Hyperbolic Discounting and Consumption

Christopher Harris   David Laibson*

February 18, 2001

Abstract. Laboratory and field studies of time preference find that discount rates are much greater in the short-run than in the long-run. Hyperbolic discount functions capture this property. This paper characterizes the behavior of hyperbolic consumers who have stochastic income and cannot trade this risk (i.e. the standard buffer-stock assumptions). We prove an equilibrium uniqueness theorem, characterize properties of the consumption function, and illustrate the consumption function with numerical simulations. We show that hyperbolic consumption functions may exhibit pathologies like discontinuities, non-monotonicities, and concavity violations. The pathologies are exacerbated as hyperbolicity increases, risk aversion falls, and income uncertainty falls. We also show that these pathologies do not arise when the model parameters are calibrated at empirically sensible benchmark values. We also discuss empirical implications of the hyperbolic model. Relative to the exponential model, the hyperbolic model better matches available evidence on credit card borrowing, consumption-income comovement, and liquid/illiquid asset accumulation.

JEL classification: C6, C73, D91, E21.

*We thank Glenn Ellison for numerous helpful suggestions. We also thank George-Marios Angeletos, Andrea Repetto, Jeremy Tobacman and Stephen Weinberg, whose ideas and work are reflected in this essay. Laura Serban provided outstanding research assistance. Laibson acknowledges financial support from the National Science Foundation (SBR-9510985), the Olin Foundation, the National Institute on Aging (R01-AG-1665) and the MacArthur Foundation.

Keywords: hyperbolic discounting, time preference, dynamic inconsistency, consumption, savings, buffer stock, Euler Relation, dynamic games, bounded variation calculus, altruistic growth.

1. Introduction

Robert Strotz (1956) first suggested that people are more impatient when they make short-run tradeoffs than when they make long-run tradeoffs.[1] Virtually every experimental study on time preference has supported Strotz's conjecture.[2] When two rewards are both far away in time, decision-makers act relatively patiently (e.g., I prefer two apples in 101 days, rather than one apple in 100 days). But when both rewards are brought forward in time, preferences exhibit a reversal, reflecting more impatience (I prefer one apple right now, rather than two apples tomorrow).[3]

Such reversals should be well understood by everyone who makes far sighted New Year's resolutions and later backtracks. We promise ourselves to exercise, diet, and quit smoking, but often postpone those virtuous behaviors when the moment arrives to make the required sacrifices. Looking to the long run we wish to act patiently, but the desire for instant gratification frequently overwhelms our good intentions.

The contrast between long-run patience and short-run impatience has been modelled with discount functions that take an approximately hyperbolic form (Ainslie 1992, Loewenstein and Prelec 1992, Laibson 1997a). Such preferences imply that the instantaneous discount rate declines as the horizon increases. This pattern of discounting sets up a conflict between today's preferences and the preferences which will be held in the future. From the perspective of period 0, the discount rate between two distant periods, t and $t+1$, is a long-term low discount rate. However, from the perspective of period t, the discount rate between t and $t+1$ is a short-term high discount rate.

Hyperbolic consumers will report a gap between what they feel they should save and what they actually save. Prescriptive saving rates will lie above actual sav-

[1]Some of Strotz's insights are anticipated by Ramsey (1928).

[2]See Ainslie 1992 and Frederick et al 2001, for reviews of the evidence for and against hyperbolic discounting.

[3]This example is from Thaler (1981).

ings rates, since short-run preferences for instantaneous gratification will undermine the consumer's desire to implement long-run patient plans. However, the hyperbolic consumer is not doomed to retire in poverty. Illiquid assets can help the hyperbolic consumer lock in the patient, welfare-enhancing course of action. Hence, the availability of illiquid assets becomes a critical determinant of household savings and welfare. However, too much illiquidity can be problematic. Consumers face substantial uninsurable labor-income risk, and need to use liquid assets to smooth their consumption. Hyperbolic agents seek an investment portfolio that strikes the right balance between commitment and flexibility.

In this paper we review and extend the literature on hyperbolic discounting and consumption. We begin our analysis of hyperbolic consumers by describing an infinite-horizon consumption problem with a single liquid asset. Using this tractable problem, we characterize equilibrium behavior. We prove a new equilibrium uniqueness theorem, characterize some properties of the consumption function, and illustrate additional properties of the consumption function with numerical simulations. We show that hyperbolic consumption functions may exhibit pathologies like discontinuities, non-monotonicities, and concavity violations. We analyze the comparative statics of these pathologies. The pathologies are exacerbated as hyperbolicity increases, risk aversion falls, and income uncertainty falls. We also show that these pathologies do not arise when the model parameters are calibrated at empirically sensible benchmark values. Finally, we review our earlier results on the Euler relation characterizing the equilibrium path (Harris and Laibson 2001a).

We then discuss simulations of savings and asset allocation choices of households who face a *lifecycle* problem with liquid assets, liquid liabilities, and illiquid assets (Angeletos, Laibson, Repetto, Tobacman, and Weinberg 2001a). These lifecycle simulations are used to compare the behavior of hyperbolic households and exponential households. Both the exponential and hyperbolic households are calibrated to hold levels of pre-retirement wealth that match observed levels of wealth reported in the

Survey of Consumer Finances. Despite the fact that this calibration imposes identical levels of total wealth for hyperbolics and exponentials, numerous differences arise.

First, the hyperbolic households invest comparatively little of their wealth in liquid assets. They hold relatively low levels of liquid wealth measured either as a fraction of labor income or as a share of total wealth. Analogously, hyperbolic households also borrow more aggressively in the revolving credit market (i.e., on credit cards). The low levels of liquid wealth and high rates of credit-card borrowing generated by hyperbolic simulations match empirical measures from the Survey of Consumer Finances much better than the results of exponential simulations.

Because the hyperbolic households have low levels of liquid assets and high levels of debt, they are unable to smooth their consumption paths in the presence of predictable changes in income. Calibrated hyperbolic simulations display substantial comovement between consumption and predictable income growth, matching empirical measures of comovement from the PSID. By contrast, calibrated exponential simulations generate too little consumption-income comovement. Similarly, hyperbolic simulations generate substantial drops in consumption around retirement, matching empirical estimates. The exponential simulations fail to replicate this pattern. All in all, the hyperbolic model matches observed consumption data better than the exponential model.

Our paper is organized in 12 sections, and readers are encouraged to pick and choose among them. Section 4 contains the most technical parts of the paper, and can be skipped by readers primarily interested in applications.

In Section 2, we discuss the hyperbolic discount function. In Section 3 we present a one-asset, infinite horizon buffer-stock consumption model, which can accommodate either exponential or hyperbolic preferences. In Section 4 we discuss existence and uniqueness of an equilibrium. In Section 5 we describe the Euler relation that characterizes the equilibrium path. In Section 6 we describe our numerical simulations of the one-asset consumption problem. In Section 7 we describe the properties of the hyperbolic consumption function, and illustrate these properties with simulations. In Section 8 we review empirical applications of the hyperbolic model. In Section 9 we discuss the level of consumer sophistication assumed in hyperbolic models. In Section 10 we describe the policy implications of the hyperbolic model. In Section 11 we discuss some important extensions of the hyperbolic model, including applications in continuous time. In Section 12 we conclude.

2. HYPERBOLIC DISCOUNTING

When researchers elicit time preferences, they ask subjects to choose among a set of delayed rewards. The largest rewards are accompanied by the greatest delays.[4] Researchers use subject choices to estimate the shape of the discount function. These estimated discount functions almost always approximate generalized hyperbolas: events τ periods away are discounted with weight $(1+\alpha\tau)^{-\gamma/\alpha}$, with $\alpha, \gamma > 0$ (Loewenstein and Prelec 1992).[5]

Figure 1 graphs the generalized hyperbolic discount function with parameters $\alpha = 4$ and $\gamma = 1$. Figure 1 also plots the standard exponential discount function, δ^τ, assuming $\delta = 0.944$ (the annual discount factor used in our simulations).

Since the discount rate represents the rate of decline of the discount function, the exponential discount function implies a constant discount rate:

$$-\frac{\frac{\partial}{\partial\tau}\left(\delta^\tau\right)}{\delta^\tau} = -\ln\delta.$$

[4] Such experiments have used a wide range of real rewards, including money, durable goods, fruit juice, sweets, video rentals, relief from noxious noise, and access to video games. For example, see Thaler 1981, Navarick 1982, Millar and Navarick 1984, King and Logue 1987, Kirby and Herrnstein 1995, Kirby and Marakovic 1995, 1996, Kirby 1997, Read et al 1996. See Ainslie 1992, Frederick et al 2001, and Angeletos et al 2001 for partial reviews of this literature. See Mulligan 1997 for a critique.

[5] Loewenstein and Prelec (1992) provide an axiomatic derivation of the generalized hyperbolic discount function. See Chung and Herrnstein (1961) for the first use of the hyperbolic discount function. The original psychology literature worked with the special cases $\frac{1}{\tau}$ and $\frac{1}{1+\alpha\tau}$. Ainslie (1992) reviews this literature.

By contrast, the hyperbolic discount function implies a discount rate that falls with the horizon, τ:

$$-\frac{\frac{\partial}{\partial\tau}\left((1+\alpha\tau)^{-\gamma/\alpha}\right)}{(1+\alpha\tau)^{-\gamma/\alpha}}=\frac{\gamma}{(1+\alpha\tau)}.$$

In the short run, the hyperbolic discount rate is γ and in the long run the discount rate converges to zero. This reflects the robust experimental finding that people are very impatient in the short run (e.g., when postponing a reward from today to tomorrow) and very patient when thinking about long-run tradeoffs (postponing a reward from 100 days to 101 days).

In order to reflect the empirical pattern of discount rates that fall with the horizon, Laibson (1997a) adopted a discrete-time discount function, $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots .\}$, which Phelps and Pollak (1968) had previously used to model intergenerational time preferences.[6] This "quasi-hyperbolic function" reflects the sharp short-run drop in valuation measured in the experimental time-preference data and has been adopted as a research tool because of its analytical tractability.[7] Figure 1 plots the particular parameterization of the quasi-hyperbolic discount function used in our simulations: $\beta = 0.7$ and $\delta = 0.957$. Using annual periods, these parameter values roughly match experimentally measured discounting patterns. Delaying an immediate reward by a year reduces the value of that reward by approximately $40\% \approx 1 - \beta\delta$. By contrast, delaying a distant reward by an additional year reduces the value of that reward by a relatively small percentage: $1 - \delta$.[8]

All forms of hyperbolic preferences induce dynamic inconsistency. Consider the discrete-time quasi-hyperbolic function. The discount factor between adjacent periods

[6] Akerlof (1991) used a similar function: $\{1, \beta, \beta, \beta, \ldots\}$.

[7] The quasi-hyperbolic discount function is only "hyperbolic" in the sense that it captures the key qualitative property of the hyperbolic functions: a faster rate of decline in the short-run than in the long-run. Laibson (1997a) adopted the phrase "quasi-hyperbolic" to emphasize the connection to the hyperbolic-discounting literature in psychology (Ainslie, 1992). O'Donoghue and Rabin (1999a) call these preferences "present biased." Krusell and Smith (2000) call these preferences "quasi-geometric."

[8] See Ainslie 1992 and Frederick et al (2000).

t and $t+1$ represents the weight placed on utils at time $t+1$ relative to the weight placed on utils at time t. From the perspective of self t, the discount factor between periods t and $t+1$ is $\beta\delta$, but the discount factor that applies between any two later periods is δ. Since we take β to be less than one, this implies a short-term discount factor that is less than the long-term discount factor.[9] From the perspective of self $t+1$, $\beta\delta$ is the relevant discount factor between periods $t+1$ and $t+2$. Hence, self t and self $t+1$ disagree about the desired level of patience which should be used to trade-off rewards in periods $t+1$ and $t+2$.

Because of this dynamic inconsistency, the hyperbolic consumer is involved in a decision which has intra-personal strategic dimensions. Early selves would like to commit later selves to honor the preferences of those early selves. Later selves do their best to maximize their own interests. Economists have modelled this situation as an intra-personal game played among the consumer's temporally situated selves (Strotz 1956). Recently, hyperbolic discount functions have been used to explain a wide range of anomalous economic choices, including procrastination, contract design, drug addiction, self deception, retirement timing, and undersaving.[10] We focus here on the implications for lifecycle savings decisions.

In the sections that follow, we analyze the 'sophisticated' version of the hyperbolic model. Sophisticated hyperbolic consumers correctly predict that later selves will not honor the preferences of early selves. By contrast, 'naive' consumers make current choices under the false belief that later selves will act in the interests of the current self. The assumption of naivety was first proposed by Strotz (1956), and has since been carefully studied by Akerlof (1991) and O'Donoghue and Rabin (1999a, 1999b,

[9]Note that a discount factor, say θ, is inversely related to the discount rate, $-\ln\theta$.

[10]For example, see Akerlof (1991), Barro (1997), Benabou and Tirole (2000), Brocas and Carrillo (2000, 2001), Carrillo and Dewatripont (2000), Carrillo and Marriotti (2000), Della Vigna and Paserman (2000), Della Vigna and Malmendier (2001) Diamond and Koszegi (1998), Gruber and Koszegi (2000), Laibson (1994,1996,1997a), Krusell et al (2000a, 2000b), O'Donoghue and Rabin (1999a, 1999b, 2000).

2000). We return to a discussion of naifs in section 9.

3. The Consumption Problem

Our benchmark model adopts the technological assumptions of standard "buffer-stock" consumption models like those originally developed by Carroll (1992, 1997) and Deaton (1991). These authors assume stochastic labor income and incomplete markets — consumers cannot borrow against uncertain future labor income. In this section we consider a stripped-down stationary version of the standard buffer-stock model. In Section 8 we discuss a more complex lifecycle model, with a richer set of institutional assumptions.

Our modeling assumptions for the stripped-down model divide naturally into four parts: the standard assumptions from the buffer-stock literature; the assumptions that make our model qualitatively hyperbolic; our equilibrium concept; and the technical assumptions that allow us to derive the Hyperbolic Euler Relation. We discuss the first three sets of assumptions below. The fourth set of assumptions is presented in Section 4.1 below.

3.1. Buffer-Stock Assumptions. During period t, the consumer has cash-on-hand $x_t \geq 0$. She chooses a consumption level $c_t \in [0, x_t]$, which rules out borrowing. Whatever the consumer does not spend is saved, $s_t = x_t - c_t \in [0, x_t]$. The gross return on her savings is fixed, $R \geq 0$, and next period she receives labor income $y_{t+1} \geq 0$. Cash-on-hand during period $t+1$ is therefore $x_{t+1} = R(x_t - c_t) + y_{t+1}$. Labor income is independently and identically distributed over time with density f.

The consumer cannot sell her uncertain stream of future labor-income payments, because of moral hazard and adverse selection, or because of prohibitions against indenturing. In other words, there is no asset market for labor.

3.2. Hyperbolic Preferences. We model an individual as a sequence of autonomous temporal selves. These selves are indexed by the respective periods, $t =$

$0, 1, 2, \ldots$, in which they control the consumption choice. Self t receives payoff

$$\mathrm{E}_t \left[U(c_t) + \beta \sum_{i=1}^{\infty} \delta^i U(c_{t+i}) \right], \tag{1}$$

where $\beta \in [0,1]$, $\delta \in [0,1)$, and $U : [0,+\infty) \to [-\infty,+\infty)$. Our model nests the standard case of exponential discounting: $\beta = 1$, $0 \le \delta < 1$. Our model also nests the quasi-hyperbolic case: $\beta < 1$, $0 \le \delta < 1$.

3.3. Equilibrium. We analyze the set of perfect equilibria in stationary Markov strategies of the intra-personal game with players (or selves) indexed by the non-negative integers. Because income is i.i.d., the only state variable is cash-on-hand x_t. We therefore restrict attention to consumption strategies C that depend only on x_t.

4. Existence and Uniqueness

This technical discussion can be skipped by readers interested primarily in applications. Such readers may wish to move immediately to Section 5.

4.1. Technical Assumptions. We make the following technical assumptions:

U1 U has domain $[0,+\infty)$ and range $[-\infty,+\infty)$;

U2 U is twice continuously differentiable on $(0,+\infty)$;

U3 $U' > 0$ on $(0,+\infty)$;

U4 there exist $0 < \underline{\rho} \le \overline{\rho} < +\infty$ such that $\underline{\rho} \le \frac{-cU''(c)}{U'(c)} \le \overline{\rho}$ for all $c \in (0,+\infty)$;

F1 f has domain $(0,+\infty)$ and range $[0,+\infty)$;

F2 f is twice continuously differentiable;

F3 there exist $0 < \underline{y} < \overline{y} < +\infty$ such that $f(y) = 0$ for all $y \notin [\underline{y}, \overline{y}]$;

D $\max\left\{\delta, \delta R^{1-\underline{\rho}}\right\} < 1.$

Assumptions (U1-U4) could be summarized by saying that U has bounded relative risk aversion (or BRRA for short). They are automatically satisfied if U has constant relative risk aversion (or CRRA for short). Assumptions (F1-F3) could be summarized by saying that f is smooth, that the support of f is compact and that 0 does not lie in the support of f. Assumption (D) ensures that the expected present discounted value of the consumer's utility stream is always well defined. Further discussion of these assumptions can be found in Harris and Laibson (2001a).

4.2. The Bellman Equation of the Hyperbolic Consumer. The intra-personal game of the hyperbolic consumer can be approached recursively as follows. Suppose that self t has current-value function W_t and continuation value function V_t, and suppose that self $t+1$ has consumption function C_{t+1} and current-value function W_{t+1}. Then it follows from the Envelope Theorem that

$$U'\left(C_{t+1}\left(x_{t+1}\right)\right) = W'_{t+1}\left(x_{t+1}\right). \tag{2}$$

Next, it follows from the definition of W_{t+1} and V_t that

$$\beta V_t\left(x_{t+1}\right) = W_{t+1}\left(x_{t+1}\right) - (1-\beta) U\left(C_{t+1}\left(x_{t+1}\right)\right). \tag{3}$$

Finally, it follows from the definition of W_t that

$$W_t\left(x_t\right) = \max_{c\in[0,x_t]} U(c) + \beta\delta \int V_t\left(R\left(x_t - c\right) + y\right) f(y)\, dy. \tag{4}$$

Hence

$$W_t\left(x_t\right) = \max_{c\in[0,x_t]} U(c) + \delta \int \left(W_{t+1} - (1-\beta) U \circ C_{t+1}\right)\left(R\left(x_t - c\right) + y\right) f(y)\, dy$$

(substituting for V_t in equation (4) using equation (3))

$$= \max_{c\in[0,x_t]} U(c) + \delta \int \left(W_{t+1} - \varepsilon U \circ g \circ W'_{t+1}\right) \left(R\left(x_t - c\right) + y\right) f(y)\, dy$$

(where $\varepsilon = 1 - \beta$ and $g = (U')^{-1}$)

$$= \left(\mathfrak{B} W_{t+1}\right)(x_t),$$

say. This is the Bellman equation of the hyperbolic consumer.

4.3. The Finite-Horizon Case: Current-Value Functions. Suppose that the intra-personal game of the hyperbolic consumer has a finite horizon $T < +\infty$. Then, in principle, the current-value functions can be shown to be unique by backwards induction. Indeed, suppose for simplicity that the consumer has no bequest motive. Then we expect $W_T = U$, $W_{T-1} = \mathfrak{B}W_T$, ..., $W_1 = \mathfrak{B}W_2$. In practice, we need to find a space of functions $\mathcal{W}$ such that, if $W_{t+1} \in \mathcal{W}$, then $W_t = \mathfrak{B}W_{t+1}$ is well defined and lies in $\mathcal{W}$. To this end, we make the following definition.

Definition 1. *The function $g : (0, +\infty) \to \mathbb{R}$ is of* **locally bounded variation** *iff there exist increasing functions $g_+ : (0, +\infty) \to \mathbb{R}$ and $g_- : (0, +\infty) \to \mathbb{R}$ such that $g = g_+ - g_-$.*

Now, let us say that two functions of locally bounded variation are equivalent iff they are equal at all points of continuity. Let $\mathcal{BV}^0_{\text{loc}}((0, +\infty))$ is the space of equivalence classes of functions of locally bounded variation, and let $\mathcal{BV}^1_{\text{loc}}((0, +\infty))$ denote the space of equivalence classes of functions W such that both W and W' are of locally bounded variation. Then the correct choice of space for our current-value function is $\mathcal{W} = \mathcal{BV}^1_{\text{loc}}((0, +\infty))$.

To see this, note first that, if $W_{t+1} \in \mathcal{BV}^1_{\text{loc}}((0, +\infty))$, then W'_{t+1} is a function of locally bounded variation. Hence W'_{t+1} is uniquely defined except at a countable set of points, and $\mathfrak{B}W_{t+1}$ is uniquely defined at all points. Second, consider the operator $\mathfrak{b}_\gamma$ given by the formula

$$(\mathfrak{b}_\gamma W_{t+1})(x_t) = U(\gamma x_t) + \delta \int \left(W_{t+1} - \varepsilon U \circ g \circ W'_{t+1}\right)(R(1-\gamma)x_t + y)\, f(y)\, dy.$$

Then

$$\mathfrak{B}W_{t+1} = \sup_{\gamma \in [0,1]} \{\mathfrak{b}_\gamma W_{t+1}\}.$$

In other words, $\mathfrak{B}W_{t+1}$ is the upper envelope of the functions $\mathfrak{b}_\gamma W_{t+1}$. Third, note that $\mathfrak{b}_\gamma W_{t+1}$ is twice continuously differentiable. Moreover, there exists a continuous function $a : (0, +\infty) \to [0, +\infty)$ such that, for all $\gamma \in [0,1]$,

$$|\mathfrak{b}_\gamma W_{t+1}|, \left|(\mathfrak{b}_\gamma W_{t+1})'\right|, \left|(\mathfrak{b}_\gamma W_{t+1})''\right| \le a$$

on $(0, +\infty)$. In particular, there exists a twice continuously differentiable convex function $\kappa : (0, +\infty) \to \mathbb{R}$ such that, for all $\gamma \in [0,1]$, $\mathfrak{b}_\gamma W_{t+1} + \kappa$ is convex. Hence

$$\mathfrak{B}W_{t+1} = \sup_{\gamma \in [0,1]} \{\mathfrak{b}_\gamma W_{t+1}\} = \sup_{\gamma \in [0,1]} \{\mathfrak{b}_\gamma W_{t+1} + \kappa\} - \kappa.$$

In other words, $\mathfrak{B}W_{t+1}$ is the difference of two convex functions. In the light of the following result, this is exactly what we need.

Proposition 2. *Suppose that $W : (0, +\infty) \to \mathbb{R}$. Then $W \in \mathcal{BV}^1_{loc}((0, +\infty))$ iff W is the difference of two convex functions.* ■

4.4. The Finite-Horizon Case: Consumption Functions. Suppose again that the intra-personal game of the hyperbolic consumer has a finite horizon $T < +\infty$, and that the consumer has no bequest motive. Then the consumption function of self T is unique, and is given by the formula

$$C_T(x_T) = x_T;$$

and, for all $1 \le t \le T-1$, the consumption function of self t is any function such that

$$C_t(x_t) \in \underset{c \in [0, x_t]}{\operatorname{argmax}} U(c) + \delta \int \left(W_{t+1} - \varepsilon U \circ g \circ W'_{t+1} \right) (R(x_t - c) + y) f(y) \, dy$$

for all $x_t \in [0, +\infty)$.

Now, $C_t = g \circ W'_t$ is uniquely defined and continuous except on a countable set of points. Since this set of points has measure zero, it is encountered with probability zero. It follows that any two consumption functions of self t are observationally equivalent. By the same token, any two equilibria are observationally equivalent. This uniqueness claim can be made precise by viewing consumption functions as elements of the space $\mathcal{BV}^0_{\mathrm{loc}}((0, +\infty))$.

4.5. The Infinite-Horizon Case: Existence. In order to establish existence in the finite-horizon case, we showed that the Bellman operator $\mathfrak{B}$ was a self-map of the space $\mathcal{BV}^1_{\mathrm{loc}}((0, +\infty))$. In order to establish existence in the infinite-horizon case, we need to strengthen this result by showing that there is a non-empty compact convex subset $\mathcal{K}$ of $\mathcal{BV}^1_{\mathrm{loc}}((0, +\infty))$ such that $\mathfrak{B}$ is a self-map of $\mathcal{K}$.

Define $\underline{V} : [0, +\infty) \to [-\infty, +\infty)$ by the formula

$$\underline{V}(x) = U(x) + \frac{\delta}{1-\delta} \int U(y) f(y) \, dy,$$

define $\overline{V} : [0, +\infty) \to [-\infty, +\infty)$ by the formula

$$\overline{V}(x) = U(x) + \sum_{t=1}^{\infty} \delta^t U \left(\sum_{s=0}^{t-1} R^s y + R^t x \right),$$

and, for all Borel measurable $V \in \left[\underline{V}, \overline{V} \right]$, define $\mathfrak{W}V : [0, +\infty) \to [-\infty, +\infty)$ by the formula

$$(\mathfrak{W}V)(x) = \max_{\gamma\in[0,1]} \left\{ U(\gamma x) + \beta\delta \int V(R(1-\gamma)x + y) f(y)\, dy \right\}.$$

Finally, put $\underline{V}^- = -(\underline{V} \wedge 0)$ and $\overline{V}^+ = +\left(\overline{V} \vee 0\right)$, define $N_1 : [0,+\infty) \to [0,+\infty)$ by the formula $N_1(x) = \underline{V}^-\left(\underline{y}\right) \vee \overline{V}^+(Rx + \overline{y})$, and define $N_2 : [0,+\infty) \to [0,+\infty)$ by the formula $N_2(x) = \frac{U'(x)}{x} \vee N_1(x)$. Then:

Theorem 3 [Global Regularity]**.** *There exist $K > 0$ such that, for all $V \in \left[\underline{V}, \overline{V}\right]$,*

1. $(1-\beta)U + \beta\underline{V} \le \mathfrak{W}\, V \le (1-\beta)U + \beta\overline{V}$,
2. $U' \le (\mathfrak{W}V)' \le U' \vee (KN_1)$ *and*
3. $(\mathfrak{W}V)'' \ge -KN_2$

on $(0,+\infty)$. ■

The required set $\mathcal{K}$ is then simply the set of $W \in \mathcal{BV}^1_{\mathrm{loc}}((0,+\infty))$ that satisfy the three estimates in this theorem.

4.6. The Infinite-Horizon Case: Uniqueness. In order to establish uniqueness in the infinite-horizon case, we begin by showing that, no matter what the initial cash-on-hand of the consumer, there exists a finite interval from which the dynamics of wealth never exit.

Theorem 4 [Absorbing Interval]**.** *Suppose that $\delta R < 1$. Then, for all $x_0 \in [0,+\infty)$, there exists $\overline{\beta}_1 \in [0,1)$ and $\overline{X} \in [x_0,+\infty)$ such that, for all $\beta \in \left[\overline{\beta}_1, 1\right]$ and all equilibria C of the infinite-horizon model,*

$$R(x - C(x)) + y \in \left[\underline{y}, \overline{X}\right]$$

for all $x \in \left[0, \overline{X}\right]$ and all $y \in \left[\underline{y}, \overline{y}\right]$. ■

We are now in a position to prove uniqueness.

Theorem 5 [Uniqueness]. *Suppose that: $\delta R < 1$, and that U is three times continuously differentiable on $(0, +\infty)$. Then, for all $x_0 \in [0, +\infty)$, there exists $\overline{\beta}_2 \in [0, 1)$ and $\overline{X} \in [x_0, +\infty)$ such that, for all $\beta \in [\overline{\beta}_2, 1]$, equilibrium is unique on $[0, \overline{X}]$.*

Notice that Theorem 5 is a local uniqueness theorem: the critical value $\overline{\beta}_2$ will in general depend on x_0. Local uniqueness is, however, all that we need: if initial cash-on-hand is x_0 and $\beta \in [\overline{\beta}_2, 1]$, then levels of cash-on-hand outside the interval $[0, \overline{X}]$ will not be observed in any equilibrium. We do not know whether Theorem 5 has a global analogue.

Proof. See the Appendix. ■

4.7. The Finite-Horizon Case: Robustness. By combining our existence and uniqueness results for the finite-horizon case with our regularity results, we can show that the equilibrium of the finite-horizon model depends continuously on the parameters U, f, β and δ. This leaves one parameter unaccounted for: T. This parameter plays an important role in empirical applications. For example, simulations of calibrated life-cycle models usually proceed by truncating the life cycle at some point. It is therefore crucial to verify that the equilibrium of the chosen model is robust with respect to the horizon chosen for the model.

The simplest way to establish robustness would be to show that there is a unique equilibrium of the infinite-horizon model. If we could show this, then it would follow at once from our regularity results that this equilibrium depended continuously on T. More precisely, note that T is chosen from the space $\mathbb{N} \cup \{\infty\}$. All the points of this space are isolated except for the point ∞, which is an accumulation point. By saying that the equilibrium depends continuously on T, we therefore mean that: there is a unique equilibrium when $T = \infty$; and, for all $\eta > 0$, there exists a $T_0 < \infty$ such that, for all $T > T_0$, the equilibrium of the model with horizon T is within η of the equilibrium of the model with horizon ∞. In other words, the choice of horizon for the model makes very little difference to the equilibrium provided that this horizon is sufficiently far into the future.

Unfortunately, the proof of Theorem 5 shows only that, if β is sufficiently close to 1, then there is a unique *stationary* equilibrium of the model. This leaves open two possibilities. First, there may be more than one stationary equilibrium if β is not close to 1. Secondly, there may be non-stationary equilibria. It may be very difficult to make progress with the first possibility: while it may be possible to identify other regions of parameter space in which there is a unique stationary equilibrium, it may not be true that there is a unique equilibrium for all choices of the parameters. After all, we are analyzing a game. It may, however, be possible to make progress with the second possibility: what is needed here is a proof that the Bellman operator is a contraction mapping. The proof of Theorem 5 falls short of this goal: it shows only that the Bellman operator is a contraction mapping when confined to the set of current-value functions of stationary equilibria.

Nonetheless, the available evidence suggests that life-cycle simulations are probably robust to the choice of horizon provided that β is sufficiently close to 1.

5. Generalized Euler Equation

In this section we discuss the hyperbolic analog of the standard Euler Relation.[11]

5.1. Heuristic derivation of the Hyperbolic Euler Relation. Suppose that C is an equilibrium consumption function. Adopt the perspective of self t. Since all future selves use the consumption function C, and since self t uses the same discount factor δ from period $t+1$ onwards, her continuation-value function V solves the

[11]The material from this section was first published in Harris and Laibson (2001).

recursive equation

$$V\left(x_{t+1}\right)=U\left(C\left(x_{t+1}\right)\right)+\mathrm{E}_{t+1}\left[\delta V\left(R\left(x_{t+1}-C\left(x_{t+1}\right)\right)+y_{t+2}\right)\right]. \tag{5}$$

Note that $V\left(x_{t+1}\right)$ is the expectation, conditional on x_{t+1}, of the present discounted value of the utility stream which starts in period $t+1$.

Self t uses discount factor $\beta\delta$ at time t. Her current-value function W therefore solves the equation

$$W\left(x_{t}\right)=U\left(C\left(x_{t}\right)\right)+\mathrm{E}_{t}\left[\beta\delta V\left(R\left(x_{t}-C\left(x_{t}\right)\right)+y_{t+1}\right)\right]. \tag{6}$$

Moreover

$$C\left(x_{t}\right)\in\underset{c\in[0,x_{t}]}{\operatorname{argmax}}\, U\left(c\right)+\mathrm{E}_{t}\left[\beta\delta V\left(R\left(x_{t}-c\right)+y_{t+1}\right)\right], \tag{7}$$

since consumption is chosen by the current self.

The first-order condition associated with (7) implies that

$$U'\left(C\left(x_{t}\right)\right)\geq\mathrm{E}_{t}\left[R\beta\delta V'\left(R\left(x_{t}-C\left(x_{t}\right)\right)+y_{t+1}\right)\right], \tag{8}$$

with equality if $C\left(x_{t}\right)<x_{t}$. The first-order condition and envelope theorem together imply that the shadow value of cash-on-hand equals the marginal utility of consumption:

$$W'\left(x_{t}\right)=U'\left(C\left(x_{t}\right)\right). \tag{9}$$

Finally, V and W are linked by the equation

$$\beta V\left(x_{t+1}\right)=W\left(x_{t+1}\right)-\left(1-\beta\right)U\left(C\left(x_{t+1}\right)\right). \tag{10}$$

These expressions can be combined to yield the Strong Hyperbolic Euler Relation. Indeed, we have

$$U'\left(C\left(x_t\right)\right) \geq \mathrm{E}_t\left[R\beta\delta V'\left(R\left(x_t - C\left(x_t\right)\right) + y_{t+1}\right)\right]$$

(this is just the first-order condition (8))

$$= \mathrm{E}_t\left[R\delta\left(W'(x_{t+1}) - (1-\beta)\,U'(C(x_{t+1}))C'(x_{t+1})\right)\right]$$

(differentiating equation (10) with respect to x_{t+1} and substituting in)

$$= \mathrm{E}_t\left[R\delta\left(U'(C(x_{t+1})) - (1-\beta)\,U'(C(x_{t+1}))C'(x_{t+1})\right)\right]$$

(from the analogue of equation (9) for self $t+1$). Rearranging yields

$$U'\left(C(x_t)\right) \geq \mathrm{E}_t\left[R\left(C'(x_{t+1})\beta\delta + (1 - C'(x_{t+1}))\delta\right)U'(C(x_{t+1}))\right], \tag{11}$$

with equality if $c < x_t$. This is the Hyperbolic Euler Relation.

When $\beta = 1$, this relation reduces to the well-known Exponential Euler Relation

$$U'\left(C(x_t)\right) \geq \mathrm{E}_t\left[R\,\delta\,U'(C(x_{t+1}))\right].$$

Intuitively, the marginal utility of consuming an additional dollar today, $U'(C_t)$, must equal the marginal utility of saving that dollar. A saved dollar grows to R dollars by next year. Utils next period are discounted with factor δ. Hence, the value of today's marginal savings is given by $\mathrm{E}_t\left[R\delta U'(C_{t+1})\right]$. The expectation operator integrates over uncertain future consumption.

The difference between the Hyperbolic Euler Relation and the Exponential Euler Relation is that, in the former, the constant exponential discount factor, δ, is replaced by the effective discount factor, namely

$$C'(x_{t+1})\beta\delta + (1 - C'(x_{t+1}))\delta.$$

This effective discount factor is a weighted average of the short-run discount factor $\beta\delta$ and the long-run discount factor δ. The respective weights are $C'(x_{t+1})$, the marginal propensity to consume out of liquid wealth, and $(1 - C'(x_{t+1}))$, the marginal propensity to consume out of liquid wealth. Since $\beta < 1$, the effective discount factor is stochastic and endogenous to the model.

In the sophisticated hyperbolic model, the effective discount factor is negatively related to the future marginal propensity to consume (MPC). To gain intuition for this effect, consider a consumer at time 0 who is thinking about saving a marginal dollar for the future. The consumer at time zero — 'self 0' — expects future selves to overconsume relative to the consumption rate that self 0 prefers those future selves to implement. Hence, on the equilibrium path, self 0 values marginal saving more than marginal consumption at any future time period. From self 0's perspective, therefore, it matters how a marginal unit of wealth at time period 1 will be divided between savings and consumption by self 1. Self 1's MPC determines this division. Since self 0 values marginal saving more than marginal consumption at time period 1, self 0 values the future less the higher the expected MPC at time period 1.

The effective discount factor in the Hyperbolic Euler Relation varies significantly with cash-on-hand. Consumers who expect to have low levels of future cash-on-hand will expect $C'(x_{t+1})$ to be close to one,[12] implying that the effective discount factor will approximately equal $\beta\delta$. Assuming that periods are annual with a standard calibration of $\beta = 0.7$ and $\delta = 0.95$, the effective discount rate would be $-\ln(0.7 \times 0.95) = 0.41$. By contrast, consumers with high levels of future cash-on-hand will

[12]Low levels of cash-on-hand imply that the agent is liquidity constrained. Hence, low levels of cash-on-hand imply a high MPC.

expect $C'(x_{t+1})$ to be close to zero,[13] implying that the effective discount factor will approximately equal δ. In this case, the effective discount rate will be $-\ln(0.95) = 0.05$. The simulations reported below confirm these claims about the shape of C.

5.2. Exact derivation. If the consumption function is discontinuous, then the derivation of the Hyperbolic Euler Relation given above is not valid. However, the consumption function is always of locally bounded variation. This property can be used to derive a weaker version of the Hyperbolic Euler Relation. This weaker version reduces to the Hyperbolic Euler Relation given above if the consumption function is Lipschitz continuous. Moreover, it can be shown that the consumption function is indeed Lipschitz continuous when β is sufficiently close to 1 (Harris and Laibson, 2001a).

6. Numerical solution and calibration of the model

We complement our theoretical analysis with numerical simulations. Numerical results help to build intuition and provide quantitative assessment of qualitative effects. In this section we describe our strategy for simulating the one-asset infinite-horizon model. The same broad strategy applies to the institutionally richer simulations that we describe in Section 8.

We calibrate our stripped-down model with the same parameter values used by Angeletos, Laibson, Repetto, Tobacman and Weinberg (2001a), hereafter ALRTW. Specifically, $\rho = 2$, $\beta = 0.7$, $\delta = 0.9571$, and $R = 1.0375$.[14]

To capture labor-income uncertainty, we adopt a shifted symmetric Beta density

[13] When the agent is not liquidity constrained, marginal consumption is approximately equal to the annuity value of marginal increments of wealth. Hence, the local slope of the consumption function is close to the real interest rate.

[14] ALRTW choose all of these parameters ex-ante, except δ. Then δ is chosen so that the simulated data matches the empirical median wealth to income ratio of 50-59 year old household heads. ALRTW also use this method to infer the preferences of exponential consumers ($\beta = 1$). They find that $\delta_{\text{exponential}} = .9437$.

with support $[\varepsilon, 1+\varepsilon]$:

$$f(Y) \propto (Y-\varepsilon)^{(a-1)}(1+\varepsilon-Y)^{(a-1)},$$

where $a > 0$ and ε is positive but close to zero. Hence, Y has mean $\frac{1}{2}+\varepsilon$. If $a > 1$ the density is bell-shaped, and continuous on $\mathbb{R}$. Moreover, if $a > 3$ then the density is twice continuously differentiable, and therefore satisfies the regularity conditions of Section 4.

We set $a = 5$, implying that $\frac{\sigma(Y)}{\bar{Y}} = 0.30$. This value is comparable to the value of $\frac{\sigma(Y)}{\bar{Y}}$ implied by standard income processes estimated from the Panel Survey of Income Dynamics. For example, ALRTW estimate a process for $\ln Y$ that has two components: an AR(1) process and i.i.d. noise.[15] Their empirically estimated process implies $\frac{\sigma(Y)}{\bar{Y}} = 0.32$.[16]

Figure 2 reports the equilibrium consumption function generated by our infinite-horizon one-asset simulation.[17] The function turns out to be continuous, monotonic, and concave. It appears smooth, except for the point at which the liquidity constraint begins to bind. In the next section, we identify cases in which these regularity properties cease to hold.

7. Properties of the consumption function

The consumption function in figure 2 is continuous, monotonic and concave. However, hyperbolic consumption functions need not have these desirable properties (Laibson

[15] Specifically, $\ln Y_t$ = [household fixed effects] + [polynomial in age] + $u_t + \eta_t$, where $u_t = \alpha u_{t-1} + \varepsilon_t$, and η_t and ε_t are white noise.

[16] This is an unconditional normalized standard deviation. The empirical conditional normalized standard deviation, $\frac{\sqrt{E_{t-1}(Y_t-\bar{Y}_t)^2}}{\bar{Y}}$, is .23.

[17] To numerically simulate our model, we adopt a numerical solution algorithm that does not interpolate between points in the state space. Specifically, our algorithm discretizes the state space, and forces the consumer to make choices that keep the state variables on the discrete partition. We believe that our algorithm successfully approximates the behavior that would arise in a continuous state space. Most importantly, we find that once our partition is made sufficiently fine, further refinement has no effect on our simulation results.

1997b, Morris and Postlewaite 1997, O'Donoghue and Rabin 1999a, Harris and Laibson 2001a, and Krusell and Smith 2000). In this section we characterize the general properties of the hyperbolic consumption function. We first discuss the kinds of pathologies that can arise. We then discuss the regularity conditions that eliminate these pathologies.

7.1. Pathologies: violations of continuity, monotonicity and concavity. To develop intuition for the existence of hyperbolic pathologies, we consider a finite-horizon version of the model of Section 3.[18] We assume that the stream of income is deterministic. We apply backwards induction arguments to solve for the equilibrium policies.

First, consider the strategy of self T. Trivially, self T sets $c_T = x_T$. Self T consumes all available cash on hand.

Now, consider the problem of self $T-1$. Self $T-1$ knows that any resources left to self T will be consumed by self T. So self $T-1$ chooses c_{T-1} to maximize

$$U(c_{T-1}) + \beta\delta U(c_T)$$

subject to the constraints

$$x_T = R\left(x_{T-1} - c_{T-1}\right) + y_T,$$

$$c_{T-1} \leq x_{T-1},$$

$$c_T = x_T.$$

The first constraint is the dynamic budget constraint. The second constraint is the liquidity constraint. The third constraint reflects the equilibrium strategy of self T.

Given this problem, it is straightforward to show that when the liquidity constraint

[18] See Laibson (1997b) for the original version of this example.

does not bind, self $T-1$ picks c_{T-1} such that

$$U'(c_{T-1}) = \beta\delta R U'(R\cdot(x_{T-1}-c_{T-1})+y_T)$$

When the liquidity constraint binds self $T-1$ sets $c_{T-1}=x_{T-1}$. Represents self $T-1's$ equilibrium policy function as $C_{T-1}(x_{T-1})$.

Now consider the problem of self $T-2$. Self $T-2$ chooses c_{T-2} to maximize

$$U(c_{T-2})+\beta\delta U(c_{T-1})+\beta\delta^2 U(c_T),$$

subject to the constraints

$$x_{T-1} = R\,(x_{T-2}-c_{T-2})+y_{T-1},$$

$$c_{T-2}\le x_{T-2},$$

$$c_{T-1} = C_{T-1}(x_{T-1})$$

$$c_T = x_T.$$

The first constraint is the dynamic budget constraint. The second constraint is the liquidity constraint. The third and fourth constraints represent the strategies of selves $T-1$ and T.

To develop intuition for the optimal policy of self $T-2$, consider the continuation value function of self $T-2$,

$$V_{T-1}(x_{T-1}) = u(C_{T-1}(x_{T-1}))+\delta u(R\,(x_{T-1}-C_T(x_{T-1}))+y_T).$$

From self $(T-2)$'s perspective, wealth at time $T-1$ has value $\beta\delta V_{T-1}(x_{T-1})$. There exists a threshold wealth level $x_{T-1}=\hat{x}$ at which the liquidity constraint for self $T-1$ ceases to bind. In the region to the left of $\hat{x}$ all marginal wealth is consumed in period $T-1$, implying,

$$V'_{T-1}(\hat{x}-) = U'(C_{T-1}(\hat{x})).$$

In the region to the right of $\hat{x}$ some marginal wealth is passed on to period T, implying,

$$V'_{T-1}(\hat{x}+) = C'_{T-1}(\hat{x}) \cdot U'(C_{T-1}(\hat{x})) + \delta R \left(1 - C'_{T-1}(\hat{x})\right) U'(C_T(\hat{x})).$$

Note that at $x_{T-1} = \hat{x}$, self $T-1$ is indifferent between marginal consumption in period $T-1$, and marginal consumption in period T. So,

$$U'(C_{T-1}(\hat{x})) = R\beta\delta U'(C_T(\hat{x})).$$

Substituting this relationship into the previous expression yields

$$V'_{T-1}(\hat{x}+) = \left[C'_{T-1}(\hat{x}) + \frac{1}{\beta}\left(1 - C'_{T-1}(\hat{x})\right)\right] U'(C_{T-1}(\hat{x}))$$

$$> U'(C_{T-1}(\hat{x})) = V'_{T-1}(\hat{x}-).$$

Hence the continuation value function V_{T-1} has a kink at $x_{T-1} = \hat{x}$. At this point the *slope* of the value function discretely rises.

This kink implies that the equilibrium consumption function of self $T-2$ will have a downward discontinuity. To understand why, note that self $T-2$ will never select a value of $c_{T-2} > 0$ such that $R(x_{T-2} - c_{T-2}) + y_{T+1} = x_{T-1} = \hat{x}$. If $x_{T-1} = \hat{x}$ *did* hold, self $T-2$ could raise her welfare by either cutting or raising consumption. If

$$U'(c_{T-2}) < \beta\delta R V'_{T-1}(\hat{x}+)$$

self $T-2$ could increase welfare by cutting consumption — with marginal cost

$U'(c_{T-2})$ — and raising saving — with marginal benefit $\beta\delta RV'_{T-1}(\hat{x}+)$. If

$$U'(c_{T-2}) \geq \beta\delta RV'_{T-1}(\hat{x}+)$$

self $T-2$ could increase welfare by raising consumption — with marginal benefit $U'(c_{T-2})$ — and lowering saving — with marginal cost $\beta\delta RV'_{T-1}(\hat{x}-) < \beta\delta RV'_{T-1}(\hat{x}+) \leq U'(c_{T-2})$.

Self $T-2$ makes equilibrium choices that avoid the region of low continuation marginal utilities — in the neighborhood to the left of $x_{T-1} = \hat{x}$ — by jumping to the region of high continuation marginal utilities to the right of $x_{T-1} = \hat{x}$. This avoidance can only be achieved with an equilibrium consumption function that has a discrete downward discontinuity. Figure 3 plots the equilibrium consumption functions for selves $T-2$, $T-1$, and T for the case in which the instantaneous utility function is isoelastic and $y_t = 1$ for all t.

Intuitively, the pathology described here arises because of a special kind of strategic interaction. Self $T-2's$ consumption function discontinuously declines because self $T-2$ has an incentive to push self $T-1$ over the wealth threshold $\hat{x}$ at which self $T-1$ has a kink in it's consumption function. Self $T-2$ is willing to discretely cut its own consumption to push $T-1$ over the $\hat{x}$ threshold, because the marginal returns to the right of $\hat{x}$ are greater than the marginal returns to the left of $\hat{x}$ *from self* $(T-2)$'s *perspective.*

If this example were extended another period, we could also demonstrate that the optimal choices of self $T-3$ will violate the Hyperbolic Euler Equation. Finally, all of these pathologies would continue to arise even if a small amount of smooth noise were added to the income process.

7.2. Sufficient conditions for continuity, monotonicity, and concavity of the consumption function. The previous subsection provides an example of the kinds of pathologies that can arise in hyperbolic models. However, these pathologies do *not* arise when the model is calibrated with empirically sensible parameter values (see Figure 2). In this section, we identify the parameter regions that generate the pathologies.

First, when β is close to one, the discontinuities and non-monotonicities vanish. Harris and Laibson (2001a) prove this claim formally. Intuitively, when β is close to one, the hyperbolic consumption function converges to the exponential consumption function, which is continuous and monotonic.[19] Likewise, when β is close to one the hyperbolic consumption function matches the concavity of the exponential consumption function. Carroll and Kimball (1996) provide sufficient conditions for exponential concavity (U in the HARA class), though they do not handle the case of binding liquidity constraints.

Figure 4 graphically demonstrates the comparative static on β. We plot the consumption functions generated by β values $\{0.1, 0.2., 0.3, \ldots, 0.7\}$.[20] The consumption functions are vertically shifted so they do not overlap. Recall that $\beta = 0.7$ corresponds to our benchmark calibration. As β falls below 0.4, the consumption function becomes increasingly irregular. However, regularity returns as β falls to zero: in a neighborhood of $\beta = 0$, the consumption function coincides with the 45 degree line.

Pathologies are also controlled by the curvature of the consumption function. Our simulation results imply that increasing ρ eliminates irregularities. Figure 5 graphically demonstrates the comparative static on ρ. We plot the consumption functions generated by ρ values $\{0.5, 0.75, 1, 1.25\}$.[21] The consumption functions are again vertically shifted. Recall that $\rho = 2$ corresponds to our benchmark calibration. As ρ falls below 1.25, the consumption function becomes increasingly irregular. The irregularities increase as ρ falls, since low curvature augments the feedback effects that engender the irregularities. Specifically, when the utility function is relatively

[19] All of the convergence results apply to an absorbing interval of x values. See section 4 for a definition and discussion of such absorbing intervals.

[20] We adopt the baseline parameter values: $a = 5$, $\delta = .9571$, $\rho = 2$, $R = 1.0375$.

[21] We adopt the baseline parameter values: $a = 5$, $\beta = .7$, $\delta = .9571$, $R = 1.0375$.

less bowed, it is relatively less costly to strategically cut consumption today to push future selves over critical wealth thresholds.

Finally, decreasing the variance of the income process increases the degree of irregularity. Figure 6 graphically demonstrates the comparative static on a. We plot the consumption functions generated by a values $\{25, 50, 100, 200, 400\}$.[22] These a values correspond to $\frac{\sigma(Y)}{\bar{Y}}$ values of $\{0.14, 0.10, 0.07, 0.05, 0.04, 0.03\}$. Recall that $a = 5$ (i.e., $\frac{\sigma(Y)}{\bar{Y}} = 0.30$) corresponds to our benchmark calibration. As a rises above 25 (i.e., $\frac{\sigma(Y)}{\bar{Y}}$ falls below 0.14), the consumption function becomes increasingly irregular. The irregularities increase as a increases, since high a values correspond to low levels of income volatility. Low volatility makes it easier for early selves to predict future wealth levels, and to strategically push later selves over critical wealth thresholds.

In summary, irregularities vanish when β is close to one, risk aversion is high, and uncertainty is high. At the benchmark calibration, the pathologies do not arise. Moreover, our model omits some sources of uncertainty that would only reinforce the regularity of our benchmark consumption functions. For example, our model omits shocks to preferences and asset return uncertainty.[23]

8. Consumption Applications

A series of papers have analyzed the positive and normative implications of the hyperbolic buffer-stock model: Laibson, Repetto, and Tobacman (1998), Laibson, Repetto, and Tobacman (2000), Angeletos, Laibson, Repetto, Tobacman, Weinberg (2001a, 2001b). These papers extend the precautionary saving models pioneered by Carroll (1992, 1997), Deaton (1991), and Zeldes (1989b).[24] We will focus our discussion on

[22] We adopt the baseline parameter values: $\beta = .7$, $\delta = .9571$, $\rho = 2$, $R = 1.0375$.

[23] Asset return uncertainty has an advantage over labor income uncertainty, since the volatility generated by noisy returns scales up with the level of wealth. With sufficient asset uncertainty, it should be possible to establish that regularity applies to the entire domain of cash-on-hand, instead of just an absorbing interval.

[24] See also Hubbard, Skinner, and Zeldes (1994, 1995), Engen, Gale, and Scholz (1994), and Gourinchas and Parker (1999).

the work of Laibson, Repetto, and Tobacman (2000), hereafter LRT, and Angeletos, Laibson, Repetto, Tobacman and Weinberg (2001a), hereafter ALRTW.

The ALRTW model incorporates most of the features of previous lifecycle simulation models and adds new features, including credit cards, time-varying household size, and illiquid assets. We summarize the key features of the ALRTW model below. A more general version of the model, and a complete description of the calibration, appears in LRT.[25]

8.1. Model summary. Households are divided into three levels of educational attainment. We only discuss simulation results for the largest group, households whose head has only a high school degree (roughly half of US households). The simulations have been replicated for households in other educational categories, and the conclusions are quantitatively similar (see LRT).

Households face a time-varying, exogenous hazard rate of survival. Households live for a maximum of 90 periods, beginning economic life at age 20, and retiring at age 63. The retirement age is calibrated to match reported retirement ages from the Panel Survey of Income Dynamics (PSID). Household composition — number of adults and non-adults — varies exogenously over the life-cycle (also calibrated to match the PSID).

Log income, $\ln Y_{it}$, is modeled as the sum of a polynomial in age and two stochastic components: an autocorrelated component and an i.i.d. component. Different processes are estimated during the working life and during retirement (using the PSID).

Households may hold liquid assets, X_t, and illiquid assets, Z_t. Since labor income is liquid wealth, $X_t + Y_t$ represents total liquid asset holdings at the beginning of period t. Credit card borrowing is modelled as a negative value for X_t. Credit card

[25]This more general model allows consumers to declare bankruptcy and allows the consumer to borrow against illiquid collateral (e.g., mortgages on housing).

borrowing must not exceed a credit limit equal to some fraction of current (average) income. Specifically, $X_t \geq -\lambda \cdot \bar{Y}_t$, where $\bar{Y}_t$ is cohort average income at age t, and $\lambda = 0.30$ (calibrated from the 1995 SCF). The real after-tax interest rate on liquid assets is 3.75%. The real interest rate on credit-card loans is 11.75%, two percentage points *below* the mean debt-weighted real interest rate reported by the Federal Reserve Board. This low value is chosen to implicitly capture the effect of bankruptcy. Actual annual bankruptcy rates of roughly one percent per year imply that the effective interest rate is at least one percentage point below the observed interest rate.

The illiquid asset generates consumption flows equal to 5% of the value of the asset ($Z_t \geq 0$). Hence, the holding return on illiquid assets is considerably higher than the return on other assets. However, the illiquid asset can only be sold with a transaction cost.

Households have isoelastic preferences with a coefficient of relative risk aversion of $\rho = 2$. Self t has instantaneous payoff function

$$u(C_t, Z_t, n_t) = n_t \cdot \frac{\left(\frac{C_t + \gamma Z_t}{n_t}\right)^{1-\rho} - 1}{1 - \rho}$$

Note that γZ_t represents the consumption flow generated by Z_t ($\gamma = 0.05$), and n_t is the effective household size, $n_t = ([\#\ \text{adults}_t] + 0.4[\#\ \text{of children}_t])$.

Households have either an exponential discount function ($\delta^t_{\text{exponential}}$) or a quasi-hyperbolic discount function ($\beta\delta^t_{\text{hyperbolic}}$, with $\beta = 0.7$). ALRTW assume that the economy is either populated exclusively by exponential households or exclusively by hyperbolic households. ALRTW pick $\delta_{\text{exponential}}$ and $\delta_{\text{hyperbolic}}$ to match empirical levels of retirement saving. Specifically, $\delta_{\text{exponential}}$ is picked so that the exponential simulations generate a median wealth to income ratio of 3.2, for individuals between ages 50 and 59. The median of 3.2, is calibrated from the SCF.[26] The hyperbolic

[26]Wealth does not include social security wealth and other defined benefit pensions, which are

discount factor, $\delta_{\text{hyperbolic}}$, is also picked to match the empirical median of 3.2.[27]

The discount factors that replicate the SCF wealth to income ratio are .9437 for the exponential model and .9571 for the hyperbolic model. Since hyperbolic consumers have two sources of discounting — β and δ — the hyperbolic δ's lie above the exponential δ's. Recall that the hyperbolic and exponential discount functions are calibrated to generate the same amount of pre-retirement wealth accumulation. In this manner the calibrations "equalize" the underlying willingness to save between the exponential and hyperbolic consumers. The calibrated long-term discount factors are sensible when compared to discount factors that have been used in similar exercises by other authors. Finally, note that these discount factors do not include mortality effects, which reduce the respective discount factors by an additional one percent on average per year.

8.2. Simulation results of ALRTW. Calibrated hyperbolic simulations — $\beta = 0.7$, $\delta = 0.957$ — generate lifecycle consumption profiles that closely match the lifecycle consumption profiles generated by calibrated exponential simulations — $\beta = 1$, $\delta = 0.944$. For example, Figure 7 compares hyperbolic and exponential consumption means over the lifecycle. These two hump-shaped profiles are very similar.[28] The only differences arise around retirement and at the very beginning and end of life. At the beginning of life, hyperbolic consumers go on a credit-card-financed spending spree,[29] leading to higher consumption than the exponentials. Around retirement,

---

already built into the model in the form of post-retirement "labor income."

[27]For calibration purposes, total wealth is measure as $X+Z+\frac{Y}{24}$, where X represents liquid assets (excluding current labor income), Z represents illiquid assets, and Y represents annual after-tax labor income. The $\frac{Y}{24}$ is included to reflect average cash-inventories used for (continuous) consumption out of labor income. If labor income is paid in equal monthly installments, $\frac{Y}{12}$, and consumption is smoothly spread over time, then average cash inventories will be $\frac{Y}{24}$.

[28]The consumption profiles roughly track the mean labor income profile. This low frequency comovement is driven by two factors. First, low income early in life holds down consumption, since consumers don't have large credit lines. Second, consumption needs peak in mid-life, when the number of adult-equivalent dependents peaks at age 47.

[29]See Gourinchas and Parker (1999) for empirical evidence on the early life consumption boom.

hyperbolic consumption falls more steeply than exponential consumption, since hyperbolic households have most of their wealth in illiquid assets, which they cannot cost-effectively sell to smooth consumption. At the end of life, hyperbolic consumers have more illiquid assets to sell, slowing down the late-life collapse in consumption.

The total wealth profiles of hyperbolics and exponentials are also similar. This correspondence is not surprising, since the hyperbolic and exponential simulations are each calibrated to match the observed level of retirement wealth accumulation in the Survey of Consumer Finances. However, the two models generate very different simulated allocations across liquid and illiquid assets. Just before retirement (age 63), the average liquid asset holding of the simulated hyperbolics households is only about \$10,000, while the exponential households have accumulated over \$45,000 in liquid wealth (1990 dollars).[30] Hyperbolics end up holding relatively little liquid wealth because liquidity tends to be splurged to satisfy the hyperbolic taste for instant gratification. Both naive and sophisticated hyperbolics will quickly spend whatever liquidity is at their disposal.

By contrast, hyperbolics hold much more *illiquid* wealth than their exponential counterparts. Just before retirement, the average illiquid asset holding of the simulated hyperbolics is \$175,000 compared with \$130,000 for the exponentials. Hyperbolics are more willing to hold illiquid wealth for two reasons. First, sophisticated hyperbolics (like the hyperbolics in these simulations), view illiquid assets as a commitment device, which they value since it prevents later selves from splurging saved wealth too quickly. Second, illiquid assets are particularly valuable to hyperbolics (both naifs and sophisticates) since hyperbolics have lower long-run discount rates than exponentials. Hence, hyperbolics place relatively greater value on the long-run stream of payoffs associated with illiquid assets.[31] Hyperbolics and exponentials

[30]For the purposes of the analysis in this subsection, simulated liquid assets are measured as $X^+ + \frac{Y}{24}$, where X^+ represents positive holdings of liquid assets (excluding current labor income).

[31]The long-run discount rate of a hyperbolic consumer, $-\ln(\delta_{\text{hyperbolic}}) = -\ln(.957) = .044$, is

dislike illiquidity for the standard reason that illiquid assets can't be used to buffer income shocks. But, this cost of illiquidity is partially offset for hyperbolics for the two reasons described above: hyperbolics value commitment and hyperbolics more highly value the long-run dividends of illiquid assets. Hence, on net, illiquidity is less costly for a hyperbolic than for an exponential consumer.

To empirically evaluate the asset allocation predictions of the hyperbolic and exponential models, ALRTW compare the simulated results to survey evidence from the Survey of Consumer Finances. For example, ALRTW analyze the percentage of households that have at least one month of liquid wealth on hand. On average, 73% of simulated exponential households hold liquid assets greater than one month of labor income. The analogous number for hyperbolics is only 40%. For comparison, 42% of households in the Survey of Consumer Finances hold liquid financial assets greater than one month of labor income.

ALRTW also evaluate the models by analyzing the simulated quantity of liquid assets as a share of total assets. In the SCF the average liquid wealth share is only 8% and neither the exponential nor hyperbolic simulations match this number, though the hyperbolic simulations are a bit closer to the mark. The average liquid wealth share for simulated hyperbolic households is 31%. The analogous exponential liquid wealth share is 50%.

Revolving credit — e.g., credit card borrowing — represents another important form of liquidity. Low levels of liquid assets are naturally associated with high levels of credit card debt. ALRTW contrast exponential and hyperbolic consumers by comparing their simulated propensities to borrow on credit cards.[32] At any point in time 51% of hyperbolic consumers borrow on their credit cards, compared to only 19% of exponentials. In the 1995 Survey of Consumer Finances, 70% of households

calibrated to lie below the long-run discount rate of an exponential consumer, $-\ln(\delta_{\text{exponential}}) = -\ln(.944) = .058$.

[32]See LRT (2000) for a much more detailed analysis of credit card borrowing.

with credit cards report that they did not fully pay their credit card bill the last time that they mailed in a payment. Hyperbolic simulations come much closer to matching these self-reports. Likewise, the simulated hyperbolic consumers borrow much more on average than the simulated exponential consumers. On average, simulated exponential households owe \$900 of interest-paying credit card debt, *including* the households with no debt. By contrast, simulated hyperbolic households owe \$3,400 of credit card debt. The actual amount of credit card debt owed per household with a credit card is approximately \$4,600 (*including* households with no debt, but *excluding* the float).[33]

Euler Equation tests have played a critical role in the empirical consumption literature since the work of Hall (1978). Many of the papers in this literature have asked whether lagged information predicts current consumption growth. In particular, many authors have tried to determine whether predictable changes in income predict changes in consumption:

$$\Delta \ln(C_{it}) = \alpha \, \mathrm{E}_{t-1} \, \Delta \ln(Y_{it}) + X_{it}\beta + \varepsilon_{it} \tag{12}$$

Here X_{it}, is a vector of control variables. The standard consumption model (without liquidity constraints), predicts $\alpha = 0$; the marginal propensity to consume out of predictable changes in income should be zero. By contrast, empirical estimates of α lie above 0 with "consensus estimates" around $\alpha = 0.2$.[34]

[33] This average balance includes households in all education categories. It is calculated on the basis of aggregate information reported by the Federal Reserve. This figure is consistent with values from a proprietary account-level data set assembled by David Gross and Nicholas Souleles (1999a, 1999b, 2000). See LRT (2000).

[34] For example, Hall and Mishkin (1982) report a statistically significant coefficient of .200, Hayashi (1985) reports a significant coefficient of .158, Altonji and Siow (1987) report an insignificant coefficient of .091, Attanasio and Weber (1993) report an insignificant coefficient of .119, Attanasio and Weber (1995) report an insignificant coefficient of .100, Shea (1995) reports a marginally significant coefficient of .888, Lusardi (1996) reports a significant coefficient of .368, Souleles (1999) reports a significant coefficient of .344, and ALRTW (2000) report a significant coefficient of .285. See Deaton (1992) and Browning and Lusardi (1996) for discussion of the excess sensitivity literature.

ALRTW estimate the standard comovement regression using *simulated* data. For the hyperbolic simulations, the coefficient on $\mathrm{E}_{t-1}\,\Delta\ln(Y_{it})$ is $\alpha = 0.17$. By contrast, the exponential simulations generate a value of $\alpha = 0.03$. Hyperbolic consumers hold more of their wealth in illiquid form than exponentials. So hyperbolics are more likely to hit liquidity constraints, raising their marginal propensity to consume out of predictable changes in income.

The hyperbolic simulations also predict income-consumption comovement around retirement. Banks et al (1998) and Bernheim et al (1997) argue that consumption anomalously falls during the mid-60's, at the same time that workers are retiring and labor income is falling. ALRTW estimate the following regression to explore the consumption drop at retirement:

$$\Delta\ln(C_{it}) = I_{it}^{\mathrm{RETIRE}}\gamma + X_{it}\beta + \varepsilon_{it}$$

Here I_{it}^{RETIRE} is a set of dummy variables that take the value of one in periods $t-1$, t, $t+1$ and $t+2$ if period t is the age of retirement, and X_{it} is a vector of control variables. Summing the coefficients on the four dummy variables (and switching signs), generates an estimate of the "excess" drop in consumption around retirement. Estimating these coefficients from the PSID, yields a statistically significant excess drop of 11.6% around retirement. The analogous drop for simulated hyperbolic consumers is 14.5%, while the drop for simulated exponential consumers is only 3.0%. Hyperbolic consumers hold relatively little liquid wealth. A drop in income at retirement translates into a substantial drop in consumption, even though retirement is an exogenous, completely predictable event.

All in all, the hyperbolic model consistently does a better job of approximating the data. Table 1 draws these findings together.

Table 1:

	Hyperbolic	Exponential	Data
% with $\frac{\text{liquid}}{Y} > \frac{1}{12}$	40%	73%	42%
mean $\frac{\text{liquid assets}}{\text{liquid + illiquid assets}}$	0.39	0.50	0.08
% borrowing on "Visa"	51%	19%	70%
mean borrowing	\$3400	\$900	\$4600
C-Y comovement	0.17	0.03	$\approx$ 0.20
% C drop at retirement	14.5%	3.0%	11.6%

Source: Angeletos et al (2001b).

9. Naifs vs. Sophisticates

Until now, we have considered the case in which early selves hold correct expectations about the preferences and behavior of later selves. Early selves anticipate that later selves will fail to maximize the patient long-run interests of early selves. When early selves hold such correct expectations, they are referred to as sophisticates (Strotz, 1956).

However, it is reasonable to imagine that early selves might mistakenly expect later selves to follow through on the early selves' best intentions. This is the naive case, discussed by Strotz (1956), Akerlof (1991), and O'Donoghue and Rabin (1999a, 1999b, 2000). Such naifs have optimistic forecasts in the sense that they believe that future selves will carry out the wishes of the current self. Under this belief, the current self constructs the sequence of actions that maximizes the preferences of the current self. The current self then implements the first action in that sequence, expecting future selves to implement the remaining actions. Instead, those future selves conduct their own optimization and therefore implement actions in conflict with the patient behavior anticipated by prior selves.

In some cases, the behavior of naive hyperbolics is very close to the behavior of sophisticated hyperbolics. For example, ALRTW have replicated their calibration and analysis under the assumption that hyperbolic consumers are naive. They find that the naive hyperbolics act effectively the same as the sophisticated hyperbolics discussed above. Hence, for the consumption applications in this paper, it does not matter whether we assume that hyperbolics are naive or sophisticated.

However, this rough equivalence does not generally hold. Ted O'Donoghue and Matthew Rabin have written a series of papers (1999a, 1999b, 2000) that examine the differences between naifs and sophisticates, developing examples where naifs and sophisticates behave in radically different ways. Their most recent paper explores the issue of retirement saving. They show that naive hyperbolics may perpetually postpone asset reallocation decisions, generating sizeable welfare costs. Each one-period postponement seems optimal, because the naif mistakenly expects some future self to undertake the reallocation.

Naifs models do not exhibit any of the pathologies that we have discussed above (e.g., non-monotonic consumption functions). If consumers do not recognize that their own preferences are dynamically inconsistent, they will not have any incentive to act strategically vis-à-vis their own future selves. However, this solution to the pathology problem requires that consumers be completely naive about their own future preferences. Any partial knowledge of future dynamic inconsistency reinstates the pathologies.

O'Donoghue and Rabin (2000) also propose an intermediate model in which decision-makers partially recognize their propensity to be hyperbolic in the future. Specifically, in this intermediate model the actor believes that future selves will have a β value equal to $\hat{\beta}$. Sophisticates hold correct expectations about the future value of β, so $\hat{\beta} = \beta$. Naifs incorrectly believe that future selves will hold preferences consistent with the long-run interests of the current self, implying $\hat{\beta} = 1$. Partial naifs, lie between these extremes, so $\beta < \hat{\beta} < 1$.

10. Normative analysis and policy implications

Welfare and policy analysis can be problematic in hyperbolic models. The crux of the difficulty is the lack of a clear welfare criterion.

The most traditional perspective has been adopted by Phelps and Pollak (1968) and Laibson (1996, 1997a). These authors take the multiple self framework literally, and simply apply the Pareto criterion for welfare analysis. If one allocation makes all selves as least as well off as another allocation, then the former allocation Pareto dominates the latter allocation. Even this very strong welfare criterion opens the door to interesting welfare analysis. It is typically the case that the equilibrium allocation in a hyperbolic model is Pareto-inferior to other feasible allocations that will not arise in equilibrium. These Pareto-dominant allocations can only be attained with a commitment technology. We turn to such commitment technologies (and the corresponding policies that support them) in the next subsection.

O'Donoghue and Rabin adopt a different approach to welfare analysis. They argue that the right welfare perspective is the long-run perspective. Specifically, they rank allocations using the welfare of an agent with no hyperbolicity (i.e., $\beta = 1$). In the long-run, all selves discount exponentially. So all past selves want future selves to discount exponentially. In this sense, $\beta = 1$ is the right discounting assumption if we adopt the preferences of some "earlier" self (say at birth). Another way to motivate this welfare criterion is to ask what discount function you would advise someone else to use. Typically, we urge others to act patiently, suggesting that we normatively discourage short-run impulsivity. In the language of these models, this advice recommends $\beta = 1$.

Recently, Caplin and Leahy (2000) have suggested another criterion. They take the multiple self framework literally, and suggest a utilitarian approach. Specifically, they argue that a sensible welfare criterion would weight the welfare of all of the selves. This approach produces challenging implications. Specifically, if later selves get roughly the same weight as early selves, then late consumption should matter much more than early consumption. To see why, consider the following two-period example. Self one cares about periods one and two (with equal weights). Self two cares only about period two. Then period two should get twice the weight of period one in the social planner's welfare function. Late consumption benefits both selves whereas early consumption benefits only self one.

At the moment there is no concensus framework for measuring welfare in multiple self models. However, the different approaches reviewed above usually give similar answers to policy questions. All of the competing welfare criteria imply that equilibrium allocations in economies without commitment typically generate savings rates that are too low (i.e., higher savings allocations would improve social welfare). This implication follows almost immediately once one adopts a welfare criterion in which $\beta = 1$ (O'Donoghue and Rabin) or once one adopts the utilitarian perspective of Caplin and Leahy. Equilibrium allocations also tend to be Pareto-inferior because the static gains of high consumption rates in the short-run (gains to the current self) tend to be overwhelmed by the dynamic losses of low savings rates in the long-run steady state (dynamic losses to the expected utility of the current self). Recall that hyperbolic consumers have low long-run discount rates. Hence, the long-run outcomes matter a great deal to the welfare of the current hyperbolic consumer (Laibson 1996). A commitment to a savings rate slightly above the equilibrium savings rate will raise the welfare of all selves.

10.1. The value of commitment. Sophisticated hyperbolic consumers are motivated to choose policies that commit the behavior of future selves. Moreover, such commitment devices can raise the welfare of all selves if the commitment locks in patient long-run behavior. Even naive consumers will benefit from commitment, though they won't appreciate these benefits at the time they are being locked into a patient behavioral regime. However, these naive agents may not mind such commitments (ex-ante) because they incorrectly expect future selves to act patiently anyway.

In a world of sophisticated hyperbolic consumers, the social planner's goal is to make commitment possible, rather than imposing it on consumers.[35] Sophisticated consumers understand the value of commitment and will adopt such commitments when it is in their interest. Hence, a 401(k), which is voluntary, might be viewed as a useful commitment device for a sophisticated hyperbolic consumer.[36] Laibson (1996) and Laibson, Repetto, and Tobacman (1998) measure the welfare consequences of providing voluntary commitment technologies, like 401(k)'s, to sophisticated hyperbolic consumers.

By contrast, in a world of unsophisticated consumers (i.e., naifs), a benevolent government may want to impose commitment on consumers.[37] Social security, with its universal coverage, and illiquid "balances" can be viewed as such a commitment.

11. EXTENSIONS

11.1. Asset uncertainty. Our simulation results reported in section 7 demonstrate that hyperbolic consumption functions become less irregular as more noise is added to the model. The analysis in section 7 explores the case in which the noise

[35] Commitment technologies typically make all selves better off.

[36] 401(k)'s are defined contribution pension accounts available in most US firms. These accounts have a penalty for "early" withdrawal (e.g., before age $59\frac{1}{2}$).

[37] Naturally, there are excellent reasons to be wary of activist government. Much political activity is directed toward rent seeking. Moreover, even a benevolent social planner needs to worry about the disincentives and distortions that arise when well-intentioned politicians tax productive activities to pay for new social programs.

comes from stochastic labor income. Another natural source of noise is the asset return process. In the analysis above we assumed that the asset return process was deterministic.

Incorporating random returns into the model will generate four likely benefits. First, when pathologies (e.g., non-monotonic consumption functions) do arise, those pathologies will probably be less pronounced when asset returns are stochastic. Second, pathologies will be less likely to arise in the first place. Third, once asset return variability is added to the model we may be able to prove more general theorems. For example, *without* asset return variability we can show that as $\beta \to 1$ the consumption function becomes monotonic and continuous on an absorbing interval of cash-on-hand. An absorbing interval is a range of cash-on-hand values which, in equilibrium, the consumer will never leave. With asset return variability we conjecture that we will be able to show that as $\beta \to 1$ the consumption function becomes monotonic and continuous on the *entire* state space. This more general theorem reflects the fact that asset return uncertainty scales up with financial wealth, in contrast to labor income uncertainty which does not scale with financial wealth. Finally, adding asset uncertainty will enable us to model multi-asset state spaces as long as each asset has some idiosyncratic variability. In this setting, we expect to be able to prove the existence, uniqueness, and regularity of equilibria using variants of the techniques developed in Harris and Laibson (2001a).

11.2. Continuous-time hyperbolic models. Continuous-time modelling provides a more robust way of eliminating pathologies like non-monotonic consumption functions (Harris and Laibson, 2001b). To motivate the continuous-time formalism, recall the discrete time set-up. In the standard discrete-time formulation of quasi-hyperbolic preferences the present consists of the single period t. The future consists of periods $t+1$, $t+2$, ... A period n steps into the future is discounted with factor δ^n and an additional discount factor β is applied to all periods except the present.

This model can be generalized in two ways. First, the present can last for any number of periods $T_t \in \{1, 2, ...\}$. Secondly, T_t can be random. The preferences in equation (13) below are a natural continuous-time analogue of this more general formulation. Specifically, the preferences of self t are given by

$$\mathrm{E}_t \left[\int_t^{t+T_t} e^{-\gamma(s-t)} U\left(c\left(s\right)\right) ds + \alpha \int_{t+T_t}^{+\infty} e^{-\gamma(s-t)} U\left(c\left(s\right)\right) ds \right], \tag{13}$$

where $\gamma \in (0, +\infty)$, $\alpha \in (0, 1]$, $U : (0, +\infty) \to \mathbb{R}$ and T_t is distributed exponentially with parameter $\lambda \in [0, +\infty)$. In other words, self t uses a stochastic discount function, namely

$$D_\lambda(t, s) = \left\{ \begin{array}{ll} e^{-\gamma(s-t)} & \text{if} \quad s \leq t + T_t \\ \alpha e^{-\gamma(s-t)} & \text{if} \quad s > t + T_t \end{array} \right\}.$$

This stochastic discount function decays exponentially at rate γ up to time $t + T_t$, drops discontinuously at $t + T_t$ to a fraction α of its level just prior to $t + T_t$, and decays exponentially at rate γ thereafter. Figure 8 plots a single realization of this discount function, with $t = 0$ and $T_t = 3.4$. Figure 9 plots the expected value of the discount function, namely

$$\mathrm{E}_t D_\lambda(t, s) = e^{-\lambda(s-t)} e^{-\gamma(s-t)} + (1 - e^{-\lambda(s-t)}) \alpha e^{-\gamma(s-t)},$$

for $\lambda \in \{0, 0.1, 1, 10, \infty\}$.

This continuous-time formalization is close to the deterministic functions used in Barro (1999) and Luttmer and Mariotti (2000). However, Harris and Laibson (2001b) assume that T_t is stochastic. The stochastic transition with constant hazard rate reduces the problem to a system of two differential equations that characterize present and future value functions.

When $\lambda = 0$ the discount function is equivalent to a standard exponential discount function. As $\lambda \to \infty$ the discount function converges to a jump function, namely

$$D_{\infty}(t,s) = \left\{ \begin{array}{lll} 1 & \text{if} & s = t \\ \alpha e^{-\gamma(s-t)} & \text{if} & s > t \end{array} \right\}.$$

This limit case is both analytically tractable and psychologically relevant. In this "instantaneous-gratification" case the present is vanishingly short. Individuals prefer consumption in the present instant discretely more than consumption in the momentarily delayed future. The lessons from this model carry over, by continuity, to the neighborhood of models in which the present is short, but not precisely instantaneous (i.e., λ large).

The instantaneous-gratification model, which is dynamically *inconsistent*, shares the same value function as a related dynamically *consistent* optimization problem with a wealth-contingent utility function. Using this partial equivalence, Harris and Laibson (2001b) prove that the hyperbolic equilibrium exists and is unique. The associated equilibrium consumption functions are continuous and monotonic in wealth. The monotonicity property relies on the condition that the long-run discount rate is weakly greater than the interest rate. For this case, all of the pathological properties of discrete-time hyperbolic models are eliminated.

12. CONCLUSIONS

We have characterized the consumption behavior of hyperbolic consumers. The hyperbolic model provides two payoffs. First, it provides an analytically tractable, parsimonious foundation with which to analyze self-control problems. Second, it is easily calibrated, providing precise numerical predictions that can be empirically evaluated in competition with mainstream models. We have shown that the hyperbolic model successfully matches empirical observations on household balance sheets and consumption choices.

Relative to exponential households, hyperbolic households hold low levels of liquid wealth measured either as a fraction of labor income or as a share of total wealth. Hyperbolic households borrow more aggressively in the revolving credit market (i.e., on credit cards), but they save more actively in illiquid assets. Because the hyperbolic households have low levels of liquid assets and high levels of credit-card debt, they are unable to smooth their consumption paths in the presence of predictable changes in income. Calibrated hyperbolic simulations explain observed levels of consumption-income comovement and the drop in consumption at retirement. Calibrated hyperbolic simulations generate 'excess-sensitivity' coefficients of approximately 0.20, very close to empirical coefficients estimated from household data.

More generally, the hyperbolic model provides a good formal foundation for the study of self-defeating behaviors. Economists usually assume that rational agents will act in their own interests. Hyperbolic agents may hold rational expectations, but they will rarely make efficient choices. Puzzling and important self-defeating behaviors like undersaving, overeating, and procrastination lose some of their mystery when analyzed with the hyperbolic model.

13. APPENDIX

Proof of Theorem 5. Fix $x_0 \in [0, +\infty)$. Suppose that W_1 and W_2 are two equilibrium current-value functions, and let S_1 and S_2 be the associated saving functions. Put $h = U \circ g$. Then

$$\begin{aligned}
(\mathfrak{B}W_1)(x) &= U(x - S_1(x)) + \delta \int (W_1 - \varepsilon h \circ W_1')(RS_1(x) + y) f(y)\, dy \\
&= U(x - S_1(x)) + \delta \int (W_2 - \varepsilon h \circ W_2')(RS_1(x) + y) f(y)\, dy \\
&\quad + \delta \int (W_1 - W_2)(RS_1(x) + y) f(y)\, dy \\
&\quad - \varepsilon\delta \int (h \circ W_1' - h \circ W_2')(RS_1(x) + y) f(y)\, dy \\
&\leq (\mathfrak{B}W_2)(x) + \delta \int (W_1 - W_2)(RS_1(x) + y) f(y)\, dy \\
&\quad - \varepsilon\delta \int (h \circ W_1' - h \circ W_2')(RS_1(x) + y) f(y)\, dy.
\end{aligned}$$

Hence, in order to obtain an upper bound for $(\mathfrak{B}W_1)(x) - (\mathfrak{B}W_2)(x)$, it suffices to estimate the expressions

$$\int (W_1 - W_2)(RS_1(x) + y) f(y)\, dy \tag{14}$$

and

$$\int (h \circ W_1' - h \circ W_2')(RS_1(x) + y) f(y)\, dy. \tag{15}$$

In doing so, we shall make use of the estimates of Theorem 3, which apply to W_1 and W_2. In particular, the constant K and the functions N_1 and N_2 used below are taken from that theorem.

Expression (14) is easy to estimate. Since S_1 is an equilibrium saving function,

$RS_1(x) + y \in [\underline{y}, \overline{X}]$ for all $x \in [\underline{y}, \overline{X}]$ and all $y \in [\underline{y}, \overline{y}]$. Hence

$$\left| \int (W_1 - W_2)(RS_1(x) + y) f(y)\, dy \right| \le \|W_1 - W_2\|_{\mathcal{C}([\underline{y}, \overline{X}])}$$

for all $x \in [\underline{y}, \overline{X}]$, where

$$\|W_1 - W_2\|_{\mathcal{C}([\underline{y}, \overline{X}])} = \sup_{x \in [\underline{y}, \overline{X}]} |W_1(x) - W_2(x)|.$$

Expression (15) requires more care. Put

$$W'_\phi(x) = (1 - \phi) W'_1(x) + \phi W'_2(x).$$

Then

$$h(W'_2(x)) - h(W'_1(x)) = \int_0^1 \left((W'_2 - W'_1) h' \circ W'_\phi \right)(x)\, d\phi$$

and

$$\begin{aligned}
& \int (h \circ W'_2 - h \circ W'_1)(RS_1(x) + y) f(y)\, dy \\
= & \int \left(\int_0^1 \left((W'_2 - W'_1) h' \circ W'_\phi \right)(RS_1(x) + y) f(y)\, d\phi \right) dy \\
= & \int_0^1 \left(\int \left((W'_2 - W'_1) h' \circ W'_\phi \right)(RS_1(x) + y) f(y)\, dy \right) d\phi.
\end{aligned}$$

Moreover

$$\begin{aligned}
& \int \left((W'_2 - W'_1) h' \circ W'_\phi \right)(RS_1(x) + y) f(y)\, dy \\
= & - \int \left((W_2 - W_1) h' \circ W'_\phi \right)(RS_1(x) + y) f'(y)\, dy \\
& - \int \left((W_2 - W_1) h'' \circ W'_\phi \right)(RS_1(x) + y) f(y) W''_\phi (RS_1(x) + dy)
\end{aligned}$$

(on integrating by parts). Hence, in order to estimate expression (15), we need to estimate the expressions

$$\int \left(\left(W_2 - W_1 \right) h' \circ W'_\phi \right) \left(RS_1 \left(x \right) + y \right) f' \left(y \right) dy \tag{16}$$

and

$$\int \left(\left(W_2 - W_1 \right) h'' \circ W'_\phi \right) \left(RS_1 \left(x \right) + y \right) f \left(y \right) W''_\phi \left(RS_1 \left(x \right) + dy \right). \tag{17}$$

Expression (16) can be estimated as follows. First, note that, since S_1 is an equilibrium saving function, $RS_1 \left(x \right) + y \in \left[\underline{y}, \overline{X} \right]$ for all $x \in \left[\underline{y}, \overline{X} \right]$ and all $y \in \left[\underline{y}, \overline{y} \right]$. Secondly, put

$$\underline{\lambda} = \min_{x \in \left[\underline{y}, \overline{X} \right]} U' \left(x \right) \text{ and } \overline{\lambda} = \max_{x \in \left[\underline{y}, \overline{X} \right]} U' \left(x \right) \vee N_1 \left(x \right).$$

Thirdly, note that, since $W'_\phi \in [U', U' \vee N_1]$ for all $\phi \in [0,1]$, $W'_\phi \in \left[\underline{\lambda}, \overline{\lambda} \right]$ for all $\phi \in [0,1]$. Hence

$$\begin{aligned} & \left| \int \left(\left(W_2 - W_1 \right) h' \circ W'_\phi \right) \left(RS_1 \left(x \right) + y \right) f' \left(y \right) dy \right| \\ \leq \; & \left\| W_1 - W_2 \right\|_{\mathcal{C}\left(\left[\underline{y}, \overline{X} \right] \right)} \left\| h' \right\|_{\mathcal{C}\left(\left[\underline{\lambda}, \overline{\lambda} \right] \right)} \left\| f' \right\|_{\mathcal{C}\left(\left[\underline{y}, \overline{y} \right] \right)} \left(\overline{y} - \underline{y} \right) \end{aligned}$$

for all $x \in \left[\underline{y}, \overline{X} \right]$.

Expression (17) can be estimated as follows. First, as in the case of expression (16), we have

$$\begin{aligned} & \left| \int \left(\left(W_2 - W_1 \right) h'' \circ W'_\phi \right) \left(RS_1 \left(x \right) + y \right) f \left(y \right) W''_\phi \left(RS_1 \left(x \right) + dy \right) \right| \\ \leq \; & \left\| W_1 - W_2 \right\|_{\mathcal{C}\left(\left[\underline{y}, \overline{X} \right] \right)} \left\| h'' \right\|_{\mathcal{C}\left(\left[\underline{\lambda}, \overline{\lambda} \right] \right)} \left\| f \right\|_{\mathcal{C}\left(\left[\underline{y}, \overline{y} \right] \right)} \left\| W''_\phi \right\|_{\mathcal{TV}\left(\left[\underline{y}, \overline{X} \right] \right)} \end{aligned}$$

for all $x \in \left[\underline{y}, \overline{X}\right]$, where $\left\|W_\phi''\right\|_{\mathcal{TV}\left(\left[\underline{y},\overline{X}\right]\right)}$ denotes the total variation of the measure W_ϕ'' on the interval $\left[\underline{y}, \overline{X}\right]$. Secondly, put

$$\mu\left(dx\right) = KN_2\left(x\right)dx \text{ and } \widetilde{W}_\phi''\left(dx\right) = W_\phi''\left(dx\right) + \mu\left(dx\right).$$

Then

$$\begin{aligned}\left\|W_\phi''\right\|_{\mathcal{TV}\left(\left[\underline{y},\overline{X}\right]\right)} &= \left\|\widetilde{W}_\phi'' - \mu\right\|_{\mathcal{TV}\left(\left[\underline{y},\overline{X}\right]\right)} \leq \left\|\widetilde{W}_\phi''\right\|_{\mathcal{TV}\left(\left[\underline{y},\overline{X}\right]\right)} + \left\|\mu\right\|_{\mathcal{TV}\left(\left[\underline{y},\overline{X}\right]\right)} \\ &= \int_{\left[\underline{y},\overline{X}\right]} \widetilde{W}_\phi''\left(dx\right) + \int_{\left[\underline{y},\overline{X}\right]} \mu\left(dx\right) = \int_{\left[\underline{y},\overline{X}\right]} W_\phi''\left(dx\right) + 2\int_{\left[\underline{y},\overline{X}\right]} \mu\left(dx\right)\end{aligned}$$

(because $\widetilde{W}_\phi''$ and μ are both positive measures, and by definition of $\widetilde{W}_\phi''$)

$$= W_\phi'\left(\overline{X}+\right) - W_\phi'\left(\underline{y}-\right) + 2\int_{\left[\underline{y},\overline{X}\right]} KN_2\left(x\right)dx \leq \overline{\lambda} - \underline{\lambda} + 2\int_{\left[\underline{y},\overline{X}\right]} KN_2\left(x\right)dx.$$

Combining our estimates for expressions (14), (16) and (17), we obtain

$$\begin{aligned}&\left(\mathfrak{B}W_1\right)\left(x\right) - \left(\mathfrak{B}W_2\right)\left(x\right) \leq \delta\left\|W_1 - W_2\right\|_{\mathcal{C}\left(\left[\underline{y},\overline{X}\right]\right)} \\ &\quad + \varepsilon\delta\left\|W_1 - W_2\right\|_{\mathcal{C}\left(\left[\underline{y},\overline{X}\right]\right)}\left\|h'\right\|_{\mathcal{C}\left(\left[\underline{\lambda},\overline{\lambda}\right]\right)}\left\|f'\right\|_{\mathcal{C}\left(\left[\underline{y},\overline{y}\right]\right)}\left(\overline{y} - \underline{y}\right) \\ &\quad + \varepsilon\delta\left\|W_1 - W_2\right\|_{\mathcal{C}\left(\left[\underline{y},\overline{X}\right]\right)}\left\|h''\right\|_{\mathcal{C}\left(\left[\underline{\lambda},\overline{\lambda}\right]\right)}\left\|f\right\|_{\mathcal{C}\left(\left[\underline{y},\overline{y}\right]\right)}\left(\overline{\lambda} - \underline{\lambda} + 2\int_{\left[\underline{y},\overline{X}\right]} KN_2\left(x\right)dx\right)\end{aligned}$$

for all $x \in \left[\underline{y}, \overline{X}\right]$. Combining this estimate with the analogous estimate for $\left(\mathfrak{B}W_2\right)\left(x\right) - \left(\mathfrak{B}W_1\right)\left(x\right)$, we obtain

$$\left\|\mathfrak{B}W_1 - \mathfrak{B}W_2\right\|_{\mathcal{C}\left(\left[\underline{y},\overline{X}\right]\right)} \leq \delta\left(1 + \varepsilon L\right)\left\|W_1 - W_2\right\|_{\mathcal{C}\left(\left[\underline{y},\overline{X}\right]\right)},$$

where

$$\begin{aligned} L &= \|h'\|_{\mathcal{C}([\underline{\lambda},\overline{\lambda}])} \|f'\|_{\mathcal{C}([\underline{y},\overline{y}])} (\overline{y}-\underline{y}) \\ &\quad + \|h''\|_{\mathcal{C}([\underline{\lambda},\overline{\lambda}])} \|f\|_{\mathcal{C}([\underline{y},\overline{y}])} \left(\overline{\lambda} - \underline{\lambda} + 2 \int_{[\underline{y},\overline{X}]} KN_2(x)\, dx \right). \end{aligned}$$

It follows that, if

$$\varepsilon < \min \left\{ 1 - \overline{\beta}_1, \frac{1-\delta}{\delta L} \right\},$$

then $\|\mathfrak{B}W_1 - \mathfrak{B}W_2\|_{\mathcal{C}([\underline{y},\overline{X}])} = 0$. In other words, W_1 and W_2 coincide on $[\underline{y}, \overline{X}]$. ■

14. References

Ainslie, George. 1992 *Picoeconomics.* Cambridge: Cambridge University Press.

Akerlof, George A. 1991. "Procrastination and Obedience." *American Economic Review. (Papers and Proceedings),* pp. 1-19.

Altonji, Joseph, and Aloysius Siow. 1987. "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data." *Quarterly Journal of Economics.* 102(2), pp. 293-328.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001a) "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation" National Bureau of Economic Research working paper.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg (2001b) "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation" Forthcoming, *Journal of Economic Perspectives*, August.

Attanasio, Orazio. 1999. "Consumption," in *Handbook of Macroeconomics.* John Taylor and Michael Woodford, eds. North Holland.

———, and Guglielmo Weber. 1993. "Consumption Growth, the Interest Rate, and Aggregation." *Review of Economic Studies.* 60:3, pp. 631-49.

———, and ———. 1995. "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence form the Consumer Expenditure Survey." *Journal of Political Economy.* 103:6, pp. 1121-57.

Banks, James, Richard Blundell and Sarah Tanner. 1998. "Is There a Retirement Puzzle?" *American Economic Review.* 88(4), pp.769-788.

Barro, Robert. 1999. "Laibson Meets Ramsey in the Neoclassical Growth Model." *Quarterly Journal of Economics*, Vol 64(4), 1125-1152.

Benabou, Roland and Jean Tirole. 2000. "Willpower and personal rules," mimeo.

Bernheim B. Douglas, Jonathan Skinner, and Steven Weinberg. 1997. "What Accounts for the Variation in Retirement Wealth Among US Households?" Working Paper 6227. Cambridge, Mass: National Bureau of Economic Research.

Blundell, Richard, Martin Browning, and Costas Meghir. 1994. "Consumer Demand and the Life-Cycle Allocation of Household Expenditures." *Review of Economic Studies.* 61, pp. 57-80.

Brocas, Isabelle and Juan Carrillo, 2000. "The Value of Information When Preferences are Dynamically Inconsistent," forthcoming in *European Economic Review*, Vol.44, 1104-1115.

Brocas, Isabelle and Juan Carrillo, 2001. "Rush and Procrastination under Interdependent Activities" *Journal of Risk and Uncertainty* (forthcoming).

Browning, Martin and Annamaria Lusardi. 1996. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature.*32, pp.1797-1855.

Carrillo, Juan and Mathias Dewatripont, 2000. "Promises, promises,...," mimeo.

Carrillo, Juan and Thomas Mariotti. 2000. "Strategic Ignorance as a Self-Disciplining Device" *Review of Economic Studies*, Vol.67, 529-544.

Carroll, Christopher D. 1992. "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence." *Brookings Papers on Economic Activity.* 2:1992, pp. 61-156.

———. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics.* 112, pp. 1-57.

_____ and Miles Kimball. 1996. "On the Concavity of the Consumption Function," *Econometrica,* 64(4), pp. *981-992*

Chung, Shin-Ho, and Richard J. Herrnstein. 1961. "Relative and Absolute Strengths of Response as a Function of Frequency of Reinforcement." *Journal of the Experimental Analysis of Animal Behavior.* 4, pp. 267-72.

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica.* 59, pp. 1221-48.

Della Vigna, Stephano and Ulrike Malmendier, 2001. "Long Term Contracts and Self Control," mimeo.

Della Vigna, Stephano and Daniel Paserman, 2000. "Job Search and Hyperbolic Discounting," mimeo.

Diamond, Peter and Botond Koszegi. 1998. "Hyperbolic discounting and retirement." MIT mimeo.

Engen, Eric, William Gale, and John Karl Scholz. 1994. "Do Saving Incentives Work?" *BPEA.* 1:1994, pp. 85-180.

Frederick, Shane, George Loewenstein, and Edward O'Donoghue, 2001. "Time Discounting: A Critical Review," mimeo.

Gourinchas, Pierre-Olivier, and Jonathan Parker. 1999. "Consumption Over the Life-Cycle." Mimeo.

Gross, David and Nicholas Souleles. 1999a. "An Empirical Analysis of Personal Bankruptcy and Delinquency." Mimeo.

________ and ______. 1999b. "How do people use credit cards?" Mimeo.

________ and ______. 2000. "Consumer Response to Changes in Credit Supply: Evidence from Credit Card Data." Mimeo.

Gruber, Jonathan and Botond Koszegi, 2000. "Is Addiction Rational? Theory and Evidence," *Quarterly Journal of Economics*, forthcoming.

Hall, Robert E. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy.* Vol. 86 (6), 971-87.

Hall, Robert E. and Frederic S. Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates From Panel Data on Households." *Econometrica.* 50(2), pp. 461-81.

Harris, Christopher and David Laibson, 2001a. "Dynamic Choices of Hyperbolic Consumers." Forthcoming *Econometrica.*

Harris, Christopher and David Laibson, 2001b. "Instantaneous Gratification," mimeo.

Hayashi, Fumio. 1985. "The Permanent Income Hypothesis and Consumption Durability: Analysis based on Japanese Panel Data." *Quarterly Journal of Economics.* 100(4), pp.1083-1113.

Hubbard, Glenn, Jonathan Skinner, and Stephen Zeldes. 1994. "The Importance of Precautionary Motives in Explaining Individual and Aggregate Saving." *Carnegie-Rochester Conference Series on Public Policy.* 40: June, pp. 59-125.

———, ——, ———. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy.* 103, pp. 360-99.

King, G. R., and A. W. Logue. (1987) "Choice in a Self-Control Paradigm with Human Subjects: Effects of Changeover Delay Duration," *Learning and Motivation,* 18, 421-438.

Kirby, Kris N. (1997) "Bidding on the Future: Evidence Against Normative Discounting of Delayed Rewards," *Journal of Experimental Psychology,* 126, 54-70.

________ and R. J. Herrnstein. (1995) "Preference Reversals Due to Myopic Discounting of Delayed Reward," *Psychological Science,* 6, 2, 83-89.

________ and Nino N. Marakovic. (1995) "Modeling Myopic Decisions: Evidence for Hyperbolic Delay-Discounting within Subjects and Amounts," *Organizational Behavior and Human Decision Processes,* 64, 1, 22-30.

________ and ________. (1996) "Delayed-Discounting Probabilistic Rewards Rates Decrease as Amounts Increase," *Psychonomic Bulletin and Review,* 3, 1, 100-104.

Krusell, Per and Anthony Smith, 2000. "Consumption and Savings Decisions with Quasi-Geometric Discounting" mimeo.

Krusell, Per, Burhanettin Kuruscu, and Anthony Smith, 2000a. "Equilibrium Welfare and Government Policy with Quasi-Geometric Discounting" mimeo.

Krusell, Per, Burhanettin Kuruscu, and Anthony Smith, 2000b. "Asset Pricing with Quasi-Geometric Discounting" mimeo.

Laibson, David I. 1994. "Self-Control and Savings." Ph.D. diss., Massachusetts Institute of Technology.

———. 1996. "Hyperbolic Discounting, Undersaving, and Savings Policy." Working Paper 5635. Cambridge, Mass: National Bureau of Economic Research.

———. 1997a. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*. 62: 2, pp. 443-478.

———. 1997b. "Hyperbolic Discount Functions and Time Preference Heterogeneity." Harvard mimeo.

———. 1998. Comments on "Personal Retirement Saving Programs and Asset Accumulation," by James M. Poterba, Steven F. Venti, and David A. Wise, in *Studies in the Economics of Aging*. David A. Wise, ed. Cambridge, Mass: NBER and the University of Chicago Press, pp. 106-24.

———, Andrea Repetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement." *BPEA*. 1, pp. 91-196.

———, ——— and ———. 2000. "A Debt Puzzle." mimeo.

Loewenstein, George, and Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics*. 57, pp. 573-98.

Lusardi, Annamaria. 1996. "Permanent Income, Current Income, and Consumption; Evidence From Two Panel Data Sets." Journal of Business and Economic Statistics. 14, pp.81-90.

Luttmer, Erzo and Thomas Mariotti. 2000. "Subjective Discount Factors," mimeo.

Millar, A., and D. J. Navarick. (1984) "Self-Control and Choice in Humans: Effects of Video Game Playing as a Positive Reinforcer," *Learning and Motivation.* 15, pp.203-218.

Morris, Stephen and Andrew Postlewaite. 1997. "Observational Implications of Nonexponential Discounting," mimeo.

Mulligan, Casey. 1997. "A Logical Economist's Argument Against Hyperbolic Discounting." University of Chicago mimeo.

Navarick, D. J. (1982) "Negative Reinforcement and choice in Humans," *Learning and Motivation.* 13, pp. 361-377.

O'Donoghue, Ted, and Matthew Rabin. 1999a. "Doing It Now or Later." *American Economic Review,* 89(1), pp. 103-124.

______ and _______ 1999b. "Incentives for Procrastinators," *Quarterly Journal of Economics*, 114(3), pp. 769-816.

______ and _______ 2000. "Choice and Procrastination," working paper.

Parker, Jonathan A. 1999. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes", *American Economic Review.*

Phelps, E. S., and R. A. Pollak. 1968. "On Second-best National Saving and Game-equilibrium Growth." *Review of Economic Studies.* 35, pp. 185-199.

Ramsey, Frank 1928. "A Mathematical Theory of Saving," *Economic Journal* 38 (December), pp. 543-559.

Rankin, Deborah M. (1993) "How to Get Ready for Retirement: Save, Save, Save," *New York Times,* (3/13/1993) p. 33.

Read, Daniel, George Loewenstein, Shobana Kalyanaraman, and Adrian Bivolaru. (1996) “Mixing Virtue and Vice: The Combined Effects of Hyperbolic Discounting and Diversification,” Working paper, Carnegie Mellon University.

Runkle, David. 1991. “Liquidity Constraints and the Permanent-Income Hypothesis: Evidence From Panel Data.” *Journal of Monetary Economics.* 27(1), pp.73-98.

Shapiro, Matthew D., and Joel Slemrod. 1995. “Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding.” American Economic Review. 85:1, pp. 274-83.

Shea, John. 1995. “Union Contracts and the Life-Cycle/Permanent Income Hypothesis.” *American Economic Review.* 85(1), pp.186-200.

Simmons Market Research Bureau. 1996. *The 1996 Study of Media and Markets.* New York.

Souleles, Nicholas. 1999. “The Response of Household Consumption to Income Tax Refunds.” *American Economic Review.*

Strotz, Robert H. 1956. “Myopia and Inconsistency in Dynamic Utility Maximization.” *Review of Economic Studies.* 23, pp. 165-180.

Thaler, Richard H. 1981. “Some Empirical Evidence on Dynamic Inconsistency.” *Economic Letters.* 8, pp. 201-207.

------- 1992. “Saving, Fungibility, and Mental Accounts,” in *The Winner’s Curse,* Princeton University Press: Princeton, NJ. pp. 107-121.

———, and H. M. Shefrin. 1981. “An Economic Theory of Self-Control.” *Journal of Political Economy* 89, pp. 392-410.

Zeldes, Stephen P. 1989a. "Consumption and Liquidity Constraints: an Empirical Investigation." *Journal of Political Economy.* 97:2, pp. 305-346.

_________. 1989b. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics.* 104:2, pp. 275-98.

Figure 1: Exponential and hyperbolic discount functions

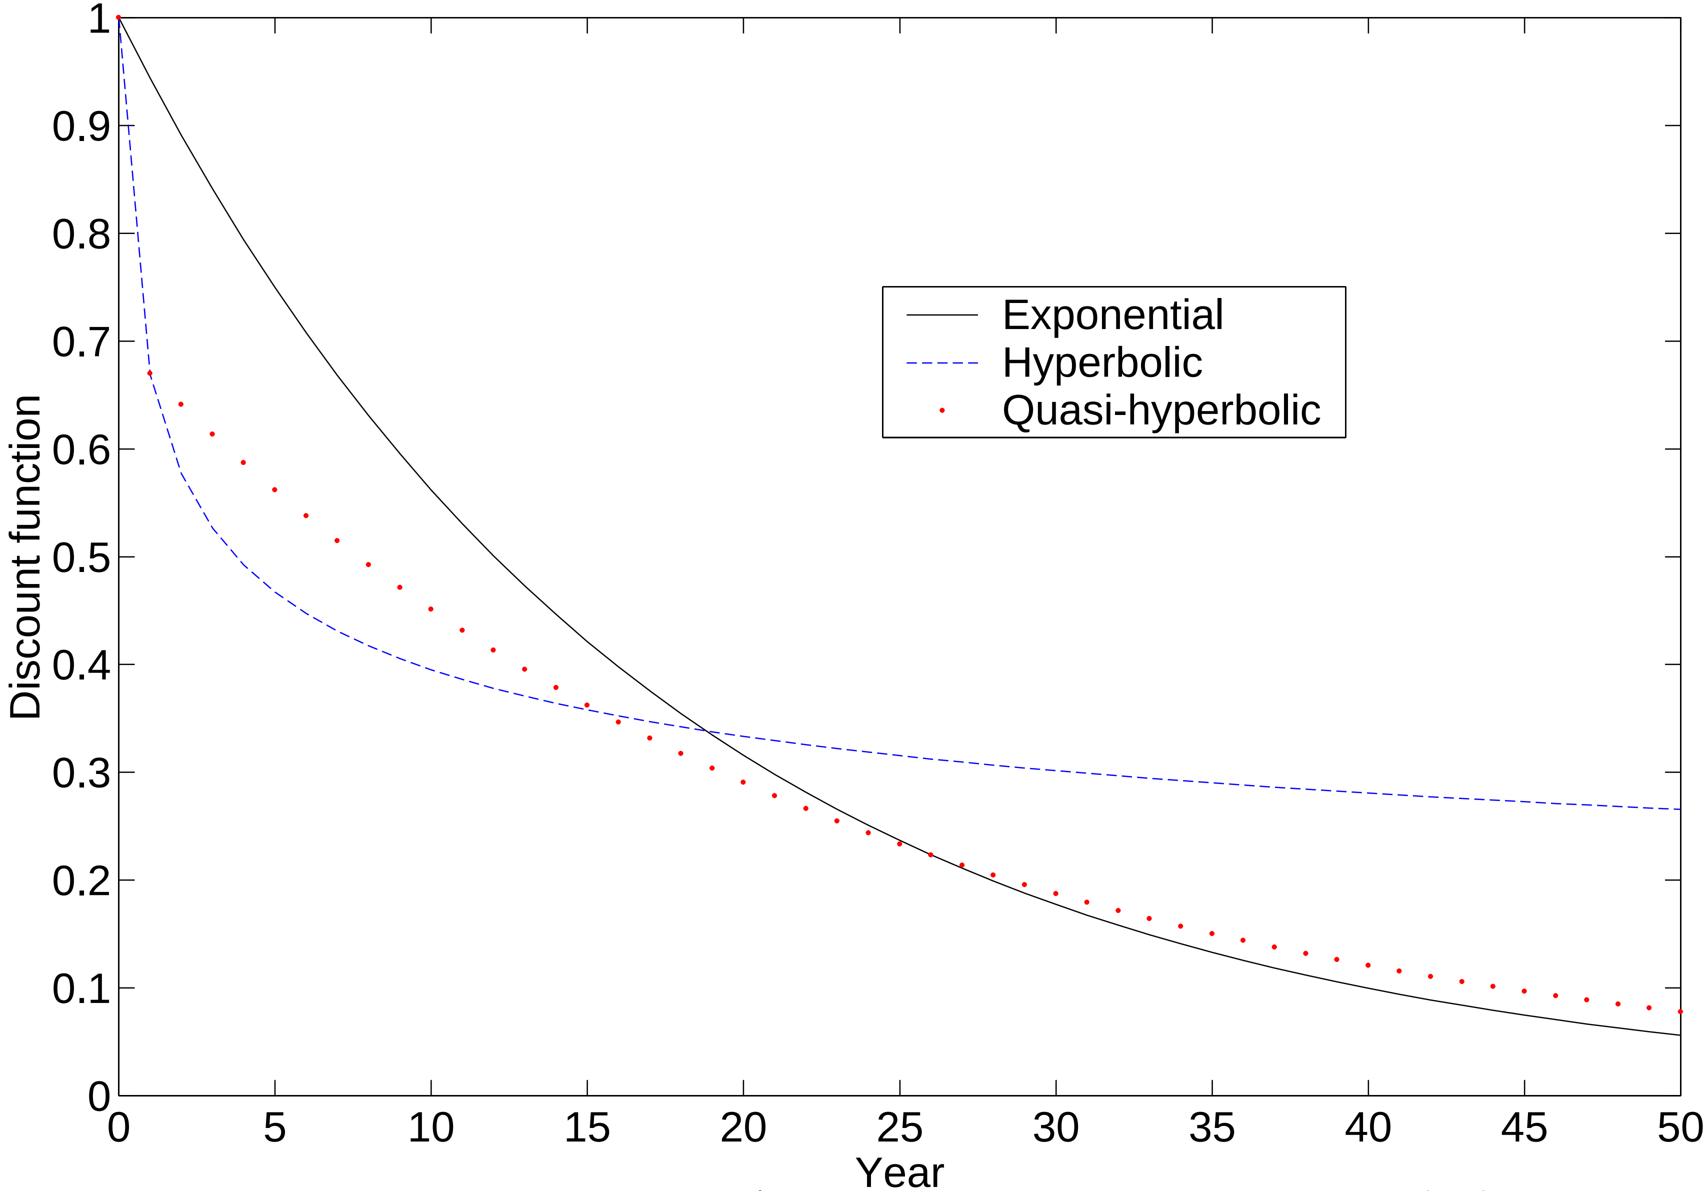

Exponential: δ^t, with δ=.944. Hyperbolic: $(1+\alpha t)^{-\gamma/\alpha}$, with α=4 and γ=1. Quasi-hyperbolic: $\{1,\beta\delta,\beta\delta^2,\beta\delta^3,...\}$, with β=.7 and δ=.957.

Figure 2: Calibrated consumption function

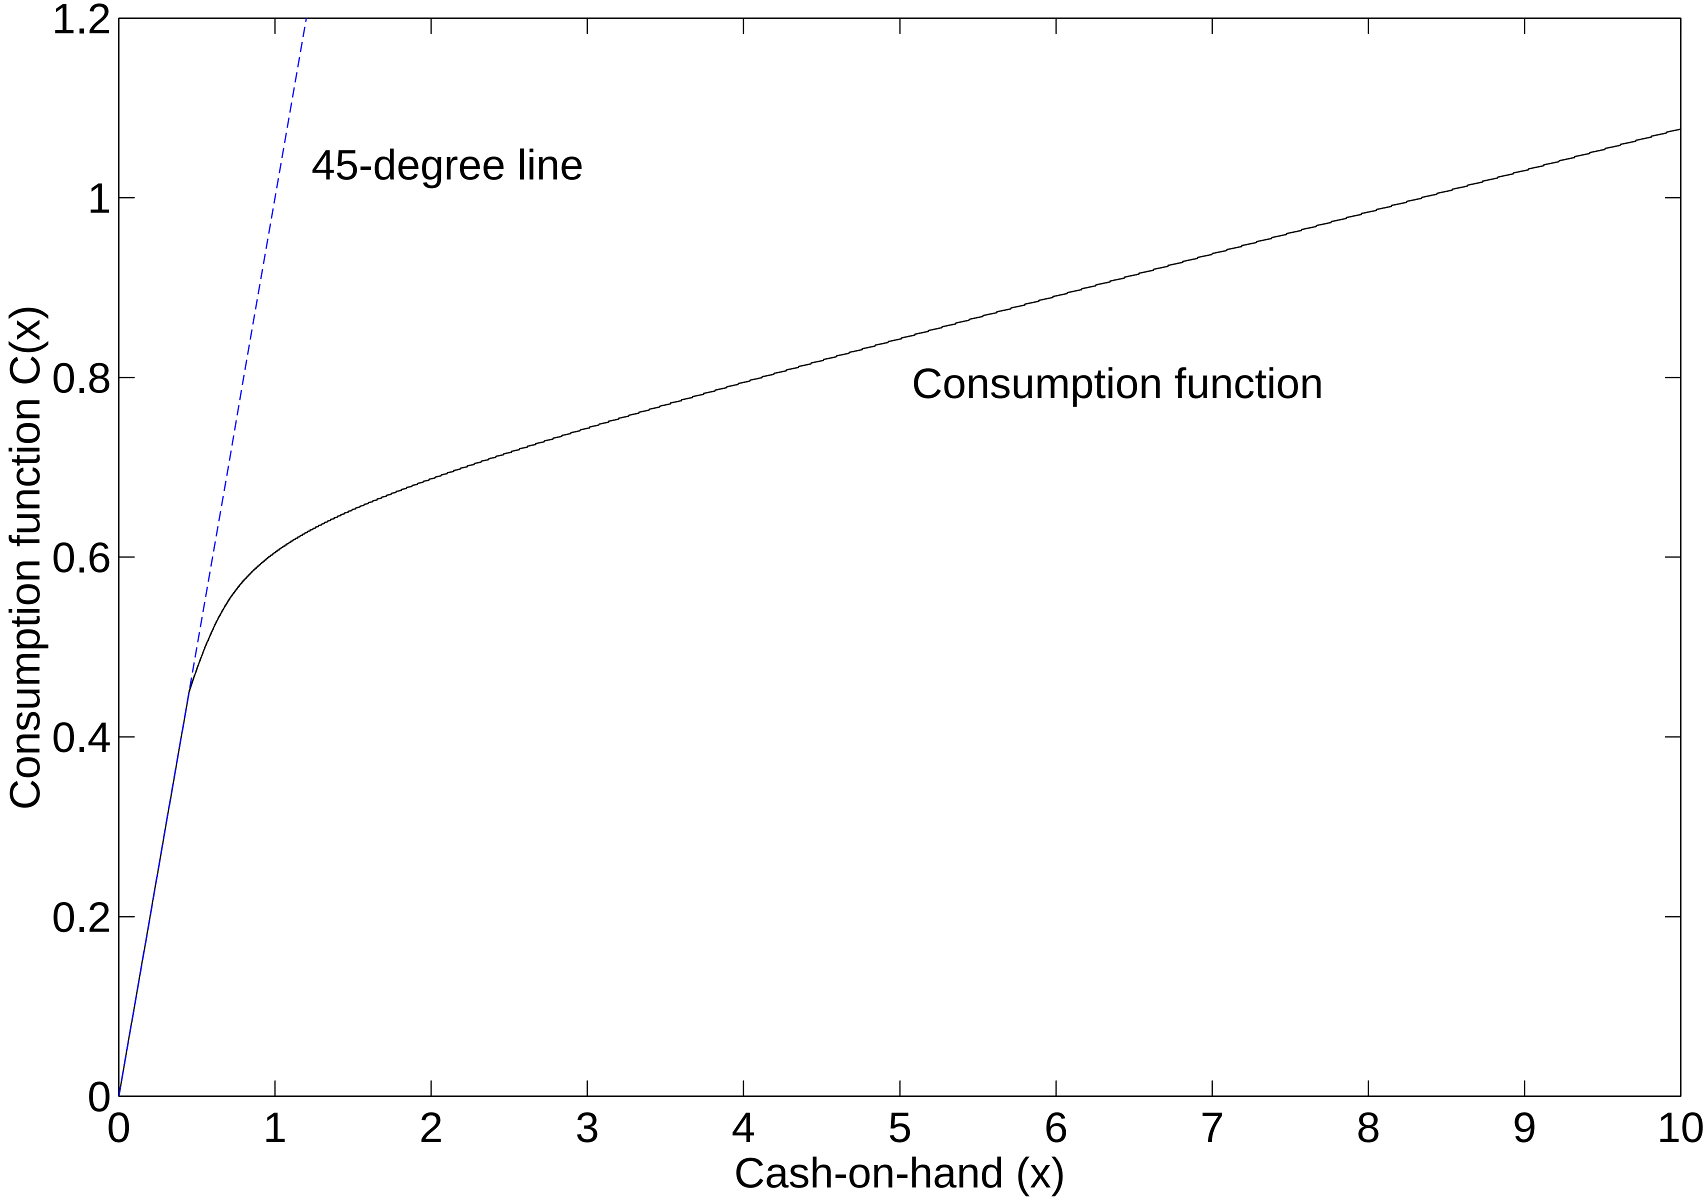

The consumption function is based on simulations in which β = .7, δ = .9571, ρ = 2, R = 1.0375, a = 5.

Figure 3: Consumption functions in periods T-2, T-1 and T

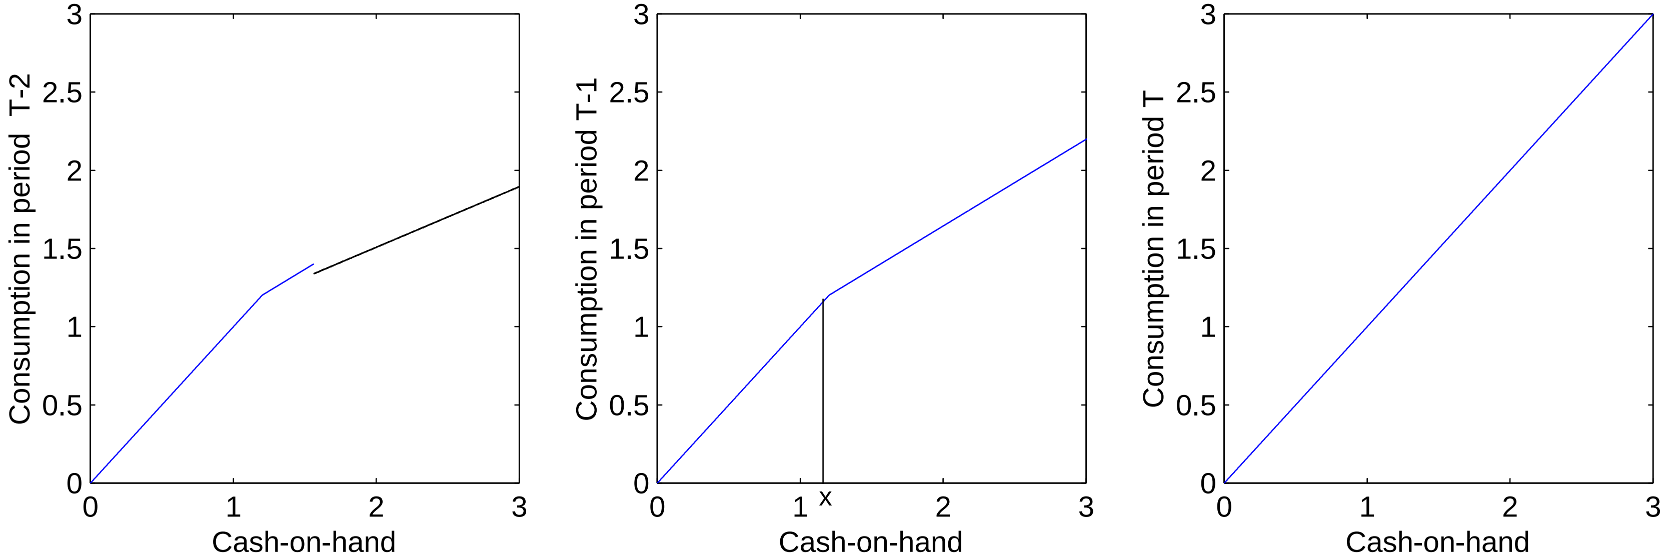

The consumption functions are based on simulations in which β = .7, δ = .9571, ρ = 2, R = 1.0375.

Figure 4: Variation in β

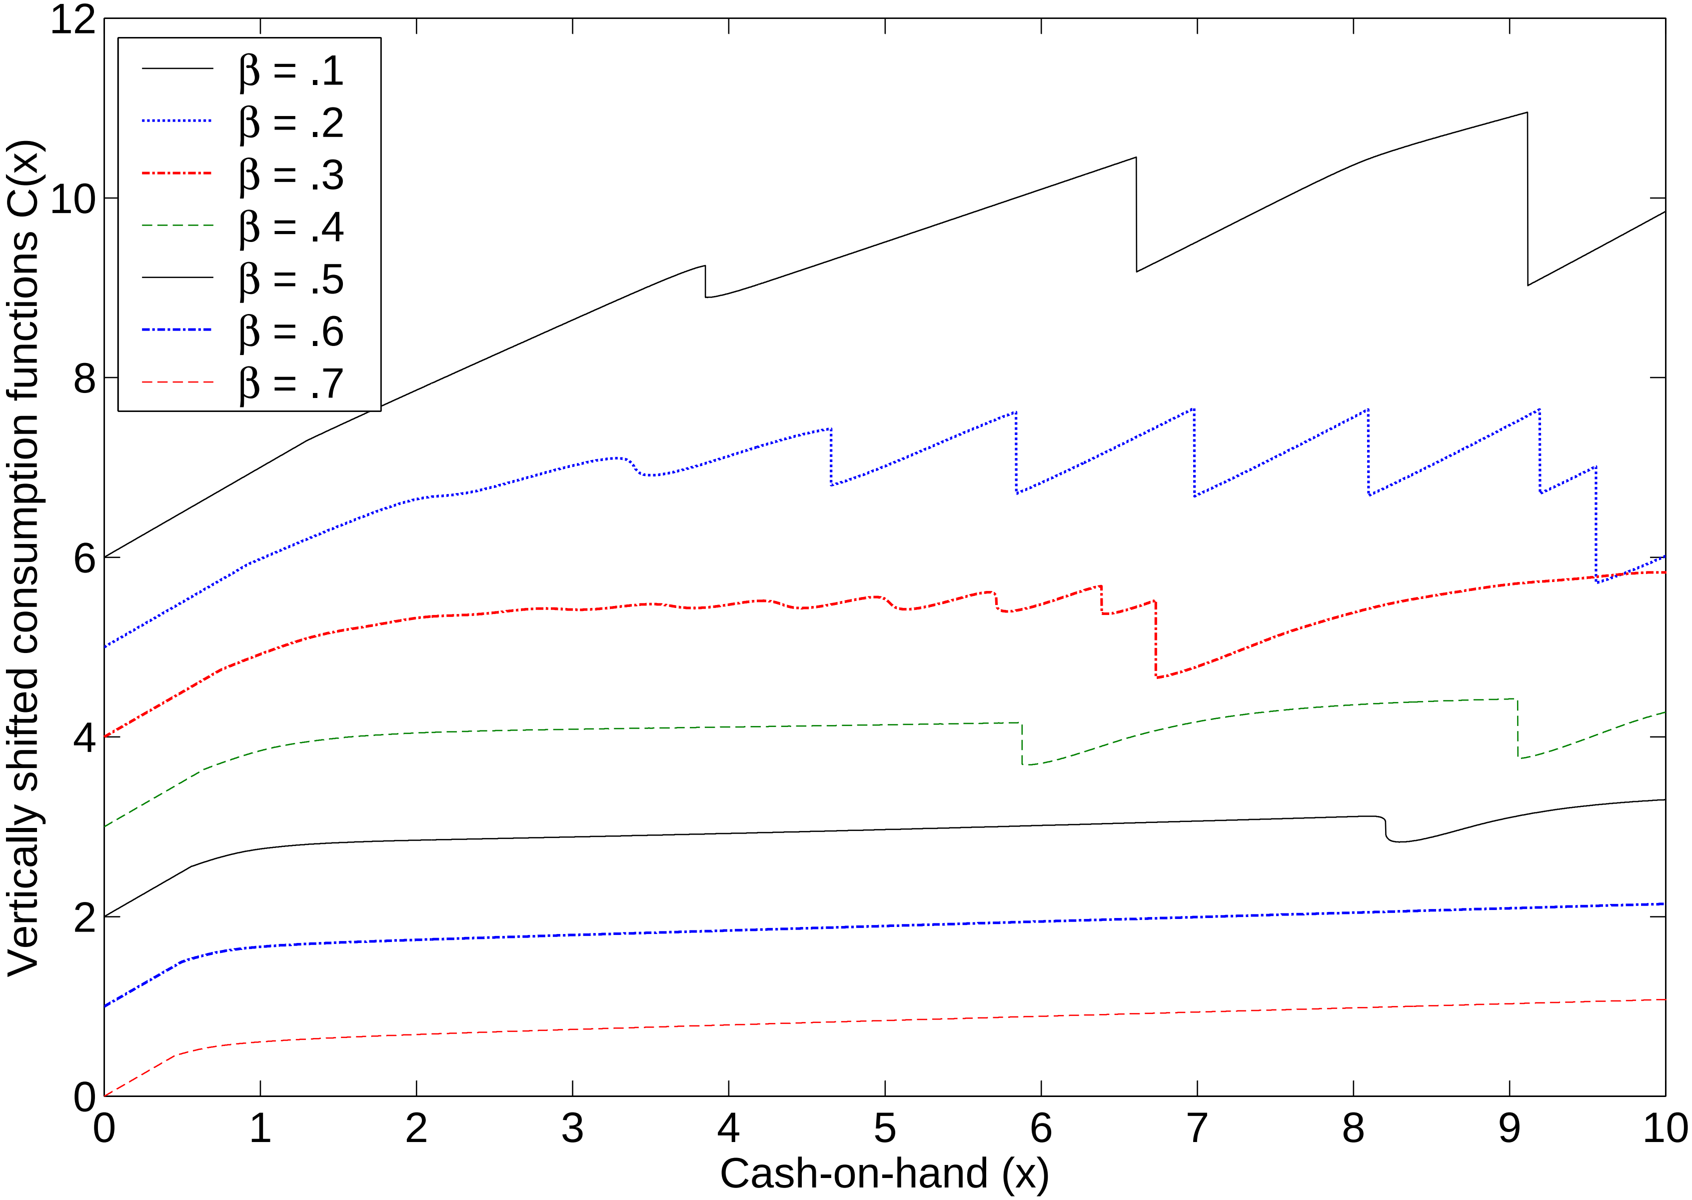

The consumption functions are based on simulations in which $\delta = .9571$, $\rho = 2$, $R = 1.0375$, $a = 5$.

Figure 5: Variation in the coefficient of relative risk aversion (ρ)

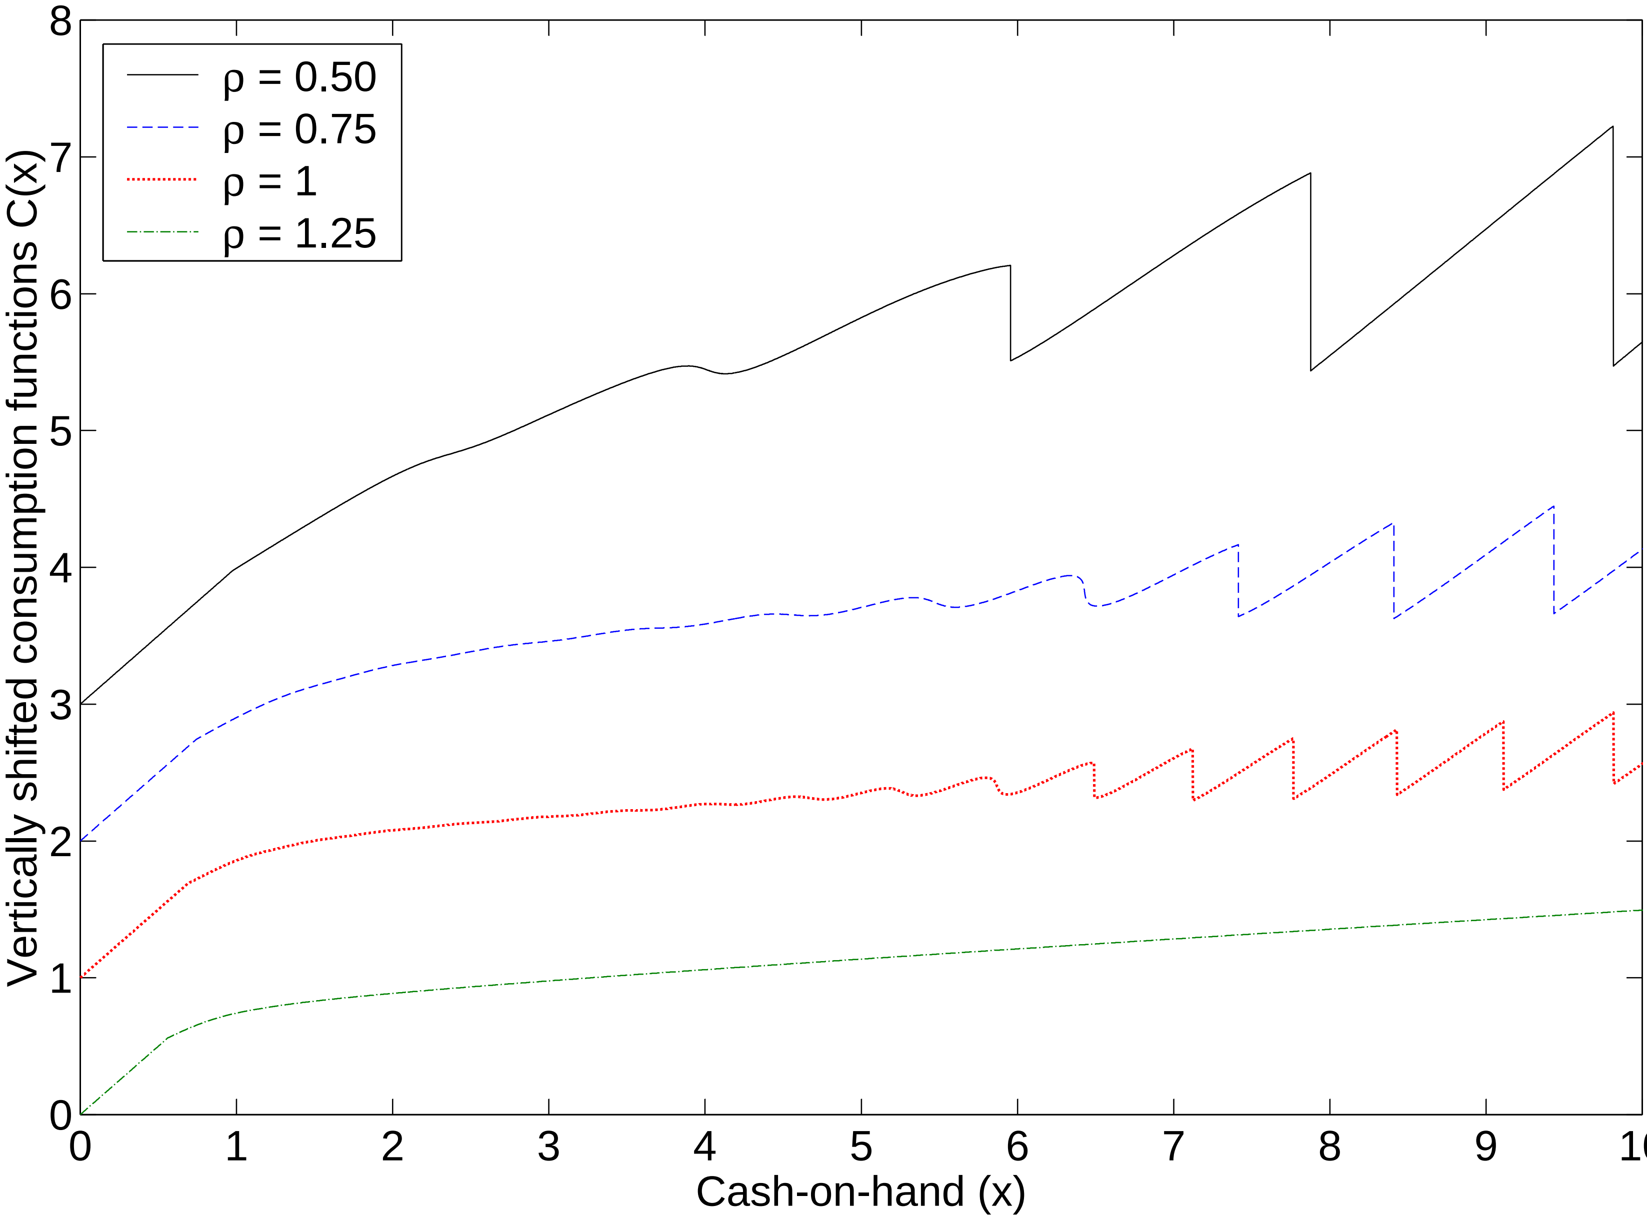

The consumption functions are based on simulations in which ρ = 2, δ = .9571, R = 1.0375, a = 5.

Figure 6: Variation in income uncertainity (a)

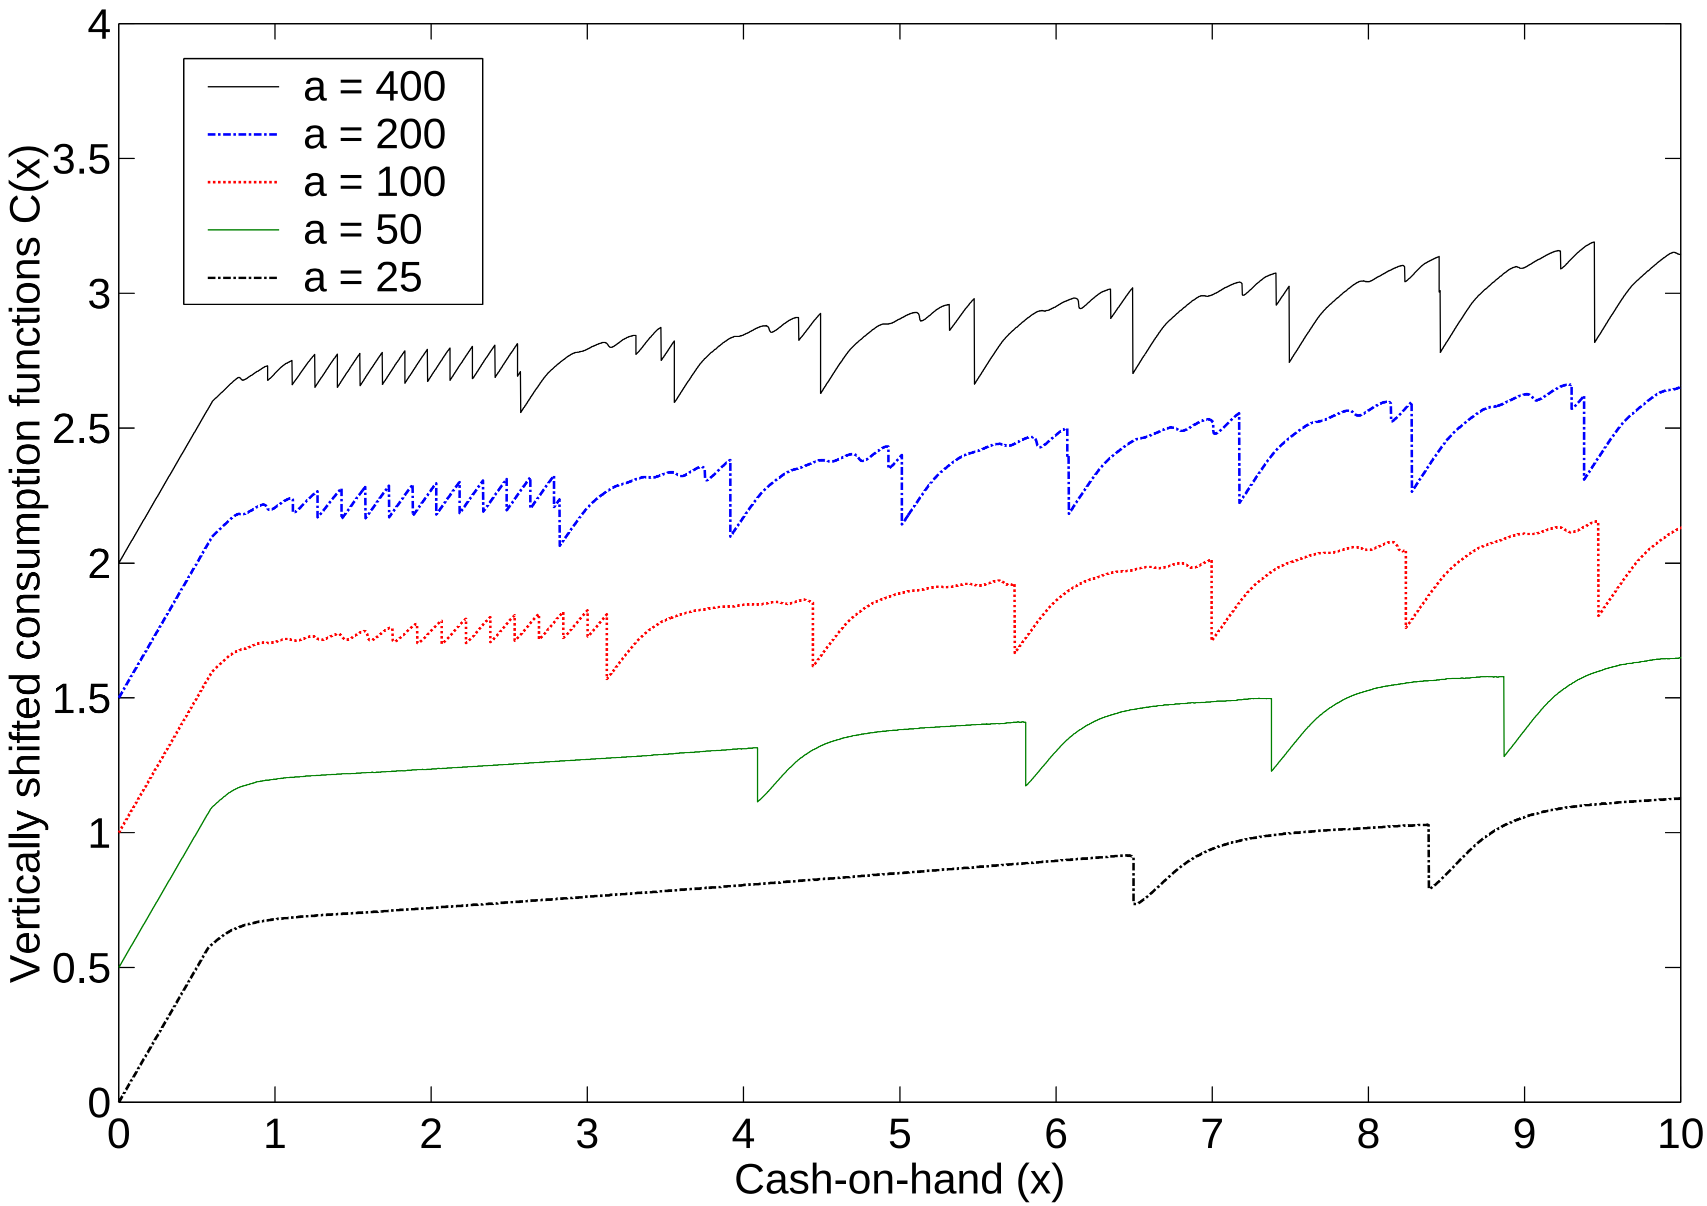

The consumption functions are based on simulations in which $\beta = .7$, $\delta = .9571$, $\rho = 2$, $R = 1.0375$.

Figure 7: Simulated mean consumption profiles
of hyperbolic and exponential households

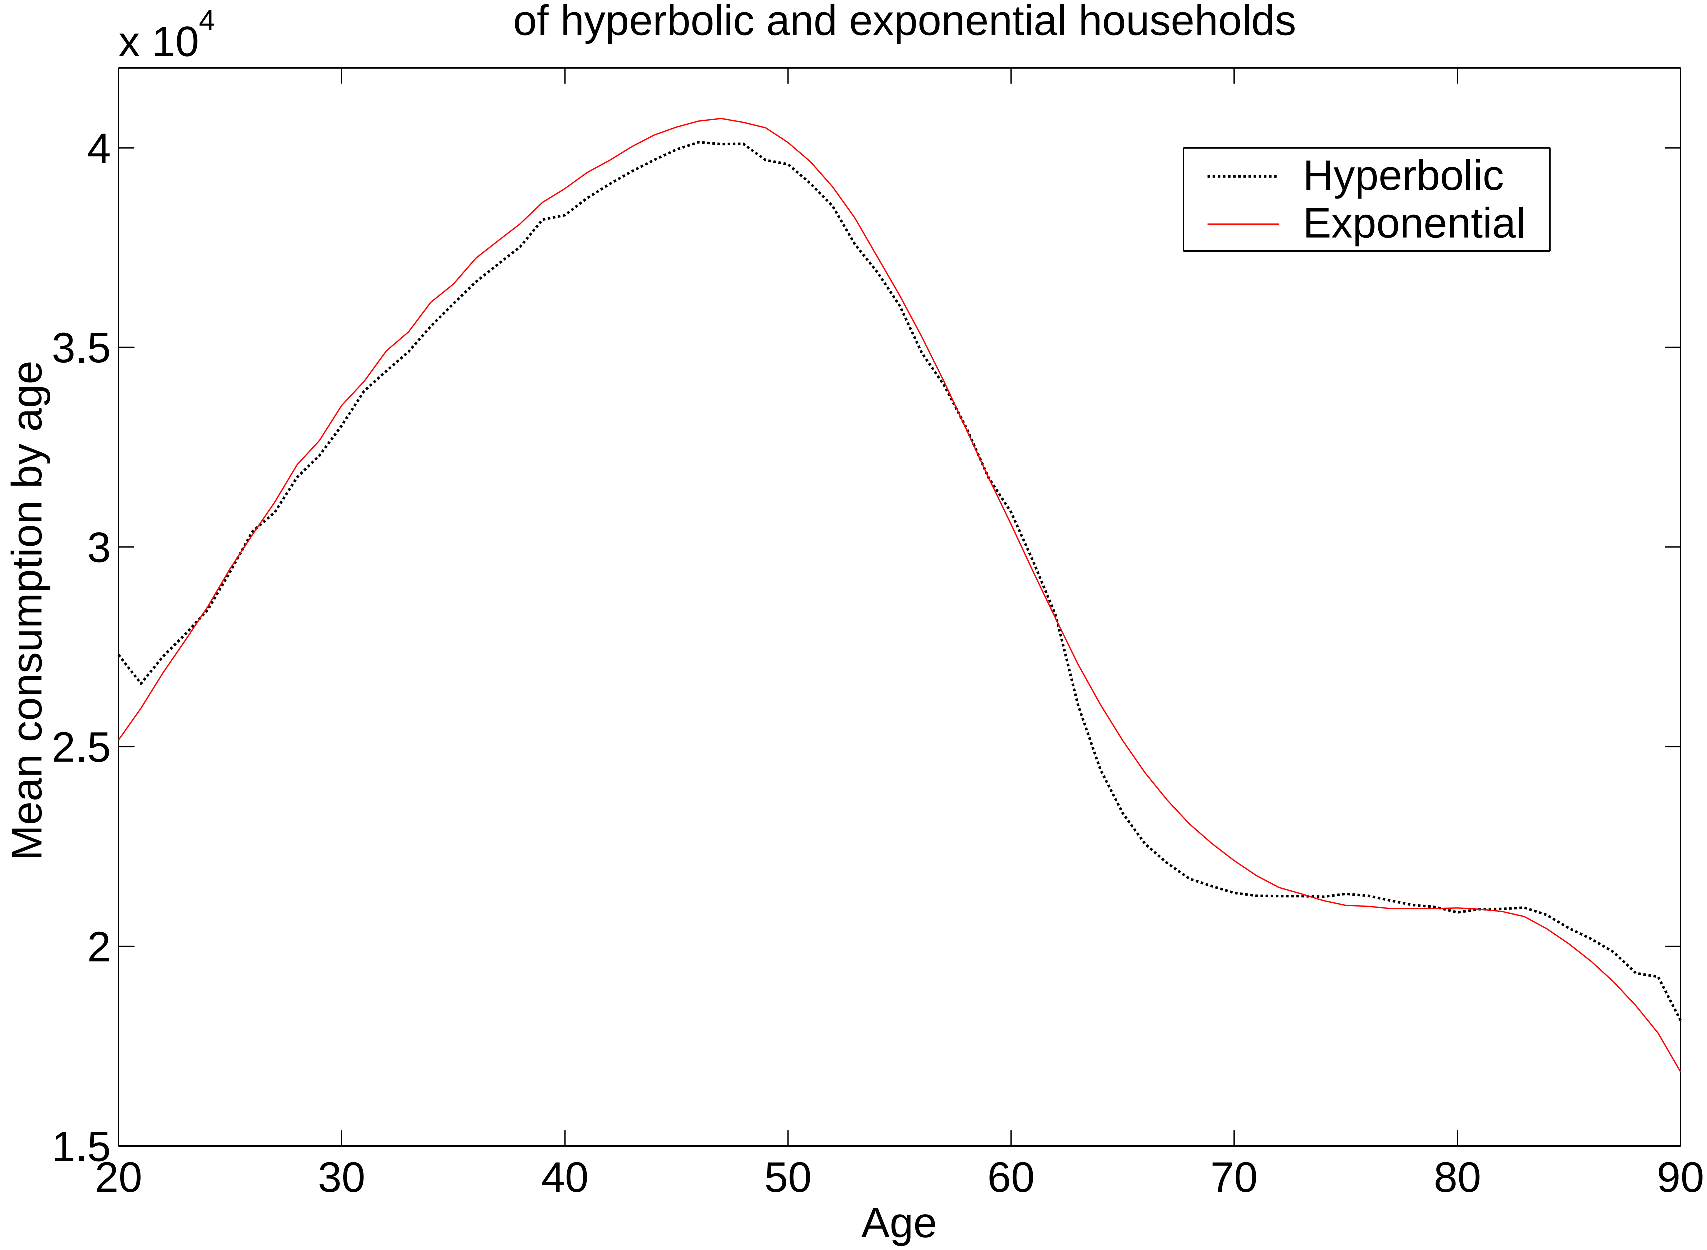

Source: Angeletos et al 2001.

Figure 8: Realization of discount function (α=0.7, γ=0.1)

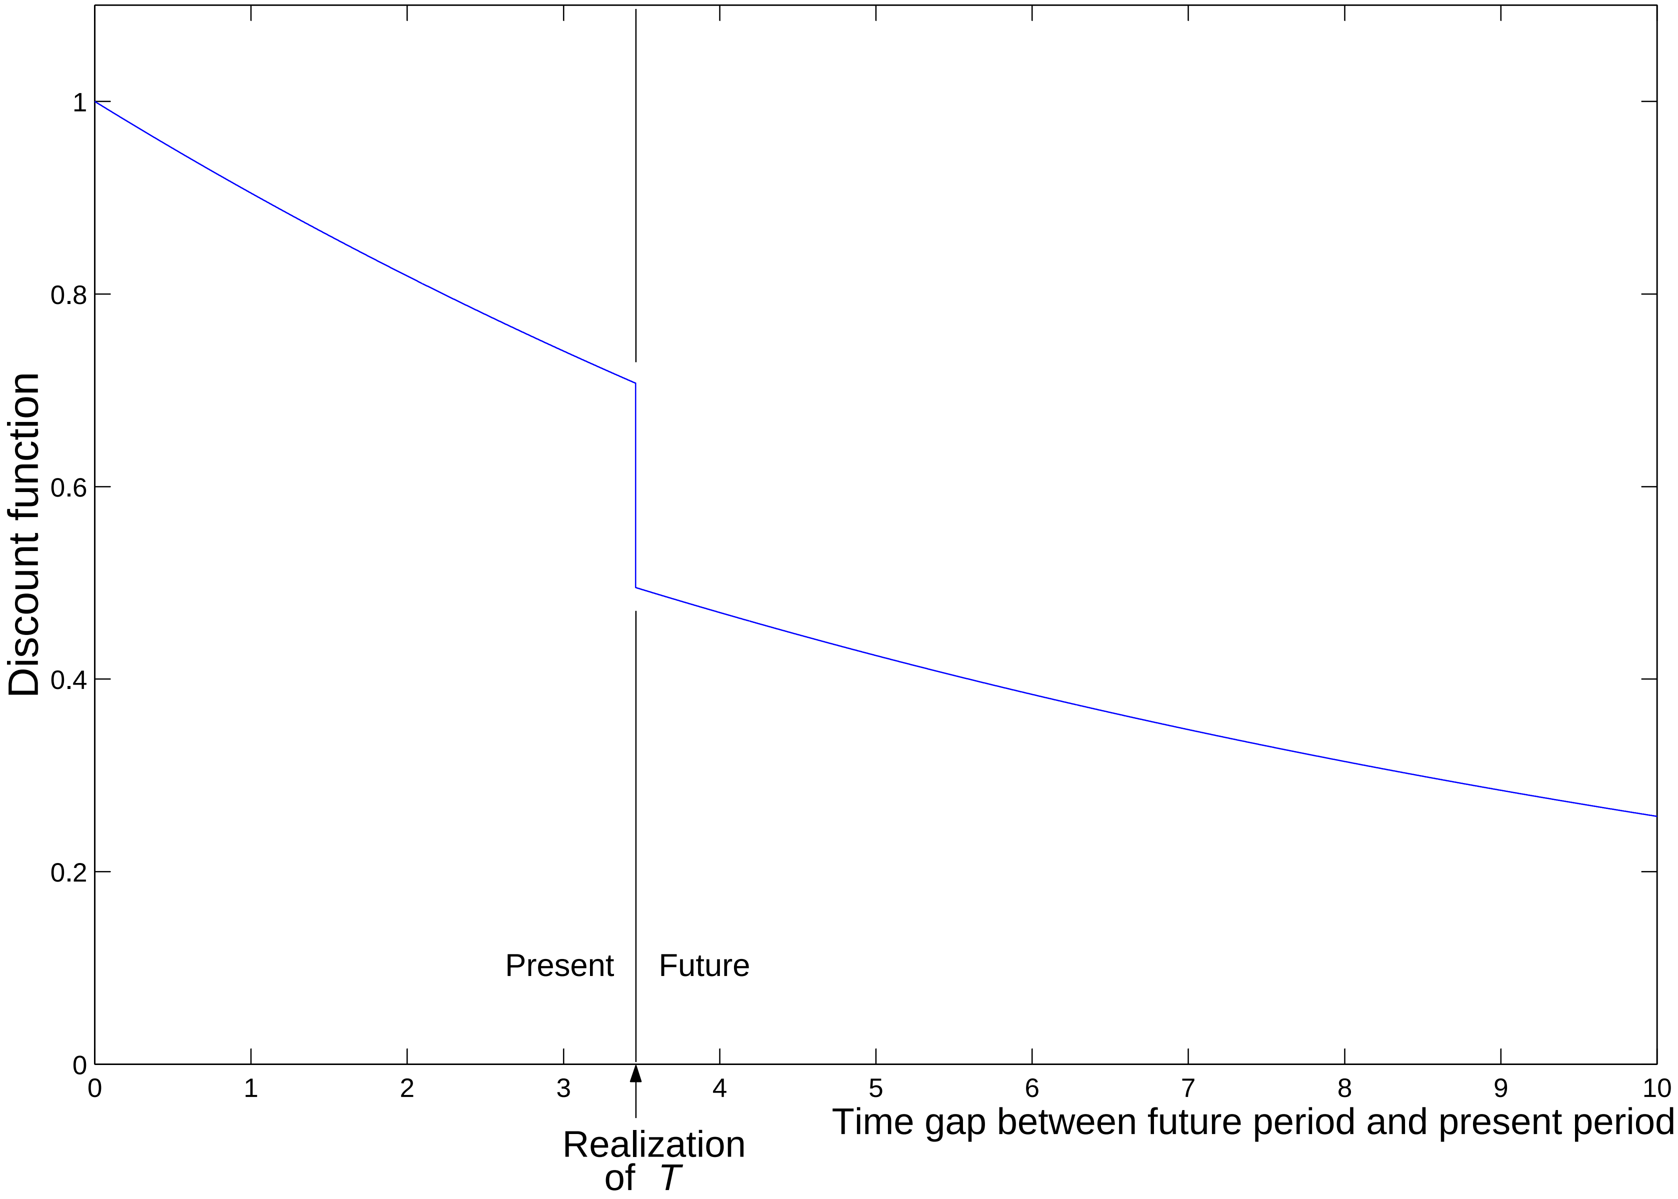

Figure 9: Expected value of discount function for $\lambda \in \{0,0.1,1,10,\infty\}$