

Self-Control, Saving, and the Low Asset Trap

B. Douglas Bernheim
Stanford University

Debraj Ray
New York University

Şevin Yeltekin
Stanford University

May 1999

Abstract

We examine a class of intertemporal consumption allocation models in which time inconsistent preferences create self-control problems. Behavior corresponds to an equilibrium of a game played by successive incarnations of the single decision-maker, and the scope for self-control is defined by the set of subgame perfect equilibria. We study the features of this set computationally. A reasonably robust feature of the equilibrium set is the existence of a threshold asset level, below which no self-discipline is possible (the individual simply consumes assets down to zero), and above which it is possible to sustain positive asset accumulation. We investigate the sensitivity of the equilibrium set to various parameters. For example, we show that when credit becomes more readily available, self-control becomes easier to impose, and fewer individuals are caught in the low asset trap.

1 Introduction

In recent years, a number of economists have questioned the suitability of the life cycle hypothesis for modeling personal saving. Their concerns fall primarily into two categories: issues related to self-control, and issues related to bounded rationality. Our focus in this paper is on the issue of self-control.

One can formalize problems of self-control in a number of different ways. Following Phelps and Pollak (1968), Laibson (1994, 1996, 1997) analyzes a class of models in which problems with self-control arise from time-inconsistent preferences. In contrast to the LCH,

Laibson’s formulation of the intertemporal planning problem assumes that an individual becomes less willing to defer gratification from period t to some period $s < t$ once period t actually arrives. As a result, the individual is typically unwilling to follow through on an optimal intertemporal plan. This approach is motivated by psychological research indicating that rates of time preference are approximately “hyperbolic.” With time-inconsistent preferences, it is natural to model behavior as an equilibrium of a dynamic game played by successive incarnations of the single decision-maker. The scope for exercising self-control is then sharply defined by the set of outcomes that can arise in subgame perfect equilibria.

To date, the literature on saving with quasi-hyperbolic discounting has focused almost entirely on Markov-perfect equilibria (that is, equilibria in which history dependence is limited to dependence on the state variable, which in this instance is assets). Though Laibson (1994) proves a “folk” theorem for this class of games, he does not provide any results for cases in which the scope for self-control is more limited. We see this as a serious gap in the literature. Equilibria with history-dependent strategies provide a way to formally depict the conscious exercise of self-control. Indeed, we believe that the structure of these equilibria correspond reasonably closely to psychologists’ descriptions of actual self-control mechanisms. In particular, Ainslee (1975) suggests that individuals keep themselves in line by construing local deviations as having global significance.¹ For example, an individual might tell herself that, if she is unable to save today, then she is the type of person who will never be able to save. This statement will literally be true if she is playing an equilibrium in which she punishes high consumption by consuming more in the future.

In this paper, we examine a class of intertemporal consumption allocation models in which time inconsistent preferences create self-control problems, and behavior corresponds to an equilibrium of a game played by successive incarnations of the single decision-maker. To determine the scope for deliberate self-control, we study the features of the set of all subgame perfect Nash equilibria (SPNE). For parametric cases, we characterize this set computationally through an algorithm based on the self-generation mapping (Abreu, 1988). Notably, we do not impose any restrictions on strategies whatsoever (for example, we do not necessarily require stationarity, or that the decision-maker to punish deviations by reverting

¹See his discussion on private side bets.

to the Markov-perfect equilibrium). We believe that there are analytical counterparts to the computational results obtained here, and we hope to include general analytical results in subsequent versions of this paper.

Our central computational result is the reasonably robust existence of a threshold asset level, below which no self-discipline is possible (the individual simply consumes assets down to zero), and above which it is possible to sustain positive asset accumulation. The threshold defines a “low asset trap,” in this sense that those with low initial assets can never accumulate much of anything, even though they would be able to accumulate substantial wealth if their initial assets were greater. This result suggests that successful self-discipline requires “priming of the pump”: it takes assets to accumulate more assets. It also suggests that the key difference between otherwise similar savers and non-savers may be initial conditions rather than differences in tastes.

We also investigate the sensitivity of the equilibrium set to changes in various parameters. One particularly interesting experiment involves making it easier to borrow (by relaxing liquidity constraints). It turns out that the relaxation of restrictions on borrowing makes self-control easier to impose, in the sense that fewer individuals are caught in the low asset trap. This result suggests that the development of credit markets during the 1970s and 1980s may not be responsible for the decline in saving, contrary to the assertion of Laibson (1996) (who, as mentioned above, focuses on Markov-perfect equilibria).

The remainder of this paper is organized as follows. Section 2 presents the model. Section 3 describes equilibria. Section 4 summarizes the computational method. Section 5 provides our computational results.

2 The Model

In this section we describe an infinite-horizon, dynamic decision problem by an individual with dynamically inconsistent (quasi-hyperbolic) preferences. An individual is modeled as a composite of independent temporary selves. In each period, a particular “self,” indexed by the time period, makes the consumption and saving decision. The preferences of the t -period self are given by

$$U_t(c_0, c_1, \dots, c_t, \dots) = (1 + \beta)u(c_t) + \sum_{\tau=1}^{\infty} \delta^\tau u(c_{t+\tau}) \quad (1)$$

where β and δ are discount parameters, and $u(c)$ is a continuous, strictly concave function. Additionally we assume that utility is bounded below: $u(0) = B \geq -\infty$.²

During its period of control, self t observes all past consumption levels $(c_0, c_1, \dots, c_{t-1})$ and the current asset level A_t . Self t then receives a constant income y and subsequently chooses a consumption level for period t which satisfies the liquidity constraint:

$$0 \leq c_t \leq A_t + y + b \quad (2)$$

where b is the borrowing limit. Self t may randomize over this choice. Since there is really only one decision-maker in this model, it is natural to assume that any randomization is public (that is, later selves can observe the probability rule governing the selection of consumption, rather than just the particular level of consumption selected). Randomization will be used to convexify the set of possible continuation payoffs. This simplifies the computational algorithm significantly.

After self t makes the consumption decision for period t , assets in period $t+1$ are determined according to the technology

$$A_{t+1} = f(A_t, c_t) = \min\{\bar{A}, \alpha(A_t + y - c_t)\}. \quad (3)$$

Note that we impose an upper bound on assets, \bar{A} . This is for computational convenience.

The initial level of assets A_0 is drawn from the asset set $\mathcal{A} = [-\alpha b, \bar{A}]$. To ensure that consumption is feasible in each period, borrowing capacity is limited by the following constraint:

$$-\alpha b + y + b \geq 0. \quad (4)$$

²If utility does not have a lower bound, then it is in general possible to sustain any feasible consumption path in a subgame perfect equilibrium. See Laibson (1994).

Note that we can rewrite this as:

$$\begin{aligned} b &\leq \frac{y}{\alpha - 1} && \text{for } \alpha > 1, \\ b &\geq \frac{y}{\alpha - 1} && \text{for } \alpha < 1. \end{aligned} \tag{5}$$

In other words, borrowing capacity must be limited to the present discounted value of income.

Henceforth, we will refer to the value of utility discounted at the rate δ (ignoring β) as quasi-utility. An upper bound on quasi-utility, $\mu(A)$, is given by the solution to the following programming problem:

$$\begin{aligned} \mu(A) &= \max \sum_{\tau=0}^{\infty} \delta^{\tau} u(c_{\tau}) && \text{subject to} \\ 0 &\leq c_t \leq A_t + y + b \\ A_{t+1} &= f(A_t, c_t) = \min\{\bar{A}, \alpha(A_t + y - c_t)\} \end{aligned} \tag{6}$$

The lower bound on discounted utility is $\lambda = \frac{B}{1 - \delta}$.

3 Equilibria

Let \mathcal{V} be the set of all correspondences $W : [-\alpha b, \bar{A}] \implies [\lambda, \mu(\bar{A})]$ satisfying, $\forall A \in [-\alpha b, \bar{A}]$, $\max \lambda \leq W(A) \leq \mu(A)$. Let \mathcal{W} be the set of all convex-valued correspondences contained in \mathcal{V} . The *self-generation mapping* Φ defines the set of possible discounted current payoffs consistent with Nash equilibrium strategies with continuation values in W .

Definition 1 $\Phi : \mathcal{W} \rightarrow \mathcal{W}$ is given by $\Phi(W)(A) = co(\Psi(W)(A))$, where $v \in \Psi(W)(A)$ iff $\exists c \in [0, A + b + y]$ and $v' \in W(f(A, c))$ such that

$$\begin{aligned} (i) \quad &v = u(c) + \delta v' \\ (ii) \quad &(1 + \beta) u(c) + \delta v' \\ &\geq \sup_{c' \in [0, A + y + b]} \{(1 + \beta) u(c) + \delta \inf w(f(A, c'))\} \end{aligned}$$

Consider a fixed point of this mapping, $W = \Phi(W)$. Each point in W corresponds to one or more subgame-perfect Nash equilibrium (SPNE).

Definition 2 Let $W_0(A) = [\lambda, \mu(A)]$ and recursively $W_t = \Phi(W_{t-1})$. Define W_∞ as follows:

$$W_\infty = \bigcap_{t=0}^{\infty} W_t$$

One can show that W_∞ is a fixed point of Φ ; it exists and is non-empty valued. Thus one can obtain SPNE by constructing W_∞ . One can also show that W_∞ contains *all* SPNE, so we can obtain a complete characterization of SPNE by constructing it. We state the pertinent results but omit the proofs; they are similar to arguments given in Abreu (1988).

Theorem 1 (i) W_∞ is a non-empty u.h.c. correspondence and (ii) $\Phi(W_\infty) = W_\infty$.

Note that the quasi-payoffs from any SPNE can be written as a sequence of correspondences W^t where $W^t(A_t)$ denotes the set of continuation payoffs following from all histories leading to asset level A_t in period t , and where $\forall t, A, W^t(A) \subseteq \Phi(W^{t+1})(A)$. Thus the following result shows that W_∞ contains all SPNE.

Theorem 2 Consider any sequence W^t s.t. $\forall t, A \in [-\alpha b, \bar{A}]$, $W^t(A) \subseteq \Phi(W^{t+1})(A)$. Then $W^t(A) \subseteq W_\infty(A)$.

4 Computational method

This section details the numerical computation of the equilibrium value set by repeated application of the Φ operator to an initial set $W_o(A) \supseteq W_\infty(A)$, $A \in [-\alpha b, \bar{A}]$. Theorem 2 implies the sequence of sets $\{W_t(A)\}_{t=0}^{\infty}$ where $W_{t+1}(A) = \Phi(W_t)(A)$ converges to $W_\infty(A)$ for each $A \in [-\alpha b, \bar{A}]$. To find the largest bounded fixed point of the set-valued map Φ , we employ a dynamic programming approach. Since $W(A)$ is convex for each A , in approximating the set W , we only need to compute the extreme values of the set (the upper and lower bounds) for each asset level.

The computational algorithm proceeds in three steps, which we describe formally below. First, we discretize the action and utility spaces. Second, given that continuation payoffs are

governed by some correspondence W , we determine the best-deviation payoffs in each state A (assuming the worst feasible punishments in the continuation set). Third, we maximize and minimize discounted utility subject to the incentive compatibility constraint and the constraint on continuation utilities. These are simply two dynamic programming problems; one maximizes, the other minimizes the value function subject to incentive compatibility constraints. The solutions to these optimization problems determine the utility bounds (\bar{w} and \underline{w} for each asset level A) for the next iteration.

For the optimization step, we think of the individual as choosing the continuation level of assets rather than current consumption. From a theoretical perspective, it does not matter whether the individual chooses continuation assets or consumption, since there is no uncertainty in the model. However, from a computational perspective, it is more convenient to model the individual as choosing continuation assets.³

The three steps in our computational algorithm are repeated until convergence is reached. The convergence criterion measures the largest difference in utility bounds for each asset level between successive approximations. We end our iterations when this difference is less than a given ϵ or,

$$\max_{A \in \mathcal{A}} \{ \max \{ |\underline{w}^t(A) - \underline{w}^{t+1}(A)|, |\bar{w}^t(A) - \bar{w}^{t+1}(A)| \} \} < \epsilon.$$

More formally, for a given set of parametric assumptions, our computational algorithm repeatedly applies the following steps until convergence is achieved:

Step 1: Initialize

- 1.1: Define an asset grid, \mathcal{A} , on $[-\alpha b, \bar{A}]$
- 1.2: Determine utility bounds $[\bar{w}(A), \underline{w}(A)]$ for each $A \in \mathcal{A}$.

Step 2: Best Deviations

- 2.1: Let $A(A_j) = \{A_i \in \mathcal{A} \mid c(A_i, A_j) \in [0, A_j + y + b]\}$ where $c(A_i, A_j) = A_j + y - A_i/\alpha$.

³If consumption remains the choice variable, then we would need to discretize the consumption set. Additionally, the technology would have to be modified to ensure that for each current asset level and consumption choice, next period's assets are in the discretized asset set.

2.2: For each $A_i \in A(A_j)$ compute

$$\tilde{P}(A_i, A_j) = (1 + \beta) U(c(A_i, A_j)) + \delta \underline{w}(A_i).$$

2.3: For each $A_j \in \mathcal{A}$ compute $P(A_j) = \max_{A_i \in A(A_j)} \tilde{P}(A_i, A_j)$.

Step 3: Optimization

3.1: Compute

$$\begin{aligned} \bar{\Phi}(\bar{w}, \underline{w})(A_j) &= \max_{A_i \in A(A_j)} \{ U(c(A_i, A_j)) + \delta \bar{w}(A_i) \} \text{ s.t.} \\ &(1 + \beta)U(c(A_i, A_j)) + \delta \bar{w}(A_i) \geq P(A_j) \end{aligned}$$

3.2: Compute

$$\begin{aligned} \underline{\Phi}(\bar{w}, \underline{w})(A_j) &= \min_{A_i \in A(A_j)} \{ U(c(A_i, A_j)) + \\ &\delta * \max\{\underline{w}(A_i), 1/\delta [P(A_j) - (1 + \beta) U(c(A_i, A_j))]\} \end{aligned}$$

Step 4: Updating.

4.1: Set $[\bar{w}(A), \underline{w}(A)] = [\underline{\Phi}(\bar{w}, \underline{w})(A), \bar{\Phi}(\bar{w}, \underline{w})(A)]$. Stop if convergence is reached; else go to Step 2.

Note that, in step 3.1, we always use $\bar{w}(A_i)$ as the continuation utility. This is appropriate since, if any continuation utility satisfies the incentive compatibility condition, then the highest feasible continuation utility also satisfies this condition. In contrast, in step 3.2, we do not necessarily use $\underline{w}(A_i)$ as the continuation utility. This is because the lowest feasible continuation utility does not necessarily satisfy the incentive compatibility condition when this condition is satisfied for a higher continuation utility. The $\max\{\bullet\}$ term in 3.2 assures that the continuation payoff is not so low as to violate incentive compatibility.

5 Computational Results

For the results reported in this section, the utility function is assumed to take the following form:

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

where σ is the relative risk aversion parameter. Our results concern cases where $\sigma < 1$, so $U(0) > -\infty$, as required. Assets take on 2501 values between $[-ab, \bar{A}]$. The upper bound on assets, \bar{A} is set to 20.0 throughout. Other parameter values are listed on the graphs depicting the results of each simulation. Throughout, we will refer to dynamically consistent preferences ($\beta = 0$) as the Ramsey case. As Laibson (1994) has established, the unique subgame perfect equilibrium for the Ramsey case corresponds to the solution of the standard dynamic programming problem.

5.1 Basic results: the low-asset trap

Figures 1 through 6 provide comparisons of the equilibrium sets for the Ramsey case and an otherwise identical case with dynamically inconsistent preferences ($\beta = 0.9$). Figures 1 and 2 depict $\bar{w}(A)$ and $\underline{w}(A)$. Figures 3 and 4 depict the corresponding consumption decisions, while figures 5 and 6 show the corresponding asset (savings) choices. In an abuse of notation, $\bar{c}(A)$ is used to denote the consumption levels on the frontier ($\bar{w}(A)$) and $\underline{c}(A)$ corresponds to the consumption levels for the utility levels $\underline{w}(A)$. Similar notation is used for savings, where $\bar{A}'(A)$ and $\underline{A}'(A)$ denote next period's asset levels on each respective frontier as a function of current assets. Note that consumption and savings corresponding to interior of the W set have not been plotted.

An inspection of figures 1 and 2 reveals the following. Not surprisingly, continuation payoffs smoothly increase with initial assets for the Ramsey case. Continuation payoffs are also monotonically increasing but below Ramsey payoffs for the dynamically inconsistent case. Moreover, these payoffs change discontinuously once assets pass through some threshold. Below this threshold, $\bar{w}(A)$ and $\underline{w}(A)$ coincide, but they diverge for higher values of A .

Figures 3 and 4 depict the associated consumption functions. Note that these functions are again smooth in the Ramsey case. In contrast, with dynamically inconsistent preferences, the consumption function associated with $\bar{w}(A)$ has a downward discontinuity at the asset threshold mentioned above. Below this threshold, the consumption functions for $\bar{w}(A)$ and $\underline{w}(A)$ coincide. Since this game has a Markov-perfect equilibrium, this implies that the Markov-perfect equilibrium path is the only subgame perfect equilibrium path if




Figure 1: $W: \beta = 0.0$

$\cdots \underline{w}(A), -\bar{w}(A)$




Figure 2: $W: \beta = 0.9$

$\cdots \underline{w}(A), -\bar{w}(A)$

initial assets are below the threshold. When initial assets are above the threshold, the decision-maker is able to impose self-discipline and achieve a strictly lower level of consumption. Returning to figure 2, note that there is no downward discontinuity in $\underline{w}(A)$ at the critical asset threshold where $\bar{w}(A)$ takes an upward jump. Based on this observation, we conjecture (but have not proven) that the worst equilibrium punishment is always the Markov-perfect equilibrium. If this is correct, then it helps us to resolve equilibrium selection. In principle, the first decision-maker could prefer the consumption path associated with $\underline{w}(A)$ to the consumption path associated with $\bar{w}(A)$, even though quasi-utility is higher in the latter case, in which case the individual would never wish to start imposing self-discipline. However, if the punishment path is the Markov-perfect equilibrium path, then the best-deviation-plus-punishment path will also be the Markov-perfect equilibrium path. The incentive compatibility condition would then tell us that the first decision maker must prefer the consumption path associated with $\bar{w}(A)$ to the consumption path associated with $\underline{w}(A)$. Figures 5 and 6 depict the associated asset (savings) functions. These functions are again smooth in the Ramsey case. Naturally, for the dynamically inconsistent case, $\bar{A}'(A)$ jumps up discontinuously at the critical asset threshold mentioned above. Note

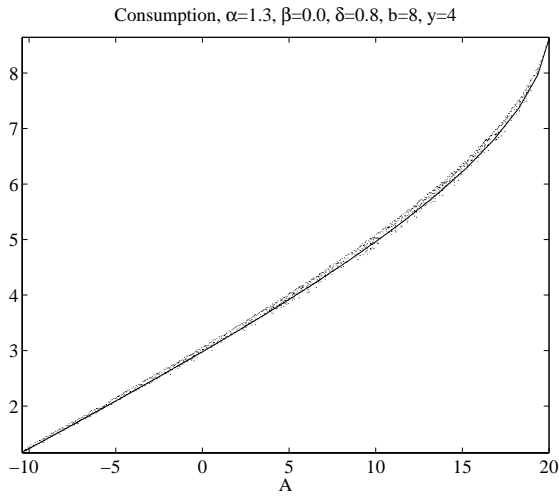


Figure 3: Consumption: $\beta = 0.0$
 $\cdots \underline{c}(A), - \bar{c}(A)$

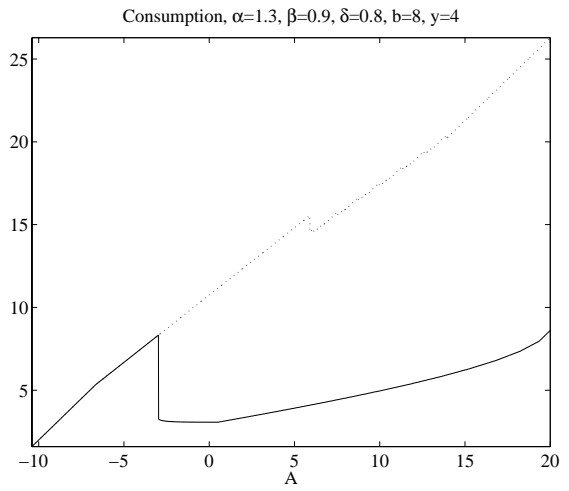


Figure 4: Consumption: $\beta = 0.9$
 $\cdots \underline{c}(A), - \bar{c}(A)$

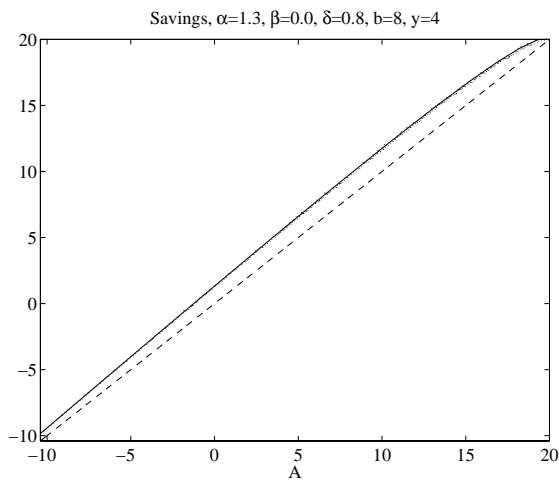


Figure 5: Savings: $\beta = 0.0$
 $\cdots \underline{A}'(A), - \bar{A}'(A)$




Figure 6: Savings: $\beta = 0.9$
 $\cdots \underline{A}'(A), - \bar{A}'(A)$

carefully that $\bar{A}'(A)$ jumps exactly to the 45 degree line. This is a robust feature of our simulations, and it appears to reflect a general principle. No degree of self-control appears to be possible unless the individual can maintain a non-decreasing asset profile. The threshold level of assets is determined as the lowest level of assets that it is just possible for the individual to sustain.

Thus, when initial assets are less than the critical threshold level, the individual necessarily consumes them down to zero in finite time. On the other hand, when initial assets are greater than the critical threshold level, the individual can impose self-discipline and accumulate or maintain additional assets indefinitely. We referred to this in the introduction as the low-asset trap.

5.2 Comparative statics

In this section, we explore the effects of varying some of the model's parameters on the set of equilibria. We begin with the parameter b . An increase in b reflects greater liquidity. Intuitively, one might expect individuals to have greater difficulty exercising self-control when borrowing is easier. Indeed, for Markov-perfect equilibria, Laibson [19**] concludes that this expectation is correct, and he suggests that the decline in saving during the 1980s may be attributable in part to credit market innovations. The issue is, however, more complicated for history-dependent subgame perfect equilibria. A greater ability to borrow implies greater temptation to defect from a desired plan, but it also implies that the consequences of defection may be more severe.

Figures 7 and 8 depict the effect of an increase in b on, respectively, continuation payoffs and asset (savings) functions. Note that an increase in liquidity (b) has two distinct effects on the savings function. First, below the critical level of initial assets that defines the low-asset trap, the level of continuation utility and the level of savings both fall with b . This is consistent with Laibson's result for Markov perfect equilibrium. However, there is also a second effect: the critical asset level that defines the low-asset trap falls with b . While those in the low-asset trap save less when credit is more readily available, fewer households get caught in this trap. In addition, asset accumulation also rises slightly with liquidity above the low-asset trap.




Figure 7: W : Low and High b
 $b = 8$, $\cdots \underline{w}(A)$, $— \bar{w}(A)$
 $b = 9$, $\bullet \underline{w}(A)$, $- - \bar{w}(A)$




Figure 8: Savings: Low and High b
 $b = 8$, $— \bar{A}'(A)$
 $b = 9$, $\cdots \bar{A}'(A)$

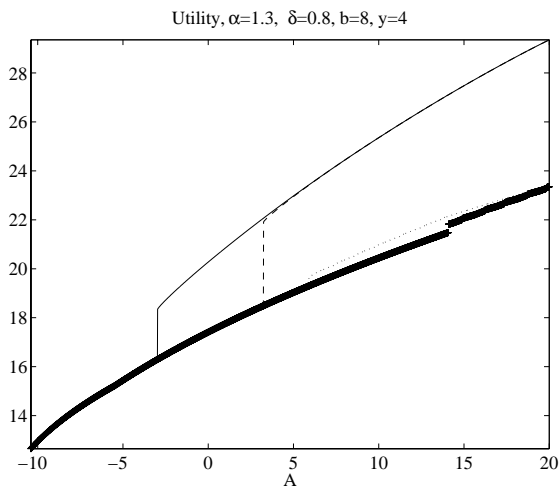


Figure 9: W : Low and High β
 $\beta = 0.9$, $\cdots \underline{w}(A)$, $— \bar{w}(A)$
 $\beta = 1.1$, $\bullet \underline{w}(A)$, $- - \bar{w}(A)$




Figure 10: Savings: Low and High β
 $\beta = 0.9$, $— \bar{A}'(A)$
 $\beta = 1.1$, $\cdots \bar{A}'(A)$

Figures 9 and 10 depict the effect of an increase in β , which corresponds to a worsening of the dynamic consistency problem. Note that this reduces the highest sustainable continuation utilities. Savings decline both below and above the low-asset trap. In addition, the critical level of assets that defines the low-asset trap rises, which implies that more individuals are caught in the trap.

Figures 11 through 14 depict the effect of an increase in the gross return on capital, α . One can think of this as corresponding to an increase in the interest rate or a reduction in the rate of capital income taxation. In effect, our simulations shed some light on the interest elasticity of saving. An especially interesting question is whether dynamic inconsistency increases or decreases this key elasticity. Consequently, we present two sets of simulation results, one for $\beta = 0$ (figures 11 and 12), and another for $\beta = 0.9$ (figures 13 and 14).




Figure 11: Savings: $\bar{A}'(A)$, $\beta = 0.0$

Figure 12: Savings: $\bar{A}'(A)$, $\beta = 0.0$

Note that a change in α has a larger effect on savings in the dynamically inconsistent case. When α is 1.1, savings decline and all individuals are trapped at the lowest level of assets.⁴ When α is equal to 1.3, those with high enough asset levels are able escape the low asset trap. For those households with debt (negative asset holdings), a higher α further increases

⁴Note that the lower bound on assets depend on α and the borrowing limit. If α is higher, than households are caught in a low-asset trap with larger debt.




Figure 13: Savings: $\bar{A}'(A)$, $\beta = 0.9$




Figure 14: Savings: $\bar{A}'(A)$, $\beta = 0.9$

their debt. Consequently, at low levels of assets, the savings line for $\alpha = 1.1$ lies above the savings line for $\alpha = 1.3$.

When α increases from 1.3 to 1.45, as depicted in figures 12 and 14, increase in savings for households caught in the low-asset trap (and near the low-asset trap) is greater than for those who are able to save at the Ramsey level. When $\alpha = 1.45$, all households are able to sustain Ramsey level of savings.

6 Conclusions

We have examined a class of intertemporal consumption allocation models in which time inconsistent preferences create self-control problems. Behavior corresponds to an equilibrium of a game played by successive incarnations of the single decision-maker, and the scope for self-control is defined by the set of subgame perfect equilibria. A reasonably robust feature of the equilibrium set is the existence of a threshold asset level, below which no self-discipline is possible (the individual simply consumes assets down to zero), and above which it is possible to sustain positive asset accumulation. We have referred to this as the low-asset trap. We have investigated the sensitivity of the equilibrium set to various param-

eters. For example, we show that when credit becomes more readily available, self-control becomes easier to impose, and fewer individuals are caught in the low asset trap.

References

- [1] Abreu, D. (1988) “Extremal Equilibria of Oligopolistic Supergames.” *Journal of Economic Theory* 39: 191-225.
- [2] Abreu, D. (1988) “On the Theory of Infinitely Repeated Games with Discounting.” *Econometrica* 56: 383-396.
- [3] Ainslee, G. (1975) “Specious Reward: A Behavioral Theory of Impulsiveness and Impulse Control” *Psychological Bulletin* 82(4): July 463-96.
- [4] Laibson, David I. (1994) “Self-control and Saving.” Massachusetts Institute of Technology mimeo.
- [5] Laibson, David I. (1996) “Hyperbolic Discount Functions, Undersaving and Savings Policy.” National Bureau of Economic Research Working Paper: 5635.
- [6] Laibson, David I. (1997) “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics* 112, n2: 0033-5533.
- [7] Phelps, E. S. and R. A. Pollak. (1968) “On Second-best National Saving and Game-Equilibrium Growth.” *Review of Economics Studies* 35: 185-199.