerical examples; see the Online Appendix for details.[25] In addition, the proposition shows that the values $L(A)$ and $H(A)$ are in fact achievable.

Notice that $L(\cdot)$ defines the worst credible self-punishments. Thus, to determine whether a particular path $(A_0, A_1, ...)$ is sustainable, we can consider a strategy that follows the path while punishing any period t deviation to $x \neq A_{t+1}$ with $L(x)$, and ask whether it is an equilibrium.

5. SELF CONTROL

We will say that there is *self-control* at asset level A if the agent is capable of strict asset accumulation starting from A in some equilibrium. Furthermore, say that there is *strong self-control* at A if the agent is capable of unbounded accumulation — i.e., $A_t \to \infty$ — along some equilibrium path starting from A. These definitions capture an essential qualitative feature of behavior: the ability to accumulate. Notice that they describe a possibility: they tell us that equilibrium accumulation is an option, not that the individual will surely exercise that option.

Now consider the absence of self-control. We say that *self control fails* at A if every equilibrium continuation asset level is strictly smaller than A, and more forcefully, that there is a *poverty trap at A* if in every equilibrium, assets decline over time from A to the lower bound B. Notice that these definitions describe an inevitability rather than a possibility.

We call a case *uniform* if either (i) there is no asset level at which self-control is possible, or (ii) there is no asset level at which self-control fails. Case (i) arises with a sufficiently high degree of time inconsistency (β small), and case (ii) arises with β sufficiently close to 1 (because some equilibrium then approximates the Ramsey program, which involves indefinite accumulation). With perfect credit markets ($B = 0$), uniformity necessarily prevails: if continuation assets x can be sustained at asset level A, then continuation assets λx can be sustained at asset level λA

[25] Public randomization is not needed to establish existence; the same argument would work without it, except that $\mathcal{V}$ would not generally be convex-valued. For related existence theorems encompassing other types of dynamic games with state variables, see Goldman (1980) and Harris (1985).

for any $\lambda > 0$. Consequently, there is no relationship between initial assets and the exercise of self-control.

When $B > 0$, non-uniformity is possible. Indeed, we have computed the equilibrium correspondence for a large collection of parameter values satisfying our assumptions, and always find non-uniformity for intermediate values of β. (See the Online Appendix.) Because non-uniformity implies that the ability to exercise self-control varies with wealth, these cases are of primary interest to us. The central issue we wish to explore is whether self-control is systematically more prevalent at high or low asset levels.

Let $X(A)$ denote the highest level of continuation assets sustainable at A. The closed-graph property of Proposition 1 guarantees that $X(A)$ is well-defined and usc, and a familiar single-crossing argument tells us that it is nondecreasing. Note that $X(A)$ isn't necessarily the value-maximizing asset choice; it could be higher. The following intuitive result tells us that the function X completely characterizes self-control.

PROPOSITION **2.** (i) *Self-control is possible at A if and only if $X(A) > A$.*

(ii) *Strong self-control is possible at A if and only if $X(A') > A'$ for all $A' \geq A$.*

(iii) *There is a poverty trap at A if and only if $X(A') < A'$ for all $A' \in (B, A]$.*

(iv) *There is uniformity if and only if $X(A) \geq A$ for all $A \geq B$, or $X(A) \leq A$ for all $A \geq B$.*

Parts (i) and (iv) are obvious. Part (iii) follows from the additional observations that X is nondecreasing and usc. To prove part (ii), we consider the problem of maximizing $\sum_{t=0}^{\infty} \delta^t u(c_t)$ starting at some initial asset value A subject to feasibility ((1), (2), and (5)) plus the additional constraint that $A_{t+1} \leq X(A_t)$. Compared to an alternative path that starts with $X(A)$ and continues along a trajectory that yields $H(X(A))$ (which we know is sustainable), this constrained-value-maximizing path yields a weakly higher value with weakly lower first-period consumption, and hence achieves a weakly higher payoff. Because this statement is true everywhere along the constrainted-value-maximizing path, that path is necessarily sustainable. And because the Ramsey path involves unbounded accumulation (under our assumptions), so too must the constrained-value-maximizing path (provided $X(A') > A'$ for all $A' \geq A$). For details, see the Appendix.

This proposition will be instrumental in visualizing the proof of our main result. It is worth emphasizing that, when (ii) holds, the value-maximizing equilibrium involves unbounded accumulation. That is noteworthy because value-maximization identifies the most attractive equilibrium from a long-run welfare perspective.[26]

[26] See the discussion of the "long-run" criterion for quasi-hyperbolic discounting in Bernheim and Rangel (2009).

6. INITIAL ASSETS AND SELF-CONTROL

Why might asset levels affect an individual's ability to exercise self-control through contingent self-reinforcement strategies, as discussed in the psychological literature? That ability necessarily depends on the differential consequences of abstaining and indulging, and hence on the severity of the consequences that would follow an impetuous act. One simple intuition is that those consequences are potentially more severe when the individual has more assets, and hence more to lose. For instance, the individual cannot decumulate assets from level B (by assumption), but decumulation may be an equilibrium path from $A > B$. Accordingly, it might be possible to accumulate assets by "threatening" decumulation as a contingent consequence starting from A, but not from B. Unfortunately, this simple intuition does not govern behavior more generally. As we have already seen (Figure 1, panel (B)), the worst possible credible self-punishment does not become monotonically worse (even with suitable renormalization) as assets grow. So, to understand the relationship between assets and self-control, we need to uncover the structure of worst self-punishments, and that is where we will begin.

6.1. **Worst Self-Punishments.** We will show that the worst self-punishments involve a temporary "binge," followed by a return to the best continuation equilibrium value available at the asset level following the binge. Formally, for any $A > B$, let $H^{-}(A)$ be the left-hand limit of H at A (which is well-defined because H is nondecreasing; see Lemma 16 in the Appendix).

PROPOSITION **3.** (i) *The worst equilibrium value at any asset level A is implemented by choosing the smallest possible equilibrium continuation asset at A; call it Y. Moreover, if $Y > B$, the associated continuation value V satisfies $V \geq H^{-}(Y)$.*

(ii) *This equilibrium value can be generated by an equilibrium path, possibly involving public randomization, which entails a return to the best continuation equilibrium after at most two periods.*

For an intuitive explanation of part (i), see Figure 2. Imagine that, following a deviation to continuation asset level A^d in period $t-1$ (not shown in the figure), the equilibrium prescribes continuation asset level A' and a period-$(t+1)$ continuation value $V' < H^{-}(A')$ (as shown in the figure). From the perspective of period t, this self-punishment involves a payoff of $\overline{P} \equiv u\left(A^d - \frac{A'}{\alpha}\right) + \beta\delta V'$ and a value of $\overline{V} \equiv u\left(A^d - \frac{A'}{\alpha}\right) + \delta V'$. The figure shows the iso-payoff curve $u\left(A^d - \frac{A}{\alpha}\right) + \beta\delta V = \overline{P}$ and the iso-value curve $u\left(A^d - \frac{A}{\alpha}\right) + \delta V = \overline{V}$. Both are upward-sloping (with higher payoff and value to the northwest), but the iso-value curve is flatter than the iso-payoff curve, because value places more weight on the future than payoff does. The iso-payoff curve necessarily remains above the $L(A)$-curve; for if it passed below L at some asset level A'', the individual could profitably deviate by playing A'' rather than A' after choosing A^d.

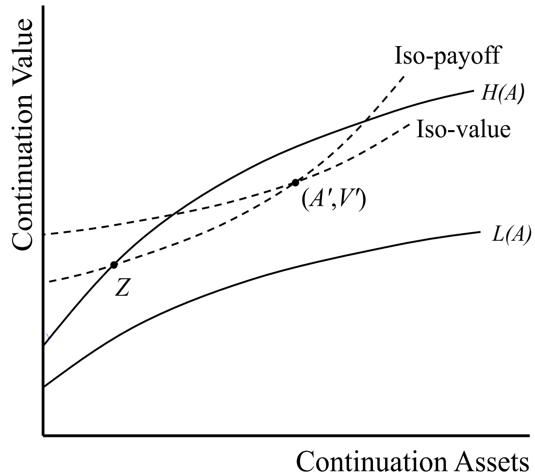

FIGURE 2. OPTIMAL SELF PUNISHMENT.

But then, moving southwest along the iso-payoff curve both preserves incentive-compatibility *and* reduces value. Consequently, the value-minimizing equilibrium continuation (and hence the worst self-punishment) must lie at a point on the upper envelope of the equilibrium value set, such as Z in the figure.

The upper envelope, in turn, is made up of points of the form $H^{-}(Y)$ (whenever $Y > B$). Part (ii) of the proposition shows how the individual can implement this envelope. We demonstrate that $H^{-}(Y)$ is a convex combination of $H(Y)$ and the value of another equilibrium that implements the lowest continuation asset *again*, this time at Y, and then follows up with the highest possible continuation value thereafter. This convex combination is achieved by publicly randomizing across the two equilibrium values. Unpacking the randomization, behavior looks like this: the individual binges for one, possibly two periods, and then returns to the best possible continuation value.

Notice that the logic of Proposition 3 hinges on time-inconsistency, which causes the iso-payoff and iso-value curves to diverge. Thus, while optimal self-punishments have a stick-and-carrot structure reminiscent of optimal penal codes for repeated games (Abreu, 1988), that structure appears here, and in somewhat more complex form, for different reasons.

When worst self-punishments are used to support the choices that achieve $H(\cdot)$, the resulting equilibrium has a natural behavioral and psychological interpretation. The individual sets a personal standard of behavior, which prescribes a level of saving for each asset level A (and, if followed, yields value $H(A)$). If she fails to meet this standard, she self-punishes. The punishment involves high consumption for one period, possibly two, which is disagreeable to the agent as viewed from the vantage point of her original deviation. But the proposition also states that she rededicates herself to her personal standards after at most two periods. In summary, a deviation causes our individual to "fall off the wagon," after which she returns to making choices consistent with achieving her best personal standards, resulting in the value H.[27]

The optimal self-reinforcement we've identified has a qualitatively simple structure, and so cannot be ruled out on the grounds of complexity or lack of realism. Furthermore, simpler punishments may have other problematic features: e.g., Markov punishments carry the unrealistic implication that self-punishments are permanent.

6.2. **The Relationship Between Wealth and Self-Control.** The possibility (illustrated in Figure 1) that an increase in wealth can render the worst self-punishment less severe, even when renormalized for asset levels, leads to the nihilistic suspicion that there might be no general connection between wealth and self-control. Nevertheless, in every numerical example we have examined (including those in the Online Appendix), either self-control is uniform, or there is a poverty trap within which self-control through credible self-reinforcement is infeasible and an asset threshold above which self-control and unbounded accumulation become possible (as in Figure 1). One can certainly imagine other patterns — e.g., that self-control is possible at low asset levels, but not at intermediate or high levels, or that it recurs intermittently as assets grow. Indeed, because Markov equilibria can have a periodic structure (see Section 7), one might conjecture that similar patterns could arise for all equilibria. Yet our main proposition rules these alternatives out, and establishes that the pattern shown in our numerical examples always prevails:

PROPOSITION **4.** *In any non-uniform case:*

(i) *There is $A_1 > B$ such that every $A \in [B, A_1)$ exhibits a poverty trap.*

(ii) *There is $A_2 \geq A_1$ such that every $A \geq A_2$ exhibits strong self-control.*

Notice that the proposition does not establish the existence of a *unique* asset threshold above which self-control is possible and below which it is not — formally, a *single* point at which $X(A)$ crosses the 45^0 line. A demonstration of this stronger property is hindered in part by the possibility that worst punishments can move in unexpected ways as assets rise.[28] From this perspective, the fact that after a finite threshold all such crossings must cease (as the proposition asserts) appears surprising, and the remainder of this section is devoted to an informal exposition of the proof.

[27] It should be noted that the best standards may not be good enough, in the sense of achieving self-control. As Proposition 4 will show, this may be true of individuals with assets close to the lower bound B, while for others with higher wealth a return to H will *also* mean a return to self-control. In both cases, Proposition 3 holds as stated.

[28] However, despite an extensive search, we have not found a numerical example with multiple crossings.

6.3. **A Sketch of the Proof of the Main Proposition.** Although borrowing constraints destroy scale-neutrality of the equilibrium sets, variants of scale-neutrality survive. One useful variant appears below as Observation 2. To state it, we define an asset level $S \geq B$ as *sustainable* if there exists an equilibrium that permits indefinite maintenance of S. It is important to appreciate that strict accumulation may not be possible at a sustainable asset level, and more subtly, an asset level that permits strict accumulation need not be sustainable.[29]

OBSERVATION **2.** *Let $S > B$ be a sustainable asset level. Define $\mu \equiv S/B > 1$. Then for any initial asset level $A \geq B$, if continuation asset A' can be supported as an equilibrium choice, so can the continuation asset $\mu A'$ starting from μA.*

To understand this result, first think of S as a new lower bound on assets. Then, given the hometheticity of utility together with the linearity of the accumulation technology, Observation 2 would obviously hold. Because S is not actually a lower bound, we must also consider deviations to asset levels below S (which have no scaled-down counterparts when B is the lower bound). Nevertheless, Lemma 8 in the Appendix shows that the continuation $L(A)$ deters such deviations.

We use Observation 2 to prove part (i) of the proposition. Recall that $X(A)$ is the largest equilibrium continuation asset level at A. By Proposition 2, we need to show that there is an asset level $A_1 > B$ such that $X(A) < A$ for all $A \in (B, A_1)$. In the formal proof, we rule out the possibility that $X(A)$ wiggles back and forth across the 45^0 line ever-more-rapidly as $A \downarrow B$. Consequently, if part (i) is false, there must be $M > B$ such that $X(A) \geq A$ for all $A \in [B, M]$. Figure 3 illustrates this case.

Because the case is non-uniform by assumption, there is some A^* at which self-control fails, and hence (by Proposition 2) $X(A^*) < A^*$. Let S be the supremum value of assets over $[B, A^*]$ for which $A \in [B, S]$ implies $X(A) \geq A$. Plainly, we must have $X(S) = S$,[30] from which it follows (with an additional argument) that S is sustainable.[31]

Now Observation 2 implies that $X(A)$ must exceed A just to the right of S: for some A' close to B, just scale up $X(A')$ to $\mu X(A')$ at $\mu A'$, where $\mu \equiv S/B$. Because the definition of S requires $X(A) < A$ just to the right of S, we have a contradiction. It therefore follows that $X(A) < A$ for all A close to B, which means there is a poverty trap, as claimed.

---

[29] With strict asset growth, the individual potentially has more to lose than with asset maintenance; hence the former may be easier to sustain.

[30] We cannot have $X(S) < S$, because X is non-decreasing, and we cannot have $X(S) > S$, because then S would not be the supremum of the set mentioned in the text.

[31] In particular, the payoff from a constant asset trajectory is no smaller than one that remains constant for a single period and never subsequently exceeds its original level. Because incentive compatibility must hold for some trajectory of the second type (given $X(S) = S$), it must hold for the first.

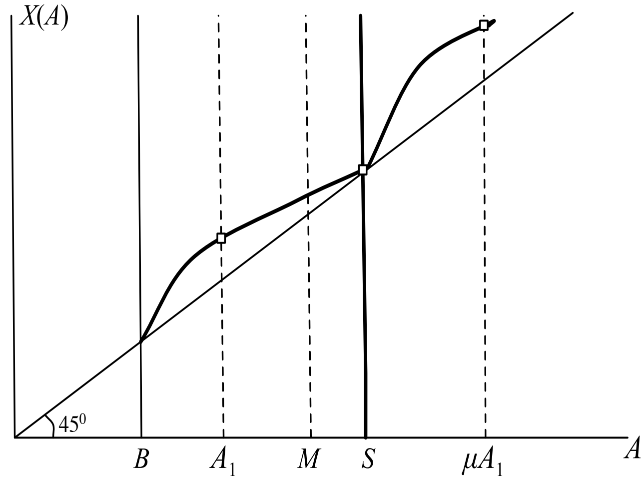

FIGURE 3. ESTABLISHING THE EXISTENCE OF A POVERTY TRAP.

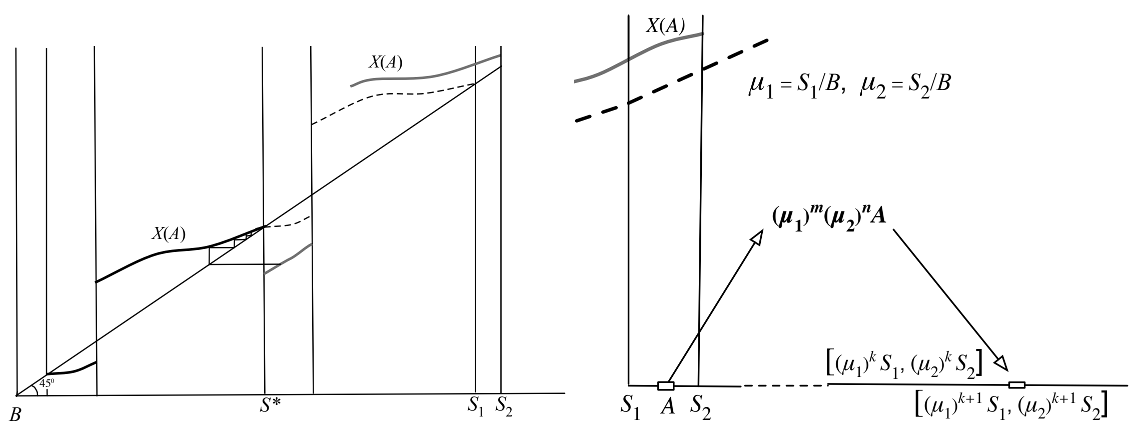

(A) Constructing the interval $[S_1, S_2]$ (B) Double scaling

FIGURE 4. THRESHOLD FOR STRONG SELF-CONTROL.

Next, we explain part (ii) of the proposition. By nonuniformity, there is certainly some value of A for which $X(A) > A$. If the same inequality holds for all $A' > A$, then by Proposition 2 (ii), strong self-control is established for A and all $A' > A$. So the case that we need to address involves $X(A') \leq A'$ for some $A' > A$. Consider Figure 4, panel (A). Focusing on the first zone over which $X(A) > A$, let S^* be the first asset level thereafter for which $X(A) = A$. As in our explanation of part (i), S^* is sustainable.

By Observation 2, the choices $X(A)$ on $[B, S^*]$ can be scaled up and replicated as equilibrium choices on $[S^*, S_1]$, where S_1 bears the same ratio to S^* as does S^* to B.[32] Figure 4 shows these choices as the dotted line within the domain $[S^*, S_1]$. Because there is a poverty trap near B, the dotted line lies below the 45^0 line just to the right of S^*. However — and this is at the heart of the argument below — that line does *not* coincide with $X(A)$ on $[S^*, S_1]$.

In the full proof, we show that worse punishments (in relative terms) are available near S^* than near B. Just to the right of S^*, one can constuct equilibria that dip into the zone to the left of S^*, and then accumulate along $X(A)$ back towards S^* (as indicated by the "stairsteps" in the figure). Because these choices — shown by the solid line to the right of S^* in Figure 4 — favor current consumption over the future, they generate even lower *values* than the equilibria that simply decline to S^*, but they earn high enough *payoffs* to be implementable. These lower values more effectively forestall deviations at even higher asset levels, and in this way greater punishment ability percolates upward from S^*. As a result, for asset levels close to S_1, the incentive constraints are relaxed and higher levels of continuation assets are implementable (see the solid line segment in this region, which lies above the dotted line). In particular, in addition to being sustainable like S^*, S_1 also permits accumulation: $X(S_1) > S_1$.

This argument implies that there is an interval just above S_1, call it (S_1, S_2), over which (a) $X(A) > A$, and (b) both S_1 and S_2 are sustainable. Part (a) follows because $X(S_1) > S_1$ and X is nondecreasing. Part (b) follows from the fact that assets just to the right of S_1 were "almost sustainable" to begin with (by virtue of Observation 2); they become sustainable given the additional punishment power that percolates upward from S^*.

Panel B of Figure 4 focuses on the interval (S_1, S_2) and higher asset levels. The following property, stated and proven formally as Lemma 19 in the Appendix, makes the key step:

OBSERVATION **3.** *Suppose that* S_1 *and* S_2 *are both sustainable, and that* $X(A) > A$ *for all* $A \in (S_1, S_2)$. *Then there exists* $\hat{A}$ *such that* $X(A) > A$ *for all* $A > \hat{A}$.

The proof of this observation is illustrated in panel B. Define $\mu_i = S_i/B$ for $i = 1, 2$. Then for all positive integers k larger than some threshold K, the intervals $(\mu_1^k S_1, \mu_2^k S_2)$ and $(\mu_1^{k+1} S_1, \mu_2^{k+1} S_2)$ *must overlap*. It is easy to see why: $\mu_2^k S_2$ is just $\mu_2^{k+1} B$, while $\mu_1^{k+1} S_1$ is $\mu_1^{k+2} B$, and for large k it must be that μ_2^{k+1} exceeds μ_1^{k+2}. Thus, we can generate any asset level $A > \mu_1^K S_1$ by simply choosing an integer $k \geq K$, an integer m between 0 and k, and $A' \in (S_1, S_2)$ so that $A = \mu_1^m \mu_2^{k-m} A'$. But $X(A') > A'$, so repeated application of Observation 2 proves that $X(A) > A$, which gives us Observation 3.

[32] The actual proof becomes considerably more complex at this point. Briefly, the domain of interest is not exactly $[S^*, S_1]$, but an interval of the form $[S_{**}, S_1]$, where S_{**} *might* coincide with S^* but generally will not. (We proceed here on the assumption that S_{**} does coincide with S^*.) There are several associated complications, and the interested reader is referred to the Appendix not just for the formalities, but also for further intuitive discussion.

Part (ii) of Proposition 2 follows immediately: because $X(A) > A$ for all A sufficiently large, the required threshold A_2 must exist.

6.4. **Renegotiation-Proofness.** A natural question is whether the individual would follow through on worst self-punishments; that is, whether they are "renegotiation-proof." In this context, with one agent "alive" at a time, it is not entirely clear how renegotiation would occur. One possibility, discussed in the psychology literature, is that an individual may initially arrive at an equilibrium strategy by modeling others.[33] In that event, any subsequent effort to change that strategy may be viewed by later selves as a deviation. Ainslie (1975) certainly recognizes this point; he notes that a decision to call off a "private side bet" might lead the individual to "perceive the bet as having been lost," and thereby jeopardize "the credibility of any similar private side bets."

Ainslie (1975) nevertheless describes one potentially feasible form of renegotiation in the context of his shoe-shining example: the individual can modify his rule for shining shoes as long as he does not do so "just before he was due to shine them again." The principle appears to be that one is always free to revise a personal rule, but not for the current period; to avoid confounding revisions and deviations, any changes must be arms-length and limited to plans for subsequent behavior. Plainly, a continuation equilibrium that reverts to $H(\cdot)$ after a single period is immune to revisions according to this criterion; indeed, it satisfies a strong form of "renegotiation-proofness" (given that the continuation is unimprovable within the entire equilibrium set). Moreover, Ainslie's reasoning arguably implies that reasonable self-punishments *must* have this structure, else they would be revised.

In this context, it is therefore noteworthy that, with an appropriate interpretation of cases in which the continuation value lies above $H^{-}(A)$ but strictly below $H(A)$, the worst self-punishment equilibria have this property. To see this, we adopt a slightly different interpretation from the one offered in Proposition 3. Under this interpretation, the agent binges for one period only, and provided there is some noise in asset returns or she can arrange small side bets (not necessarily fair to her, so that any risk-neutral second party would accept such a bet), she can return to the highest continuation value function in the very next period.[34]

Other notions of renegotiation may also be applicable to this setting. Any approach to this issue requires (a) an understanding of the set of agents that might engage in renegotiation at any

[33] For example, a classic experiment shows that "children's patterns and magnitude of self-reinforcement closely matched those of the model to whom they had been exposed. Adults generally served as more powerful modeling stimuli than peers in transmitting self-reinforcing responses" (Bandura and Kupers (1964); see also Bandura (1971, 1976)).

[34] As an example, the equilibrium strategy might specify that, in addition to consuming slightly more than $A - \frac{Y}{\alpha}$, the individual also makes a small wager with another party, which has the effect of leaving her with continuation assets of either Y or $Y - \varepsilon$, with appropriate probabilities. The wager need not be fair. One can also smooth out expected continuation values by introducing a small amount of noise in the return α to assets. While we do not formally consider a stochastic model, the same arguments in the proof of Proposition 3 go through.

date, and (b) a description of the various paths over which they might negotiate. For part (b), we follow Bernheim and Ray (1989) and Farrell and Maskin (1989): we consider the collection of all continuation paths for a given equilibrium and ask if it is internally consistent, in the sense that the renegotiating agents would never wish collectively to switch from one of these paths to another.[35] This requirement is known as *weak renegotiation-proofness* (WRP). It seems particularly appropriate in the present context if one thinks of the individual as inheriting behavioral principles (i.e., an equilibrium) by modeling mentors, but retaining the ability to choose (and renegotiate) a starting point within that equilibrium.

We will focus on an equilibrium that achieves $H(\cdot)$ from all asset levels. This equilibrium exhibits accumulation whenever self-control is feasible, and unbounded accumulation whenever strong self-control is feasible (see Lemma 15 in the Appendix). Consider two selections $X^*(A)$ and $Y^*(A)$ from the equilibrium asset choice correspondence, and two "plans" built from these selections. Plan 1 applies X^* repeatedly and generates the highest value $H(A)$. Plan 2 applies Y^* *once* at A to obtain the lowest possible equilibrium continuation asset $Y^*(A)$, and then uses a randomization device to switch back to Plan 1 at $Y^*(A)$, or restart Plan 2 at $Y^*(A)$. The randomization probability is chosen so that the expected value at $Y^*(A)$ is exactly equal to the continuation value of the worst punishment at A. By Proposition 3, Plan 2 implements the value $L(A)$ at A.[36]

Initially, we will take the (reasonable) position that an agent must be present to renegotiate, which limits the set of negotiators to the period t self. However, according to one view, the standard model of quasi-hyperbolic discounting is simply a reduced form for a process involving repeated Nash bargaining between two long-lived selves, an impulsive "doer" who cares only about current consumption, and a patient "planner" with standard Ramsey preferences with the parameter β capturing the relative weight assigned to the patient self.[37] Under the WRP concept, our equilibrium set is vulnerable to renegotiation only if there is some history at which *both* the impulsive and patient selves would like to switch from one equilibrium plan to another. Yet that is never the case. Plan 1 provides a higher payoff to the patient self, while Plan 2 provides higher payoff to the impulsive self. Because the two selves do not agree, our equilibrium set is WRP.

If instead we treat each date t self as a single, distinct entity, "renegotiation" becomes a simple question of whether any date t self ever prefers to switch between plans. According to the

[35] Bernheim and Ray (1989) and Farrell and Maskin (1989) actually apply their concept to collections of equilibrium *payoffs*. Given the potentially confusing distinction between payoffs and values in the current setting, we apply the same notion to equilibrium paths or outcomes.

[36] By Proposition 3, if H is continuous at Y and $Y > B$, no randomization is needed and the move back to Plan 1 is deterministic.

[37] The doer-planner model was originally formulated by Thaler and Shefrin (1981); see also Fudenberg and Levine (2006). For the relationship to quasi-hyperbolic discounting, see Bernheim (2009).

incentive constraints that uphold the equilibrium, we have (for Plans 1 and 2 respectively),

(12) $$u\left(A - \frac{X^*(A)}{\alpha}\right) + \beta\delta H(X^*(A)) \geq D(A), \text{ and}$$

(13) $$u\left(A - \frac{Y^*(A)}{\alpha}\right) + \beta\delta V \geq D(A),$$

where $D(A)$ is the "best" deviation payoff at A,[38] and V is the continuation value associated with the worst punishment, obtained from future randomization between the values $H(Y^*(A))$ and $L(Y^*(A))$, as described above. Notice that, when both constraints bind (as they often do), the date t self is indifferent between the two plans, and therefore has no incentive to "renegotiate." Once again, the equilibrium set is WRP. Here, the key insight is that the path used to punish a deviation by the $(t-1)$-self has the flavor of something the *punisher* — that is, the t-self — would like to follow, because it involves greater consumption at date t.

What if renegotiation at date t also encompasses future selves? When the set of negotiating parties expands (without altering the feasible set), successful renegotiation becomes more difficult. Accordingly, such considerations would strengthen our conclusions.[39]

7. MARKOV EQUILIBRIUM

As we have mentioned, most of the literature on quasi-hyperbolic discounting focuses on Markov-perfect equilibria, in which choices depend only on the current asset level A, and not on how that level was reached (e.g., whether the individual exercised restraint from a lower asset level or splurged from a higher asset level). Because a Markov-perfect equilibrium provides no scope for contingent self-reinforcement, it cannot capture the phenomena described in the psychological literature (see Section 2).

That said, our definition of self-control is not specific to contingent self-reinforcement; it only concerns the ability to accumulate assets. It is therefore of interest to ask whether Markov equilibria, as a class, manifest the same patterns highlighted in our main result, Proposition 4. They do not. The following proposition establishes that, unlike the set of subgame-perfect equilibria, the set of Markov equilibria never exhibits "non-uniformity."

---

[38] That is, $D(A)$ is the supremum over all payoffs in which every deviation to an alternative asset choice is "punished" by the lowest equilibrium value available at that asset. The function $D(A)$ is formally defined in the Appendix, where we deal with various technicalities arising from lack of the continuity in the value correspondence.

[39] To illustrate, suppose the renegotiating set at date t consists of the date t self and the date $t+1$ self. Successful renegotiation then requires both of the agents to prefer one available continuation plan over another. Notice that the best value function $H(A)$ is now a WRP equilibrium even when the incentive constraint (13) does not bind (because the date $t+1$ self unambiguously prefers Plan 1 over Plan 2).

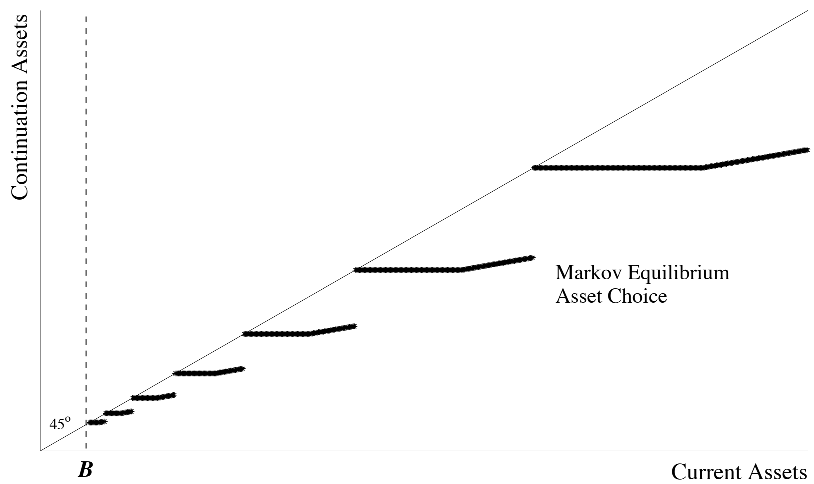

FIGURE 5. MARKOV EQUILIBRIUM WITH PERIODIC MAINTENANCE.

PROPOSITION **5.** *Suppose there exists a Markov equilibrium ϕ with $\phi(A) > A$ for some $A \geq B$. Then there exists a Markov equilibrium ϕ' with $\phi'(A) \geq A$ for all $A \geq B$, with strict inequality if $\beta \neq \beta^* \equiv (1-\delta)/[\delta(\alpha-1)]$.*

The proof of this result proceeds as follows. It is easy to identify a value β^* at which $\phi(A) = A$ is a Markov equilibrium, and to show that, for $\beta > \beta^*$, there are linear Markov equilibria with strict accumulation. Hence, for those values, the Markov set is uniform. Focusing then on $\beta < \beta^*$, we note that, if there is an equilibrium with $\phi(A) > A$ for some $A > B$, we can create a scaled-down version ϕ' such that $\phi'(B) > B$. It is easy to check that ϕ' must be non-decreasing. Accordingly, if it ever passes below the 45^o line, there must be some asset level S at which $\phi'(S) = S$. But then, from S, the individual could decumulate slightly, and count on his future selves to accumulate back to S. With $\beta < \beta^*$, that alternative necessarily yields a higher payoff than choosing S, which contradicts the supposition that ϕ' is an equilibrium. Therefore, the Markov set must also be uniform for $\beta < \beta^*$. See the Appendix for details.

Despite this negative result, it is worth noting that Markov equilibria can give rise to asset traps. Specifically, in a (weakly) decumulating Markov equilibrium,[40] there may be levels of A at which asset *maintenance* is possible. Figure 5 shows one such case (which reflects actual numerical examples). To see how such equilibria arise, suppose that β is relatively small, so that starting just above the lower bound B, the individual exhausts liquid assets. As the initial asset level increases, the subsequent (proportional) decline in assets is greater, and hence the (normalized) value of the trajectory is lower. Consequently, as long as β is not too small, there comes a point (call it S) at which asset maintenance yields equivalent value, in which case one can construct

[40] It can be shown that weakly decumulating Markov equilibria exist for $\beta < \beta^*$.

a Markov equilibrium in which S is chosen from S. Moreover, for asset levels above S, the game "scales up" by the factor $\mu = S/B$, just as described in Observation 2. Applying this logic recursively, we generate an infinite sequence of asset levels $\{B, S, \mu S, \mu^2 S, \ldots\}$ for which maintenance is possible, each of which acts as an asset trap.[41]

8. Some Additional Implications

In this section, we explore the broader implications of our analysis for behavior and policy (aside from the benefits of "priming the pump" for those caught in the poverty trap). We touch on four topics: the effect on saving of easier access to credit; the demand for external commitment devices; the design of accounts to promote saving; and the observed variation in marginal propensities to consume from wealth across classes of resource claims.

8.1. **The Effects of Easier Access to Credit.** The decline in saving rates among U.S. households during the latter part of 20th century is sometimes attributed, at least in part, to institutional developments that progressively improved access to credit.[42] Conventional theory predicts that more abundant (and cheaper) credit reduces aggregate saving. In contrast, our model has more nuanced implications.

Conceptually, comparative statics with respect to the level of the borrowing constraint (B) are straightforward. Although the constraint destroys scale neutrality, a change in B simply rescales the equilibrium set. Thus, we can reinterpret Proposition 4 as showing that there are two values, μ' and μ'', with $1 < \mu' \leq \mu'' < \infty$, such that a poverty trap exists whenever $A/B < \mu'$, while unbounded accumulation is possible whenever $A/B > \mu''$.

It follows that the effect on saving of relaxing the credit limit depends on the level of initial assets. The direct effect of such a relaxation is to reduce B, e.g., from B_1 to $B_2 < B_1$, thereby increasing the ratio A/B for everyone. That change may allow an individual to escape the poverty trap (i.e., if $A/B_1 < \mu' < A/B_2$), and may even enable him to accumulate assets indefinitely (i.e., if in addition $A/B_2 > \mu''$). However, there is also a downside to easy credit: those whose assets remain below $\mu' B_2$ will slide even further into poverty. In any given context, either the first effect

[41] Yet another alternative is to focus on subgame-perfect equilibria supported by Markov-reversion. Because this alternative involves permanent punishments (in which the individual never "climbs back on the wagon"), it strikes us as unappealing. That said, reversion to a cyclical Markov equilibrium, such as the one exhibited in the text, cannot give rise to the patterns highlighted in Proposition 4. Specifically, it is easy to show that any equilibrium of this type starting from an asset level in $[\mu^k S, \mu^{k+1} S]$ for any $k > 0$ can be scaled down to a starting point in any $[\mu^j S, \mu^{j+1} S]$ with $j < k$. It follows that the equilibrium set supportable by reversion to a cyclical Markov equilibrium is also cyclical. That said, there may also be strictly decumulating Markov equilibria, and our computational analysis has identified cases in which reversion to those equilibria supports the qualitative pattern highlighted in Proposition 4; see the Online Appendix. However, whether reversion to the *worst* Markov equilibrium generally yields a counterpart to Proposition 4 is an open question.

[42] See, e.g., Bacchetta and Gerlach (1997), Ludvigson (1999), Parker (2000) and Glick and Lansing (2011).

or the second may be more prevalent. Notably, Karlan and Zinman (2010) report the results of a field experiment showing that expanded access to costly consumer credit in South Africa on average improved economic self-sufficiency, intra-household control, community status, and overall optimism. Similarly, Dobbie and Skiba (2013) show that larger payday loans lead to lower rates of default.

8.2. **The Demand for Commitment Devices.** As noted in Section 1, a demand for precommitment has been documented for poor households in developing countries. However, there is surprisingly little evidence that this demand is more widespread,[43] and as a result nagging doubts about the importance of (sophisticated) time inconsistency persist. Skeptics wonder why, if time inconsistency is so prevalent, the market provides few commitment devices, and why unambiguous examples in the field are so difficult to find.

Our analysis provides a potential resolution to this puzzle. Because full-precommitment is neither possible nor desirable (due to the value of flexibility), people must rely to some extent on internal mechanisms for self-control. Significantly, the use of external commitments may undermine the efficacy of those internal mechanisms by rendering ineffective the individual's personal rules. As an illustration, consider an external commitment that "locks up" certain assets in an illiquid savings account. The direct effect of that commitment is to increase B, the lower bound on net worth, say from B_1 to $B_2 > B_1$. The impact on saving is then the same as for a tightening of the credit constraint. In particular, defining μ' and μ'' as above, if $A/B_1 > \mu'' > A/B_2$, the external commitment would render unbounded accumulation infeasible, and if $\mu' > A/B_2$, it would induce the individual to deplete all of her other (liquid) assets. Accordingly, people may have powerful reasons to avoid (partial) external commitments.

In our model, the people who value external commitments are those who are asset-poor relative to their credit limits. The asset-rich would rather save on their own. By the same reasoning, if we assume B is a constant fraction of permanent income, the *income*-rich would exhibit a desire for external commitment, while the income-poor would prefer to rely on internal mechanisms. To be sure, the income-rich may also be asset-rich, so that the net effect is ambiguous. Nevertheless, the theory yields empirical predictions that are, in principle, testable.

8.3. **Designing Accounts to Promote Saving.** Some policies encourage saving by providing special accounts for specific purposes, such as retirement, education, or medical expenses. Virtually all such accounts entail commitments, the nature of which differs considerably across programs. For example, the degree to which savings are "locked up" until retirement varies across

[43] Studies documenting a demand for precommitment in *developed* countries are scarce. Exceptions include Ariely and Wertenbroch (2002) on homework assignments, Beshears et al. (2011) on commitment savings devices in the U.S., and Houser et al. (2010) on a laboratory experiment in which subjects gain relevant experience. Gine, Karlan, and Zinman (2010) write that "there is little field evidence on the demand for or effectiveness of such commitment devices." For recent surveys, see Bryan, Karlan, and Nelson (2010) and DellaVigna (2009).

pension programs. For social security and many private plans (especially of the defined benefit variety), lock-up is absolute. IRAs impose a moderate early withdrawal penalty of 10%. For 401(k)s and 403(b)s, the same 10% penalty applies, but employers can also impose additional restrictions and, as an example, often limit early withdrawals to funds contributed by the employee. After retirement, the lock-up continues in a modified form for social security and many private plans: income is paid out at a specified rate, or investment in annuities is mandated. In contrast, IRAs and many other private plans effectively unlock the funds at retirement, making them highly liquid. In addition, participants in retirement savings programs often precommit to contributions. For social security and many private plans, contributions are inflexible. For 401(k)s and 403(b)s, they are adjustable, but only with a significant delay (e.g., a pay period). Only IRA contributions are fully flexible.

Our analysis potentially sheds light on the ways in which the commitment features of special savings accounts affect saving. Caution is warranted, inasmuch as our model lacks a retirement period, and therefore maps imperfectly to a realistic life-cycle planning problem. Still, one can interpret it as providing a stylized representation of saving decisions during the accumulation phase of the life cycle.

An asset lock-in has both an upside and a downside. The upside is that it can compensate for the absence of self-control when assets are low; the downside is that it can undermine internal self-control mechanisms when assets are high. Because these effects materialize at different asset levels, it is in principle possible to design programs that capitalize on the upside while avoiding the downside, for example by locking up all funds until some personally chosen asset target is achieved, and then removing the lock (irreversibly), making all funds liquid. Pilot programs with such features have indeed been tested in developing countries.[44]

Formalizing the preceding intuition is not entirely straightforward. In our simple model, lock-up would prevent people with low assets from decumulating, but it would not necessarily enable them to employ personal rules that support contributions to the account in the first place (and might undermine that ability). Furthermore, there is an obviously superior policy alternative that achieves the Ramsey outcome: require participants to select contributions and withdrawals one period in advance.

Our intuition concerning account design is nevertheless borne out in a slightly more elaborate model that incorporates preference shocks (e.g., reflecting transient needs associated with illnesses requiring costly medical care). In such cases, an exclusive reliance on external commitment is problematic. Suppose in particular that flow utility is given by

$$u(c, \eta) = \eta \frac{c^{1-\sigma}}{1-\sigma},$$

[44] See Ashraf, Karlan, and Yin (2006), as well as Karlan et al. (2010).

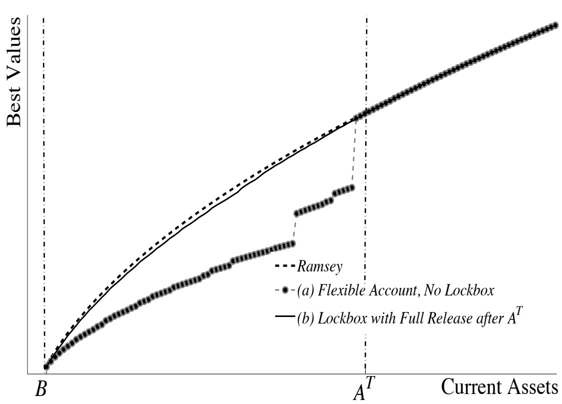

(A) Overall

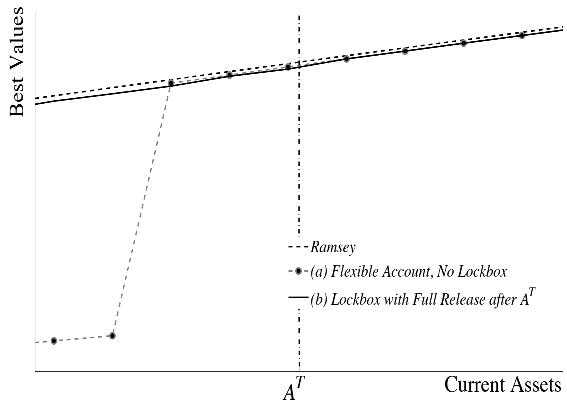

(B) Zoom

FIGURE 6. EQUILIBRIUM VALUES: LOCKBOX WITH UNLOCKING.

where η is an iid random variable realized at the outset of each period. If the distribution of η encompasses sufficiently low values, the individual will contribute to a lock-up account in some states of nature even when assets are low. Moreover, committing to contributions one period in advance sacrifices the individual's ability to condition consumption on the realization of η, and consequently does not deliver the generalized Ramsey solution.

Due to the complexity of the extended model, we analyze it computationally; see the Online Appendix for details. Numerical solutions generally confirm our intuition. Figures 6 and 7 depict results for an illustrative case.[45] Figure 6, panel (A), shows the highest achievable equilibrium value as a function of initial assets for two policy regimes: in the first, labeled (a), only a standard savings account is available; in the second, labeled (b), the individual has access to an account that locks up principal but unlocks once an appropriate target is reached. For policy (a), the highest value jumps upward when the individual has sufficient assets to save on her own. We've chosen the lockbox target (labeled A^T) to be slightly higher than the jump point: with a lower target, she would deplete her assets once the account is unlocked.

At "low" asset levels below the jump point, the individual fares better with the lock-up account than with the standard account. For asset levels above the lockbox target, the two curves must obviously coincide (because the account is unlocked). Notice that the lock-up account allows the individual to achieve values close to those of the generalized Ramsey solution (also shown). In our example, it doesn't quite reach that ideal; see panel (B) (which magnifies the value functions around the asset target).

[45] For the parameters, we take $A = 1.3$, $\sigma = 0.5$, $\delta = 0.8$, and $\beta = 0.4$. The taste shock η takes two values, 0.8 (with probability 0.3) and 1.1 (with probability 0.7). The horizontal axis starts at $B = 0.5$, and υ is taken to be a tiny number.

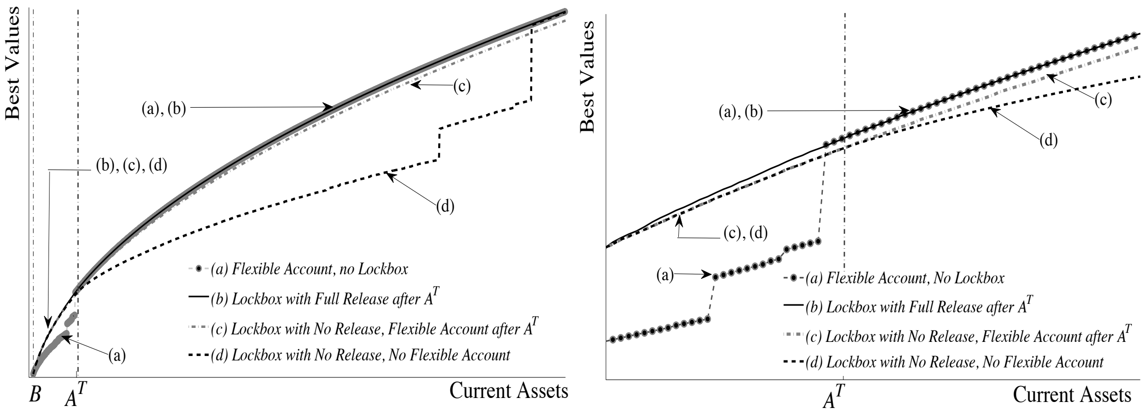

(A) Overall (B) Zoom

FIGURE 7. ALTERNATIVE LOCKBOX REGIMES.

Critically, *both* the lock-up *and* the subsequent release are important. Figure 7 replicates the highest equilibrium value functions for the standard and lock-up accounts, and adds lines for two additional policy regimes. In both, principal in the special account remains locked up forever, there is no release, but interest is always accessible. Before the asset target is reached, all savings are channeled into the lockbox and added to the existing principal. With policy (d), there is a finite asset target, after which additional savings are not put into the lockbox, but instead stored in a fully flexible, standard account with access to principal and interest, while the principal in the lockbox account continues to remain inaccessible. With policy (c), a flexible account is not accessible at any asset level and all savings are treated as contributions to the lockbox account.[46] Figure 7 shows that policy (c) reduces equilibrium value relative to policy (b) (and panel (B) magnifies the region around the jump point for clarity); though it promotes saving, it does not allow the individual's self-control mechanisms to take over. She fares even worse with policy (d); though it helps her achieve the target A^T, its subsequent effect is to scale up the credit constraint from B to A^T, thereby creating a new "poverty trap" at that level. In short, though a permanent lockbox is helpful, only one that is eventually dismantled, such as policy (b), can reap the benefits of unleashing effective personal rules.

8.4. **Asset-Specific Marginal Propensities to Consume.** It has been observed that marginal propensities to consume might differ across classes of resource claims (e.g., between income flows and liquid assets); see Hatsopoulos, Krugman, and Poterba (1989), Thaler (1990) and Laibson (1997). Our model provides a potential explanation. Recall from Section 2.1 that the

[46] Because it is easier to accumulate assets in a lockbox account than in a standard account, the lack of access to a standard account is not consequential.

lower bound on assets (B) may reflect limitations on the ability to borrow against future earned income, so that, e.g., $B = \lambda\frac{\alpha}{\alpha-1}y$ for some $\lambda \in (0, 1)$. An increase in financial assets then leaves B unchanged and increases A_t/B, potentially enhancing self-control according to our main result. As a consequence, the marginal propensity to consume out of an unforeseen change in financial assets could be quite low. In contrast, an equivalent increase in human capital, $\frac{\alpha}{\alpha-1}y$, leaves B/y unchanged, so A_t/B falls, potentially undermining self-control. Therefore, the marginal propensity to consume from an unforeseen change in permanent income could be quite high. More generally, as long as B is an *increasing* function of permanent income, the marginal propensity to consume will tend to be higher for permanent income than for liquid assets. This provides a new perspective on the "excess sensitivity" of consumption to income.

9. Conclusion

If people fundamentally differ in their capacities for exercising self-control, then those who are more impulsive are more likely to deplete their assets. Yet as we emphasize, there may also be a strong feedback effect from indigence to poor self-control. Indeed, our main result shows that poverty can undermine the ability to exercise self-control through contingent self-reinforcement, while wealth can enhance that ability. While there are many other explanations for the persistence of poverty, the endogenous failure of self-control emerges as a potentially important contributory factor.

Our analysis leaves a number of important questions unanswered. Some are technical; e.g., whether the two thresholds identified in our main result can differ, creating an intermediate zone in which the individual can avoid depleting her assets, but cannot accumulate past some critical asset level. Others open issues involve policy implications. For instance, while our analysis points to some intriguing relationships between external commitment devices (such as fixed deposit or lockbox retirement accounts) and the efficacy of contingent self-reinforcement, a thorough analysis of that topic is beyond the scope of the present paper.

Finally, at the broadest level, this paper is a contribution to the behavioral economics of poverty, a subject that has recently been the focus of much empirical work but little theoretical analysis. Self-control is one of several pertinent behavioral considerations; others include internally and socially generated aspirations, the reliance on role models, decisions to acquire detailed knowledge about rates of return from investments in health and education, and other types of informational and psychological distortions that are traceable to the conditions of poverty. Which of these considerations tend to amplify initial conditions, and what types of interventions will promote convergence by nullifying those conditions? Progress towards answering those questions would be of immense significance.

REFERENCES

Ainslie. 1991. "Derivation of "Rational" Economic Behavior from Hyperbolic Discount Curves." *American Economic Review* 81:134–140.

Ainslie, G. 1975. "Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control." *Psychological Bulletin* 82:463–496.

Akerlof, G. 1991. "Proscrastination and Obedience." *The American Economic Review* 81:1–19.

Aleem, I. 1993. "Imperfect Information, Screening, and the Costs of Informal Lending: A Study of a Rural Credit Market in Pakistan." In *The Economics of Rural Organization: Theory, Practice and Policy*, edited by A. Braverman K. Hoff and J. Stiglitz, chap. 7. Oxford University Press (for the World Bank).

Ali, S. N. 2011. "Learning Self-Control." *Quarterly Journal of Economics* 126(2):857–893.

Aliber, M. 2001. "Rotating Savings and Credit Associations and the Pursuit of Self-Discipline." *African Review of Money Finance and Banking* :51–72.

Ambec, S. and N. Treich. 2007. "Roscas as Financial Agreements to Cope with Self-Control Problems." *Journal of Development Economics* 82:120–137.

Anderson, S. and J-M. Baland. 2002. "The Economics of Roscas and Intrahousehold Resource Allocation,." *Quarterly Journal of Economics* 117:963–995.

Appadurai, A. 2004. "The Capacity to Aspire." In *Culture and Public Action*, edited by V. Rao and M. Walton. Stanford, CA: Stanford University Press.

Ariely, D. and K. Wertenbroch. 2002. "Procrastination, Deadlines, and Performance: Self-Control by Pre-Commitment." *Psychological Science* 13:219–224.

Ashraf, N., D. Karlan, and W. Yin. 2006. "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines." *Quarterly Journal of Economics* 121 (2):635–672.

Bacchetta, P. and S. Gerlach. 1997. "Consumption and Credit Constraints: International Evidence." *Journal of Monetary Economics* 40:207–238.

Banarjee, A. and E. Duflo. 2011. *Poor Economics: A Radical Rethinking of the Way to Fight Global Poverty*. Public Affairs.

Bandura, A. 1971. *Social Learning Theory*. General Learning Press: New York.

———. 1976. "Self-reinforcement: Theoretical and Methodological Considerations." *Behaviorism* 4(2):135–155.

Bandura, A. and C. J. Kupers. 1964. "Transmission of Patterns of Self-Reinforcement Through Modeling." *Journal of Abnormal and Social Psychology* 69 (1):1–9.

Banerjee, A. and S. Mullainathan. 2010. "The Shape of Temptation: Implications for the Economic Lives of the Poor." Working Paper 15973, National Bureau of Economic Research.

Basu, K. 2011. "Hyperbolic Discounting and the Sustainability of Rotational Savings Arrangements." *American Economic Journal: Microeconomics* 3 (4):143–171.

Benabou, R. and J. Tirole. 2004. "Willpower and Personal Rules." *Journal of Political Economy* 112:848–886.

Benartzi, S. and R. Thaler. 2004. "Save More Tomorrow: Using Behavioral Economics to Increase Employee Savings." *Journal of Political Economy* 112:S164–S187.

Benhabib, J. and A. Bisin. 2001. "Self-Control and Consumption-Saving Decisions: Cognitive Perspectives." Mimeo., Department of Economics, New York University.

Bernheim, B. D. 2009. "Behavioral Welfare Economics." *Journal of the European Economics Association* 7 (2-3):267–319.

Bernheim, B. D. and A. Rangel. 2009. "Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics* 124:51–104.

Bernheim, B. D. and D. Ray. 1989. "Collective Dynamic Consistency in Repeated Games." *Games and Economic Behavior* 1:295–326.

Bernheim, B. D., D. Ray, and S. Yeltekin. 1999. "Self-Control, Saving, and the Low Asset Trap." URL `http://www.econ.nyu.edu/user/debraj/Papers/selfcontrol.pdf`.

Beshears, J., J. Choi, D. Laibson, and B. Madrian. 2011. "Self Control and Liquidity: How to Design a Commitment Contract." Working Paper, Department of Economics, Harvard University.

Bryan, G., D. Karlan, and S. Nelson. 2010. "Commitment Devices." *Annual Review of Economics* 2:671–698.

de Mel, S., D. McKenzie, and C. Woodruff. 2008. "Returns to Capital in Microenterprises: Evidence from a Field Experiment." *Quarterly Journal of Economics* 123:1329–1372.

DellaVigna, S. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47:315–372.

Dobbie, W. and P. M. Skiba. 2013. "Information Asymmetries in Consumer Credit Markets: Evidence from Payday Lending." *American Economic Journal: Applied Economics* 5:256–282.

Duflo, E., M. Kremer, and J. Robinson. 2011. "Nudging Farmers to Use Fertilizer: Theory and ExExperiment Evidence from Kenya." *American Economic Review* 101 (6):2350–2390.

Farrell, J. and E. Maskin. 1989. "Renegotiation in Repeated Games." *Games and Economic Behavior* 1:327–360.

Fudenberg, D. and D. Levine. 2006. "A Dual-Self Model of Impulse Control." *The American Economic Review* 96 (5):1449–1476.

Genicot, G. and D. Ray. 2006. "Contracts and Externalities: How Things Fall Apart." *Journal of Economic Theory* 131 (1):71–100.

Gine, X., D. Karlan, and J. Zinman. 2010. "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation." *American Economic Journal: Applied Economics* 2:213–225.

Glick, R. and K. Lansing. 2011. "Consumers and the Economy, Part I: Household Credit and Personal Saving." Economic letter, Federal Reserve Bank of San Francisco. URL http://www.frbsf.org/publications/economics/letter/2011/el2011-01.html.

Goldman, S. M. 1980. "Consistent Plans." *Review of Economic Studies* 47:533–537.

Goldstein, M. and C. Udry. 1999. "Agricultural Innovation and Resource Management in Ghana." URL http://www.econ.yale.edu/~cru2//pdf/finalrep.pdf. Mimeo., Yale Economics.

Gugerty, M. 2007. "You Cant Save Alone: Commitment and Rotating Savings and Credit Associations in Kenya." *Economic Development and Cultural Change* 55:251–282.

Harris, C. 1985. "Existence and Characterization of Perfect Equilibria in Games of Perfect Infomation." *Econometrica* 53 (3):613–628.

Hatsopoulos, G., P. Krugman, and J. Poterba. 1989. "Overconsumption: The Challenge to U.S. Economic Policy." American Business Conference.

Houser, D., D. Schunk, J. Winter, and E. Xiao. 2010. "Temptation and Commitment in the Laboratory." Working Paper No. 488, Institute for Empirical Research in Economics, University of Zurich.

Karlan, D., McConnell, S. M., Mullainathan, and J. Zinman. 2010. "Getting to the Top of Mind: How Reminders Increase Saving." Working Paper, Department of Economics, Yale University.

Karlan, D. and J. Zinman. 2010. "Expanding Credit Access: Using Randomized Supply Decisions to Estimate the Impacts." *Review of Financial Studies* 23:433–464.

Kazdin, A. E. 2012. *Behavior Modification in Applied Settings, Seventh Ed.* Waveland Press.

Khera, R. 2011. "Indias Public Distribution System: Utilisation and Impact." *Journal of Development Studies* 47 (7):1038–1060.

Klein, B. and K. B. Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance." *Journal of Political Economy* 89 (4):615–641.

Laibson, D. 1994. "Self-Control and Saving." Mimeo., Department of Economics, Harvard University.

———. 1996. "Hyperbolic Discount Functions, Undersaving and Savings Policy." NBER Working Paper w5635, NBER.

———. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112:443–477.

Ludvigson, S. 1999. "Consumption and Credit: A Model of Time-Varying Liquidity Constraints." *Review of Economics and Statistics* 81:434–447.

Mahoney, M. J., N. G. Moura, and T. C. Wade. 1973. "Relative Efficacy of Self-reward, Self-punishment, and Self-monitoring Techniques for Weight Loss." *Journal of Consulting and Clinical Psychology* 40 (3):404–407.

Mischel, W. 1973. “Toward a Cognitive Social Learning Reconceptualization of Personality.” *Pscyhological Review* 80 (4):252–283.

Moav, O. and Z. Neeman. 2012. “Saving Rates and Poverty: The Role of Conspicuous Consumption and Human Capital.” *Economic Journal* 122:933–956.

O’Donoghue, T. and M. Rabin. 1999. “Doing It Now or Later.” *American Economic Review* 89:103–124.

Parker, J. 2000. “Spendthrift in America? On Two Decades of Decline in the U.S. Saving Rate.” *NBER Macroeconomics Annual* 14:317–387.

Phelps, E. and R. Pollak. 1968. “On Second-Best National Saving and Game Equilibrium Growth.” *Review of Economic Studies* 35:185–199.

Rachlin, H. 1974. “Self-Control.” *Behaviorism* 2:94–107.

Ray, D. 2006. “Aspirations, Poverty and Economic Change.” In *What Have We Learnt About Poverty*, edited by A. Banerjee, R. Bènabou, and D. Mookherjee. Oxford: Oxford University Press.

Rehm, L. P. 1977. “A Self-Control Model of Depression.” *Behavior Therapy* 8:787–804.

Shipton, P. 1992. “The Rope and the Box: Group Savings in the Gambia.” Mimeo., Department of Anthropology, Boston University.

Strotz, R. 1956. “Myopia and Inconsistency in Dynamic Utility Maximization.” *Review of Economic Studies* 23:165–180.

Thaler, R. 1990. “Anomalies: Saving, Fungibility, and Mental Accounts.” *Journal of Economic Perspectives* 4:193–205.

Thaler, R. and H. Shefrin. 1981. “An Economic Theory of Self-Control.” *Journal of Political Economy* 89:392–406.

Udry, C. and S. Anagol. 2006. “The Return to Capital in Ghana.” *American Economic Review* 96:388–393.