

NBER WORKING PAPER SERIES

A DEBT PUZZLE

David Laibson
Andrea Repetto
Jeremy Tobacman

Working Paper 7879
<http://www.nber.org/papers/w7879>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2000

We have benefited from the insights of Daron Acemoglu, Alberto Alesina, Orazio Attanasio, Robert Barro, Martin Browning, John Campbell, Christopher Carroll, José de Gregorio, Eduardo Engel, Benjamin Friedman, Edward Glaeser, Cristóbal Huneus, Greg Mankiw, Julio Rotemberg, Andrei Shleifer, Nicholas Souleles, Richard Zeckhauser, and Steve Zeldes. Marios Angeletos, Eddie Nikolova, and Stephen Weinberg provided excellent research assistance. Laibson acknowledges the financial support of the National Science Foundation (SBR-9510985), the National Institute on Aging (R01-AG-16605), the MacArthur Foundation and the Olin Foundation; and Repetto, from DID-Universidad de Chile, FONDECYT (1990004), Fundación Andes, and an institutional grant to CEA from the Hewlett Foundation. The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

© 2000 by David Laibson, Andrea Repetto, and Jeremy Tobacman. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Debt Puzzle

David Laibson, Andrea Repetto, and Jeremy Tobacman

NBER Working Paper No. 7879

September 2000

JEL No. D91, E21

ABSTRACT

Over 60% of US households with credit cards are currently borrowing -- i.e., paying interest -- on those cards. We attempt to reconcile the high rate of credit card borrowing with observed levels of life cycle wealth accumulation. We simulate a lifecycle model with five properties that create demand for credit card borrowing. First, the calibrated labor income path slopes upward early in life. Second, income has transitory shocks. Third, consumers invest actively in an illiquid asset, which is sufficiently illiquid that it can not be used to smooth transitory income shocks. Fourth, consumers may declare bankruptcy, reducing the effective cost of credit card borrowing. Fifth, households have relatively more dependents early in the life-cycle. Our calibrated model predicts that 20% of the population will borrow on their credit card at any point in time, far less than the observed rate of over 60%. We identify a resolution to this puzzle: hyperbolic time preferences. Simulated hyperbolic consumers borrow actively in the revolving credit card market and accumulate relatively large stocks of illiquid wealth, matching observed data.

David Laibson
Department of Economics
Littauer M-14
Harvard University
Cambridge, MA 02138
and NBER
dlaibson@harvard.edu

Andrea Repetto
Department of Industrial Engineering
Universidad de Chile
República 701
Santiago, Chile
arepetto@dii.uchile.cl

Jeremy Tobacman
Department of Economics
Harvard University
Cambridge, MA 02138
tobacman@fas.harvard.edu

1. INTRODUCTION

At year-end 1998, the Federal Reserve reported that U.S. consumers held approximately \$500 billion in credit card debt. This total only includes debt on which consumers pay interest — not the “float”.¹ Dividing this debt over 102 million U.S. households², yields average debt of approximately \$5,000 per household. Moreover, this average overlooks the fact that many households do not have access to credit. If we restrict attention to the 80% of households with credit cards³, average debt per household rises to over \$6,000. Survey evidence implies that this debt is spread over a large population of debtors. At any given point in time, at least 63% of all households with credit cards are borrowing (i.e., paying interest) on those cards.⁴ These publicly available credit card statistics match the analysis of David Gross and Nicholas Souleles (1999a, 1999b, 2000), who have assembled a proprietary data set that contains a representative sample of several hundred thousand credit card accounts from several different credit card issuers.

This borrowing comes at substantial cost. Despite the rise of teaser interest rates and the high level of competition in the credit card industry, the average debt-weighted credit card interest rate has been approximately 16% in the last five years, implying a real interest rate of 14%.⁵ Within the population of households with a credit card, average interest payments per year exceed \$1,000. This average *includes* households with no interest payments.

This paper attempts to explain credit card borrowing with a standard life-cycle model. Our model has five realistic properties that make credit card borrowing appealing to our simulated consumers.⁶ First, our calibrated labor income path follows a trajectory that is upward sloping

¹The actual total was \$586 billion, but this includes approximately \$80 billion dollars in float. Board of Governors of the Federal Reserve System.

²U.S. Census Bureau, 1998.

³SCF, 1995 cross-section.

⁴The SCF 1995 cross-section implies that 63% of households are borrowing at any point in time, but credit card borrowing in the SCF suffers from dramatic underreporting, perhaps because credit card borrowing is stigmatized.

⁵Board of Governors of the Federal Reserve System. This is a debt-weighted interest rate that includes teaser rates.

⁶Many of these motives for borrowing are elegantly theoretically analyzed by Brito and Hartley (1995). Like them, we build a model that gives consumers a strong motive to borrow on credit cards to smooth intertemporal variation in consumption.

early in life. Second, our income path has transitory income shocks. Third, we introduce an illiquid asset that attracts substantial investment, but is sufficiently illiquid that it can not be used to smooth transitory income shocks. Fourth, we give consumers the opportunity to declare bankruptcy, making credit card borrowing less costly. Fifth, our simulated households have relatively more dependents early in the life-cycle. Despite these institutional features, we are unable to match the actual frequency of credit card borrowing. At any point in time, less than 20% of our simulated consumers hold credit card debt.

The intuition for this result is straightforward. Our simulated model cannot simultaneously match actual levels of credit card borrowing and actual levels of mid-life wealth accumulation. Even if one ignores private and public defined-benefit pension wealth, the median U.S. household enters retirement with assets roughly equal to three times annual pre-retirement labor income. Restricting attention to households with heads between the ages of 50 and 59, actual median net wealth per household is \$149,401.⁷ To match this magnitude of retirement wealth accumulation, we need to calibrate our simulations with low exponential discount rates ($\approx .05$). But, to match actual household credit card borrowing, we need high exponential discount rates ($\approx .18$). Hence, the paper identifies a life-cycle puzzle, which we call the Debt Puzzle. Consumers do not act consistently, acting patiently when it comes to retirement accumulation, and impatiently in the credit card market.

Our simulations show that hyperbolic time preferences may resolve the Debt Puzzle. Intuition for this result comes from the Euler Equation for hyperbolic economies (Harris and Laibson, 2000). This hyperbolic Euler Equation implies that consumers act as if they have endogenous time preferences, acting when liquidity constrained like exponential consumers with a discount *rate* close to .40. However, hyperbolic consumers act patiently when accumulating *illiquid* wealth, because

⁷ June 1999 dollars. This number is the mean of the inflation-adjusted medians from the past four SCF surveys. This net wealth calculation includes all real and financial wealth (e.g., home equity and money market account) as well as all claims on defined contribution pension plans (e.g., 401(k)). The measure does not include Social Security wealth and claims on defined benefit pension plans.

illiquid wealth generates utility flows over long horizons. Hence, our hyperbolic model can explain why the median household borrows aggressively on credit cards, but still manages to accumulate substantial stocks of primarily *illiquid* wealth by retirement.⁸

The rest of the paper formalizes our analysis. In section 2, we present evidence on the proportion of households borrowing on their credit cards. In section 3 we present our benchmark model, which can accommodate either exponential or hyperbolic preferences. In section 4 we provide some analytic approximations that help us evaluate the model's predictions and provide intuition for the simulations that follow. In section 5 we calibrate the model. In section 6 we present our simulation results. In section 7 we present additional simulation results which evaluate the robustness of our conclusions. In section 8 we conclude.

2. CREDIT CARD BORROWING

Eighty percent of households surveyed in the 1995 Survey of Consumer Finances (SCF)⁹ report having a credit card. Of the households with a card, 63% report carrying over a balance the last time that they paid their credit card bill.¹⁰ The average self-reported unpaid balance is \$1,715. The median is \$343. Both this mean and median are calculated on the population of households with credit cards, including households with zero balances. Table 1 reports these statistics for the entire population and for subgroups conditioned on age and educational status.

An average balance of \$1,715 may seem large, but it almost surely reflects dramatic *underreporting* among household respondents to the SCF. The Federal Reserve requires that banks report information on their portfolios of revolving credit loans, excluding loans to businesses. At year-end

⁸We do not explain another credit card puzzle which has recently been documented by Morrison (1998) and Gross and Souleles (1999b). These authors show that a fraction of households (approximately 33%) simultaneously carry credit card debt and hold liquid wealth which exceeds one month of income. Our model predicts that consumers will carry credit card debt and simultaneously hold illiquid wealth, but our model explicitly rules out the phenomenon that Morrison (1998) and Gross and Souleles (2000) document. In addition, the model does not explain why consumers carry credit card debt at high interest rates, rather than switching to low interest rate cards (Ausubel, 1991).

⁹Survey of Consumer Finances, Federal Reserve Board.

¹⁰Specifically, respondents answer the following question: "After the last payment was made on this account, roughly what was the balance still owed on this account?" The answers to this question are used to determine the incidence and level of credit card borrowing.

1995, the total portfolio of loans was \$464 billion. Once the float of approximately \$80 billion is removed, the total falls to approximately \$384 billion. Dividing among the 81 million U.S. households with credit cards, implies average debt per card-holding household of over \$4,500, roughly *three* times as large as the self-reported average from the 1995 SCF. For year-end 1998, the Federal Reserve numbers imply average debt per card-holding household of over \$6,000.

These numbers match values from a proprietary account-level data set assembled by David Gross and Nicholas Souleles (1999a, 1999b, 2000). The Gross-Souleles data contains several hundred thousand representative credit card account statements provided by several large banks. The Federal Reserve figures and Gross-Souleles figures are reported directly by banks and are hence more reliable than household survey evidence which is the raw material for the SCF. Moreover, the Federal Reserve and Gross-Souleles numbers match each other, reinforcing the conclusion that average debt per card-holding household is approximately \$6,000.

Because of the drastic SCF underreporting of the *magnitude* of revolving credit, we focus our analysis on the fraction of households that report carrying over a balance the last time that they paid their credit card bill (e.g., 63% in 1995).¹¹ We believe that this fraction is probably downward biased, but we believe that this bias is relatively minor when compared to the SCF bias for debt magnitudes. The principal goal of this paper will be to determine if standard economic models can match the observed 63% rate of credit card borrowing.

We also analyze the lifecycle pattern of the fraction of households borrowing. Figure 1 plots the estimated age-contingent fraction of married households that carry revolving credit. We plot profiles for household heads in three educational categories: no high school diploma (NHS), high school graduate (HS), college graduate (COLL).¹² To construct these profiles we have eliminated cohort and business cycle effects by including cohort dummies and regional unemployment rates as control variables. The profiles in Figure 1 are estimated using splines with knots at ages 35,

¹¹The 1998 SCF survey has recently become available, and reports credit card borrowing behavior which is little different from the behavior reported in the 1995 survey.

¹²The household's educational status is determined by the educational attainment of the household head.

50, 65, and 80. A full description of the estimation procedure is provided in the appendix. For households in the HS group, we find that 72.5% borrow on their credit cards at age 20. The percent borrowing peaks at age 35 at 81.5%. This rate is relatively flat between ages 35 and 50, and then drops to 51.8% at age 80, and rises to 64.6% by age 90. Households in the NHS group borrow most frequently and COLL households borrow least frequently, but all three groups borrow at roughly similar rates. Indeed, the most striking property of the profiles in Figure 1 is the uniformly high rate of borrowing.

The identification strategy described in the preceding paragraph attributes time trends to age and cohort effects, and assumes that the unemployment rate captures cyclical fluctuations. The estimated age profiles are quite sensitive to these identification assumptions. When we replace the cohort dummies with time dummies, we find that the fraction of households borrowing tends to fall over the lifecycle. This pattern is reflected in Table 1, which reports the raw data from the 1995 SCF. The sensitivity in the estimated lifecycle profiles, leads us to be agnostic about the appropriate identification approach. We believe that cohort effects exist — reflecting habits of behavior and social norms fixed at a relatively young age — and we believe that time effects exist — reflecting society-wide changes in technology and borrowing norms. We can not simultaneously include cohort, age, and time effects in our estimation because these three variables are collinear. John Ameriks and Stephen Zeldes (2000) offer a particularly clear discussion of these identification issues. We will consistently report our cohort-adjusted estimates, since we have greater faith in these results. We urge readers who are skeptical about the identification of age effects to focus on the raw, unadjusted lifecycle averages reported in Table 1. Specifically, 68%, 70%, and 53% of households in the NHS, HS, and COLL groups reported that they were currently borrowing on their credit cards (i.e., paying interest) at the time of the 1995 SCF.

Finally, we are also interested in the relationship between wealth-holding and borrowing. Table 2 reports borrowing frequencies in the 1995 SCF tabulated by age and educational status contingent wealth quartiles. As expected, borrowing declines with wealth (holding age fixed), but this decline

is surprisingly small among the younger cohorts. Consider the 40-49 year-old households in the HS group: 86% of the households in the bottom wealth quartile report that they are borrowing on their credit cards, compared to 50% of the households in the top quartile. By any measure, borrowing is not confined to the bottom half of the wealth distribution.

Using the simulations that follow, we ask whether standard calibrated lifecycle models can match these stylized facts on the frequency of credit card borrowing.

3. MODEL

We model the complex set of constraints and stochastic income events that consumers face. Our framework is based on the simulation literature pioneered by Carroll (1992, 1997), Deaton (1991), and Zeldes (1989) and extended by Hubbard, Skinner, and Zeldes (1994), Engen, Gale, and Scholz (1994), Gourinchas and Parker (1999), and Laibson, Repetto and Tobacman (1998). We discuss the conceptual features of our model in this section and calibrate the model in Section 5.

Our simulations adopt most of the features of previous lifecycle simulation models. We extend the existing literature by enabling households to borrow on credit cards, including a time-varying number of dependent adults and children in the household, allowing the household to invest in a partially illiquid asset, and allowing the household to declare bankruptcy. We divide the presentation of the model into eight domains: 1) demographics, 2) income from transfers and wages, 3) liquid assets and non-collateralized debt, 4) illiquid assets and collateralized debt, 5) budget constraints, 6) bankruptcy, 7) preferences, and 8) equilibrium.

3.1. Demographics. The economy is populated by households who face a time-varying, exogenous hazard rate of survival s_t , where t indexes age. Households live for a maximum of $T + N$ periods, where T and N are exogenous variables that represent respectively the maximum length of pre-retirement life and the maximum length of retirement. If a household is alive at age $20 \leq t \leq T$, then the household is in the workforce. If a household is alive at age $T < t \leq T + N$, then the household is retired. We assume that economic life begins at age 20 and do not model consump-

tion decisions before this date. We assume that household composition — number of adults and non-adults — varies over the life-cycle. Households always contain a household head and a spouse, but the number of adult dependents and non-adult dependents varies.

Our population is divided into three education categories: consumers with no high school diploma, graduates of high school, and graduates of college. We assume education is exogenous, and assign a different working life (T), retirement duration (N), household composition, and labor income process to each education category.

3.2. Income from transfers and wages. Let Y_t represent all after-tax income from transfers and wages. Hence, Y_t includes labor income, inheritances, private defined-benefit pensions, and all government transfers. Since we assume labor is supplied inelastically, Y_t is exogenous. Let $y_t \equiv \ln(Y_t)$. We refer to y_t as “labor income,” to simplify exposition. During working life ($20 \leq t \leq T$):

$$y_t = f^W(t) + u_t + \nu_t^W \quad (1)$$

where $f^W(t)$ is a cubic polynomial in age, u_t is a Markov-process, and ν_t^W is iid and normally distributed, $N(0, \sigma_{\nu, W}^2)$. During retirement ($T < t \leq T + N$):

$$y_t = f^R(t) + \nu_t^R \quad (2)$$

where $f^R(t)$ is linear in age, and ν_t^R is iid and normally distributed, $N(0, \sigma_{\nu, R}^2)$. The parameters of the labor income process vary across education categories.

3.3. Liquid assets and non-collateralized debt. Let X_t represent liquid asset holdings at the beginning of period t , excluding current labor income. So $X_t + Y_t$ represents total liquid asset holdings at the beginning of period t . To model non-collateralized borrowing — i.e., credit card borrowing — we permit X_t to lie below zero, but we introduce a credit limit equal to some fraction

of current (average) income

$$X_t \geq -\lambda \cdot \bar{Y}_t$$

where \bar{Y}_t is average income at age t for the appropriate education group.

3.4. Illiquid assets and collateralized debt. Let Z_t represent illiquid asset holdings at age t . The illiquid asset generates two sources of returns: capital gains and consumption flows. We assume that in all periods Z is bounded below by zero.

$$Z_t \geq 0$$

The household borrows to invest in Z , and we represent such collateralized debt as D , where D is normalized to be positive. Let $I^Z \geq 0$ represent new investments into Z and let $\psi(I^Z)$ represent transaction costs generated by that investment. We assume that each new investment is paid for with a down-payment of exactly $\mu \cdot I^Z$, implying that investment of magnitude I^Z generates new debt equal to $(1 - \mu) \cdot I^Z$.

3.5. Dynamic and static budget constraints. Let I_t^X represent net investment into the liquid asset X , during period t . Recall that I_t^Z represents net investment into the illiquid asset Z , during period t . Let I_t^D represent net repayment of debt, D , during period t . Hence the dynamic budget constraints are given by,

$$X_{t+1} = R^X \cdot (X_t + I_t^X) \tag{3}$$

$$Z_{t+1} = R^Z \cdot (Z_t + I_t^Z) \tag{4}$$

$$D_{t+1} = R^D \cdot (D_t - I_t^D) \tag{5}$$

where R^X , R^Z , and R^D are the real interest rates, respectively, on liquid wealth, illiquid wealth, and debt. We assume that the interest rate on liquid wealth depends on whether the consumer is

borrowing or saving in her liquid accounts. We interpret liquid borrowing as credit card debt.

$$R^X = \begin{cases} R^{CC} & \text{if } X_t + I_t^X < 0 \\ R & \text{if } X_t + I_t^X > 0 \end{cases}$$

Naturally, R^{CC} is the interest rate on credit card debt, and R represents the interest rate on positive stocks of liquid wealth. The within-period budget constraint is:

$$C_t = Y_t - I_t^X - I_t^Z - I_t^D - \psi(I_t^Z)$$

For computational tractability, we have made an additional restriction, which eliminates one choice variable. Specifically, we assume that the debt contract is structured so that a proportion $\Delta = .10$ of D_t is paid off between periods. Hence, we require that debt repayments, I_t^D , be set such that

$$D_{t+1} = (1 - \Delta) \cdot D_t + R^D \cdot (1 - \mu) \cdot I_t^Z \quad (6)$$

Combining Equation 6 with Equation 5 implies that I_t^D is fully determined by the other variables in the model. Hence, the state variables at the beginning of period t are liquid wealth ($X_t + Y_t$), illiquid wealth (Z_t), collateralized debt (D_t), and the value of the Markov process (u_t). The non-redundant choice variables are net investment in liquid wealth (I_t^X) and net investment in illiquid wealth (I_t^Z). Consumption is calculated as a residual.

3.6. Bankruptcy. For some of our simulations we will allow households to declare bankruptcy. If a consumer declares bankruptcy in period t , we assume the following consequences: consumption drops permanently to some level which is proportional to the expected value of permanent income (where permanent income is evaluated at the date at which bankruptcy is declared), X drops permanently to zero, Z drops permanently to $\min\{Z^{\text{Bankruptcy}}, Z_t - D_t\}$, and D drops permanently to zero.

3.7. Preferences. We use standard preferences in our benchmark model. The instantaneous utility function is characterized by constant relative risk aversion and the discount function is exponential (δ^t).

We also analyze an alternative model that has hyperbolic discount functions, but is otherwise identical to the benchmark model. Hyperbolic time preferences imply that from today's perspective discount rates are higher in the short-run than in the long-run. Experimental data support this intuition. When researchers use subject choices to estimate the shape of the discount function, the estimates consistently approximate generalized hyperbolas: events τ periods away are discounted with factor $(1 + \alpha\tau)^{-\gamma/\alpha}$, with $\alpha, \gamma > 0$.¹³

Figure 2 graphs the standard exponential discount function (assuming $\delta = .944$), the generalized hyperbolic discount function (assuming $\alpha = 4$, and $\gamma = 1$), and the quasi-hyperbolic discount function, which is an analytically convenient approximation of the generalized hyperbola. The quasi-hyperbolic function is a discrete time function with values $\{1, \beta \cdot \delta, \beta \cdot \delta^2, \beta \cdot \delta^3, \dots\}$. Figure 2 plots the case of $\beta = .7$ and $\delta = .956$.¹⁴ When $0 < \beta < 1$ the quasi-hyperbolic discount structure mimics the qualitative property of the hyperbolic discount function, while maintaining most of the analytical tractability of the exponential discount function.

Quasi-hyperbolic and hyperbolic preferences induce dynamically inconsistent preferences. Consider the discrete-time quasi-hyperbolic function. Note that the discount factor between adjacent periods n and $n + 1$ represents the weight placed on utils at time $n + 1$ relative to the weight placed on utils at time n . 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

¹³See Loewenstein and Prelec (1992) for an axiomatic derivation of this discount function. See Chung and Herrnstein (1961) for the first use of the hyperbolic discount function. Laboratory experiments have been done with a wide range of real rewards, including money, durable goods, fruit juice, sweets, video rentals, relief from noxious noise, and access to video games. See Ainslie (1992) for a partial review of this literature. See Mulligan (1997) for a critique.

¹⁴This discount function was first analyzed by Phelps and Pollak (1968). However, their use of this structure was motivated in a different way. Their application is one of imperfect intergenerational altruism, and the discount factors apply to non-overlapping generations of a dynasty. Following Laibson (1997a) we apply this discount function to an intra-personal problem. Like Laibson (1997a) we assume the horizon is finite. Phelps and Pollak assume an infinite horizon which admits a continuum of equilibria (Laibson 1994). The particular parameter values used in this example correspond to the calibration used in this paper for households with a high school educated head.

less than one, this implies a short-term discount rate that is greater than the long-term discount rate. 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 at time $t + 1$.

Because of the 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. Recently, hyperbolic discount functions have been used to explain a wide range of anomalous economic choices, including procrastination, contract design, drug addiction, retirement timing, and undersaving.¹⁵

To analyze the decisions of an agent with dynamically inconsistent preferences, we must specify the preferences of all of the temporally distinct selves. We index these selves by their lifecycle position, $t \in \{20, 21, \dots, T + N - 1, T + N\}$. Self t has instantaneous payoff function

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

and continuation payoffs given by:

$$\beta \sum_{i=1}^{T+N-t} \delta^i \left(\prod_{j=1}^{i-1} s_{t+j} \right) [s_{t+i} \cdot u(C_{t+i}, Z_{t+i}, n_{t+i}) + (1 - s_{t+i}) \cdot B(X_{t+i}, Z_{t+i}, D_{t+i})]. \quad (7)$$

Note that n_t is the effective household size,

$$n_t = ([\# \text{ adults}_t] + \kappa[\# \text{ of children}_t]),$$

ρ is the coefficient of relative risk aversion, γZ_t represents the consumption flow generated by

¹⁵See Akerlof (1991), Barro (1997), Diamond and Koszegi (1998), Laibson (1994,1996,1997a), O'Donoghue and Rabin (1997a, 1997b).

Z_t , s_{t+1} is the probability of surviving to age $t + 1$ conditional on being alive at age t , and $B(\cdot)$ represents the payoff in the death state, which incorporates a bequest motive. The first expression in the bracketed term in Equation 7 represents utility flows that arise in period $t + i$ if the household survives to age $t + i$. The second expression in the bracketed term represents termination payoffs in period $t + i$ which arise if the household dies between period $t + i - 1$ and $t + i$.

3.8. Equilibrium: When $\beta < 1$ the household has dynamically inconsistent preferences, and hence the consumption problem can not be treated as a straightforward dynamic optimization problem. Late selves will not implement the policies that are optimal from the perspective of early selves. Following the work of Strotz (1956) we model consumption choices as an intra-personal game. Selves $\{20, 21, \dots, T + N - 1, T + N\}$ are the players in this game. Taking the strategies of other selves as given, self t picks a strategy for time t that is optimal from its perspective. This strategy is a mapping from the (Markov) state variables, $\{t, X + Y, Z, D, u\}$, to the non-redundant choice variables $\{I^X, I^Z\}$. An equilibrium is a fixed point in the strategy space, such that all strategies are optimal given the strategies of the other players. We solve for the equilibrium strategies using a numerically implemented backwards induction algorithm.

Our choice of the quasi-hyperbolic discount function simplifies the induction algorithm. Let $V_{t,t+1}(X_{t+1} + Y_{t+1}, Z_{t+1}, D_{t+1}, u_{t+1})$ represent the time $t + 1$ continuation payoff function of self t . Then the objective function of self t is:

$$u(C_t, Z_t, n_t) + \beta \delta E_t V_{t,t+1}(\Lambda_{t+1}) \quad (8)$$

where Λ_{t+1} represents the vector of state variables: $\{X_{t+1} + Y_{t+1}, Z_{t+1}, D_{t+1}, u_{t+1}\}$. Self t chooses I_t^X and I_t^Z , which jointly define C_t , maximizing this expression. The sequence of continuation payoff functions is defined recursively:

$$V_{t-1,t}(\Lambda_t) = s_t[u(C_t, Z_t, n_t) + \delta E_t V_{t,t+1}(\Lambda_{t+1})] + (1 - s_t)E_t B(\Lambda_t) \quad (9)$$

where s_t is the probability of surviving to age t conditional on being alive at age $t - 1$ and C_t is the consumption chosen by self t . The induction continues in this way. Note that dynamic inconsistency in preferences is reflected in the fact that a β factor appears in Equation 8 — reflecting self t 's discount factor between periods t and $t + 1$ — but does not appear in Equation 9, since self $t - 1$ does not use the β factor to discount between periods t and $t + 1$.

Equations 8 and 9 jointly define a functional equation which is not a contraction mapping. Hence, the standard dynamic programming results do not apply to this problem. Specifically, V does not inherit concavity from u , the objective function is not single-peaked, and the policy functions are in general discontinuous and non-monotonic.¹⁶ We have adopted a numerically efficient solution algorithm — based on local grid searches — which iterates our functional equation in the presence of these non-standard properties.

Our equilibrium definition has a major shortcoming: we adopt the standard economic assumption of unlimited problem-solving sophistication. The consumers in our model solve perfectly a complex backwards induction problem when making their consumption and asset allocation choices. We are not satisfied with this extreme assumption, but view it as a reasonable starting point for analysis.¹⁷

4. ANALYTIC APPROXIMATIONS

4.1. Exponential case: $\beta = 1$. Consider a stripped-down version of our benchmark model. Specifically, set $\beta = 1$, assume that labor income is iid, eliminate the illiquid asset, and eliminate time-varying mortality and household size effects. It is possible to use the standard Euler Equation to impute a value for the discount rate, $-\ln(\delta)$. The exponential Euler Equation is:

$$u'(C_t) = E_t R \delta u'(C_{t+1})$$

¹⁶See Laibson (1997b).

¹⁷Another reasonable starting point is the model of “naïf” behavior first proposed by Robert Strotz (1956) and more recently studied by Akerlof (1991), and O’Donoghue and Rabin (1997a, 1997b). These authors propose that decision makers with dynamically inconsistent preferences make current choices under the false belief that later selves will act in the interests of the current self.

The second order approximation of this equation is:

$$E_t \Delta \ln(C_{t+1}) = \frac{1}{\rho} (r + \ln(\delta)) + \frac{\rho}{2} V_t [\Delta \ln(C_{t+1})],$$

which can be rearranged to yield

$$\begin{aligned} \text{discount rate} &= -\ln(\delta) \\ &= -\rho E_t \Delta \ln(C_{t+1}) + r + \frac{\rho^2}{2} V_t [\Delta \ln(C_{t+1})] \end{aligned} \tag{10}$$

To impute the value of the discount rate, we need to evaluate $E_t \Delta \ln(C_{t+1})$, r , ρ , and $V_t [\Delta \ln(C_{t+1})]$.

We will do this for a typical household.

Consider only U.S. households which have access to a line of revolving credit and have a 45-year-old head. Order these households by the expected one-year rate of consumption growth. Survey data implies that the median household should expect flat consumption between ages 45 and 46.¹⁸ It is reasonable to assume that this median household holds credit card debt, as credit card borrowing peaks in frequency and magnitude for households with 45-year-old heads. Over three-quarters of households with 45-year-old heads and credit cards have credit card debt.¹⁹ Hence, for our analysis, the appropriate real interest rate is the real credit card borrowing rate, $r = r^{cc} \approx .14$.²⁰ We will consider a range of values for ρ . Finally, the conditional variance of consumption growth can be represented as a proportion of the conditional variance of income growth. When income is a random walk, the conditional variance of consumption growth is approximately equal to the conditional variance of income growth. We assume that the conditional variance of consumption growth is half of the conditional variance of income growth, implying that the conditional variance of consumption growth is .025. This value is consistent with our calibrated simulation results. The lack of consumption smoothing is also consistent with the fact that the typical household is

¹⁸E.g., Gourinchas and Parker (1999).

¹⁹SCF, 1995 cross-section.

²⁰See section 5 for details on the calibration of interest rates.

borrowing in the credit card market, a portfolio decision that suggests low levels of liquid wealth accumulation and hence necessarily imperfect consumption smoothing.²¹

We are now in a position to evaluate $-\ln(\delta)$. Figure 3 plots $-\ln(\delta)$ on the y-axis, against ρ on the x-axis. The solid line reflects the assumptions described in the previous paragraph. The line is monotonically increasing with a minimum of 0.14 (at $\rho = 0$). For reasons that we describe below, this value turns out to be anomalously high. In anticipation of this problem, we have plotted a second line in Figure 3 (the dashed line), which reflects more aggressive assumptions that lower our envelope of discount rates. Specifically, we raise $E_t \Delta \ln(C_{t+1})$, lower r , and lower $V_t[\Delta \ln(C_{t+1})]$ in an effort to make the discount rate, $-\ln(\delta)$, as low as possible. For this second line, we set $E_t \Delta \ln(C_{t+1}) = .01$, $r = .13$, and $V_t[\Delta \ln(C_{t+1})] = .015$. We believe that these assumptions are inappropriate, but they serve to identify a lower bound for the discount rate envelope. This second plotted line begins at a discount rate of 0.13 (at $\rho = 0$), and then falls slightly to a minimum of 0.127 (at $\rho = .67$), before rising monotonically thereafter. Hence, whatever assumptions we make, we are unable to generate implied discount rates below thirteen percent.

Our first set of analytical approximations turn out to match closely the results of our numerical simulations of exponential consumers (assuming $\rho = 2$). When we choose the exponential discount rate so our simulations match a credit card borrowing frequency of .70 (the empirical frequency for the HS households), we end up selecting an exponential discount rate of .179, close to the value predicted by Equation 10:

$$\begin{aligned} -\rho E_t \Delta \ln(C_{t+1}) + r + \frac{\rho^2}{2} V_t[\Delta \ln(C_{t+1})] &= -2 \cdot 0 + .14 + \frac{2^2}{2} \cdot (.025) \\ &= .19 \end{aligned}$$

Such high discount rates are problematic. Observed household wealth accumulation profiles

²¹It is not optimal for consumers with exponential or hyperbolic time preferences to simultaneously hold credit card debt (at a real interest rate of 14%) and hold positive liquid assets (at a real interest of approximately 4%). See Morrison (1998) and Gross and Souleles (1999) for evidence that some consumers do engage in such transparently irrational behavior.

can only be explained with much *lower* discount rates. For example, the median U.S. household accumulates total pre-retirement wealth equal to 3.34 times after-tax income.²² To calibrate lifetime consumption and wealth profiles, most authors have used discount rates that lie below 0.10. Engen, Gale, and Scholz (1994) calibrate their model with a discount rate of 4% ($\rho = 3$). Hubbard, Skinner and Zeldes (1995) calibrate their simulations with a discount rate of 3% ($\rho = 3$). Gourinchas and Parker estimate a discount rate of 4% ($\rho = .5$). Laibson, Repetto, and Tobacman (1998) estimate two central discount rates: 4% ($\rho = 1$) and 6% ($\rho = 3$). Engen, Gale, and Uccello (1999) calibrate their model with a discount rate of 0% and 3% ($\rho = 3$).²³

Hence, these observations suggest a puzzle. Consumers act impatiently in the credit market but act patiently when accumulating for retirement. We call this the Debt Puzzle. In Sections 6 and 7 we extend this argument by analyzing the fully general model which incorporates many of the rich institutional details that complicate real-world decisions.

4.2. Hyperbolic case: $\beta < 1$. The discussion above only applies to exponential consumers. As Harris and Laibson (2000) have shown, making the discount function hyperbolic generates an important modification of the Euler Equation. To derive this Hyperbolic Euler Equation²⁴, recall that the current self chooses C according to:

$$C^* = \operatorname{argmax}_C u(C) + \beta \delta E_{\Omega} [V(R \cdot (X + Y - C) + Y_{+1})].$$

where $V(\cdot)$ is the continuation payoff function, and for simplicity the horizon is infinite, implying that $V(\cdot)$ does not depend on time. Recall from above that $V(\cdot)$ has the recursive property,

$$V(X + Y) \equiv u(C^*) + \delta E_{\Omega} [V(R \cdot (X + Y - C^*) + Y_{+1})]$$

²²SCF, 1995 survey. Our definition of wealth includes all assets except claims on defined contribution pension plans. For a detailed list of the assets that we include, see the section on model calibration.

²³All of these papers assume real interest rates (on positive savings) of 1-5 percent. Naturally, substantially higher interest rates would justify substantially higher discount rates, but historical data pin the interest rate down.

²⁴An heuristic derivation follows, which assumes differentiability of the value and consumption functions. For a fully general derivation, see Harris and Laibson (2000).

where Ω represents the current information set. Finally, represent the welfare of the current self as:

$$W(X + Y) \equiv u(C^*) + \beta\delta E_{\Omega} [V(R \cdot (X + Y - C^*) + Y_{+1})]$$

Then the envelope theorem (ET) implies:

$$W'(X + Y) = u'(C^*) \quad (\text{ET})$$

Moreover, the first-order-condition (FOC) in the current self's problem implies

$$u'(C^*) = R\beta\delta E_{\Omega} [V'(R \cdot (X + Y - C^*) + Y_{+1})]. \quad (\text{FOC})$$

Finally, $V(\cdot)$ and $W(\cdot)$ are linked by the identity

$$\beta V(X + Y) = W(X + Y) - (1 - \beta)u(C^*). \quad (\text{By def.})$$

Using these relationships it follows that

$$\begin{aligned} u'(C_t) &= R\beta\delta E_t [V'(X_{t+1} + Y_{t+1})] && \text{by the FOC} \\ &= R\delta E_t \left[W'(X_{t+1} + Y_{t+1}) - (1 - \beta)u'(C_{t+1}) \frac{\partial C_{t+1}}{\partial X_{t+1}} \right] && \text{by definition} \\ &= R\delta E_t \left[u'(C_{t+1}) - (1 - \beta)u'(C_{t+1}) \frac{\partial C_{t+1}}{\partial X_{t+1}} \right] && \text{by the ET} \end{aligned}$$

Note that the partial derivative of consumption with respect to cash-on-hand can be equivalently represented as either $\frac{\partial C_{t+1}}{\partial X_{t+1}}$ or $\frac{\partial C_{t+1}}{\partial (X_{t+1} + Y_{t+1})}$. Rearranging the last equation yields:

$$u'(C_t) = E_t R \left[\beta\delta \left(\frac{\partial C_{t+1}}{\partial X_{t+1}} \right) + \delta \left(1 - \frac{\partial C_{t+1}}{\partial X_{t+1}} \right) \right] u'(C_{t+1})$$

This equation is identical to the exponential case, except that the exponential discount factor, δ , is replaced by the endogenous effective discount factor

$$\left[\beta\delta \left(\frac{\partial C_{t+1}}{\partial X_{t+1}} \right) + \delta \left(1 - \frac{\partial C_{t+1}}{\partial X_{t+1}} \right) \right].$$

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 $\frac{\partial C_{t+1}}{\partial X_{t+1}}$, the marginal propensity to consume, and $\left(1 - \frac{\partial C_{t+1}}{\partial X_{t+1}} \right)$. The effective discount factor is stochastic and endogenous to the model.

When consumers are liquidity constrained, the marginal propensity to consume, $\frac{\partial C_{t+1}}{\partial X_{t+1}}$, is approximately equal to unity. In this case, the effective discount factor is approximately equal to $\beta\delta$. Assuming that $\beta = 0.7$ and $\delta = 0.95$ (a conservative calibration of the quasi-hyperbolic discount function when each period is a year)²⁵ the effective discount rate will approximately equal $-\ln(0.7 \times 0.95) = 0.41$.

Hyperbolic consumers have an incentive to keep themselves liquidity constrained (Laibson, 1997a). By storing wealth in illiquid form, hyperbolic consumers prevent themselves from overspending in the future. Early selves intentionally try to constrain the consumption of future selves. This has the effect of raising the future marginal propensity to consume out of the (constrained) stock of liquid wealth. The high marginal propensity to consume generates high effective discount rates ($\approx .41$), explaining why hyperbolics are frequently willing to borrow on credit cards.

Hyperbolics recognize that illiquid wealth will be spent much more slowly than liquid wealth. Illiquid wealth — e.g., housing — generates marginal utility flows for many periods in the future. The consumer discounts utility flows τ periods away with factor $\beta\delta^\tau$. When discounting consumption increments over long-horizons, a hyperbolic consumer uses an effective discount rate of

$$\lim_{\tau \rightarrow \infty} \left[\ln(\beta\delta^\tau)^{\frac{1}{\tau}} \right] = \lim_{\tau \rightarrow \infty} \left[-\frac{1}{\tau} \ln(\beta) - \ln(\delta) \right] = -\ln(\delta).$$

Hence, illiquid wealth accumulation is primarily driven by δ , not β , implying that the consumer

²⁵See Laibson (1997a).

accumulates illiquid wealth as if she had a discount rate of $-\ln(\delta) = .05$.

With the potential for effective discount rates of 41% per year, the model predicts widespread borrowing on credit cards at 15% – 20% annual interest rates. However, the hyperbolic model simultaneously predicts that most consumers will accumulate large stocks of illiquid wealth, basing accumulation decisions on a relatively low discount rate of .05.

5. CALIBRATION

In this section we discuss our calibration decisions for both our benchmark models and for variations to that benchmark. Most of our calibration decisions are standard for the consumption literature except for the second to last subsection which discusses calibration of preferences.

5.1. Demographics. We use education group population weights 0.25, 0.50, and 0.25 (no-high school, high school, college) which roughly match the actual proportions in the PSID.

Consumers live for a maximum of 90 years ($T + N$), though they do not enter the work force or make economically meaningful decisions in our model until age 20. The conditional hazard rates of survival are taken from the life tables of the U.S. National Center for Health Statistics (1993). These tables report the probability of living to age $t + 1$, conditional on having lived to age t . This one-year survival probability is close to one through age 70, dropping to 96.3% by age 80, and 67.6% by age 89.

Following Engen, Gale and Scholz (1994), we use the survival rates for a single individual even though the “consumers” in our model are actually multi-person households. Conceptually our model assumes that surviving households always have two non-dependent adults (e.g., a head of household and a spouse) and an exogenously age-varying number of dependents — including adult dependents and non-adult dependents.²⁶

²⁶Our “single individual” mortality assumption engenders two subtle biases that go in opposite directions. First, our approach may yield *too much* simulated retirement saving because our model implicitly rules out insurance effects that arise when spouses have independent mortality outcomes (in real life an N-person marriage creates a partial annuity which becomes perfect as N goes to infinity). Second, our mortality assumption yields a bias which implies *too little* simulated retirement saving, because widows and widowers have expenses that fall by less than 50% when

To calibrate the age-varying number of dependents, we use the Panel Study of Income Dynamics (PSID), and condition on households with a head and a spouse. The measure of children in the household includes all children between 0 and 17; it does not include the head or spouse even if either or both of them is younger than 18. It includes all children whether or not they are actually children of the head or spouse. The number of dependent adults represents the actual number of members 18 years of age and older, excluding head and spouse.

To construct effective household size, we smooth the observed profiles of dependent children and dependent adults. These smooth profiles are computed, for each educational category, as follows. First, we dropped households with heads younger than 20 or older than 90. Second, we restricted the sample to households with a head and a spouse. Finally, we estimated the following nonlinear regression model, using nonlinear least squares

$$x_{it} = \beta_0 \exp(\beta_1 \cdot \text{age}_{it} - \beta_2 \cdot \text{age}_{it}^2) + \varepsilon_{it}. \quad (11)$$

Note that x_{it} represents either the number of children or the number of dependent adults in household i at date t , and the errors ε_{it} represent i.i.d. noise. We picked this particular function because it captures the shape of the observed profiles, and because it predicts a positive number of children and dependent adults for every age.

In Table 3 we report the estimated coefficients and their standard errors. In Figure 4 we plot the smoothed profiles for the number of children and the number of dependent adults of the three education groups. To construct these profiles we set ε_{it} equal to zero. The life cycle pattern of the variables is considerably different across education groups. The profiles are lower and slightly steeper for college educated individuals, and the peak in the number of children occurs two to three years later.

Following Blundell et al (1994), we define effective household size as the number of adults plus

their spouses die.

0.4 times the number of children.²⁷ We assume that the total number of adults is equal to two (head and spouse) plus the number of predicted dependent adults. As expected, our predicted measure of effective household size exhibits a hump shape pattern. Furthermore, like empirical profiles of consumption (Gourinchas and Parker, 1999), family size peaks in the mid to late 40's.

5.2. Income from transfers and wages. We define income as after-tax non-asset income. Our definition includes labor income, bequests, lump-sum windfalls, and government transfers such as AFDC, SSI, workers' compensation and unemployment insurance. This definition is broader than the one used by Engen and Gale (1993) — who use only labor earnings — and the one used by Hubbard et al (1994 and 1995) — who only add unemployment insurance payments to labor income.

The sample of households is taken from the Panel Study of Income Dynamics (PSID). We use the family files for the interview years between 1983 and 1990, since these are the only PSID sample years that include bequests and other lump-sum windfalls, as well as federal taxes. We exclude all households whose head is younger than 20 years of age, that report annual income less than \$1000 (in 1990 dollars, deflated by the CPI for urban consumers), or that have any crucial variable missing.²⁸ To calculate pre-retirement income we follow the approach of Bernheim et al (1997), who define a year as pre-retirement if anyone in the household worked 1500 hours or more in that year or in any subsequent year. A household is retired if no member works more than 500 hours in the current year or in any year in the future.

We estimate the regression equation:

$$y_{it} = HS_{it} + \text{polynomial}(\text{age}_{it}) + TE_t + CE_i + \xi_{it} \quad (12)$$

²⁷There exist other adult equivalence scales. For instance, Attanasio (1998) uses the official OECD scale, which gives weight 1 to the first adult, 0.67 to the following adults, and 0.43 to each child. Using empirical data, Deaton and Muellbauer (1986) estimate that children cost their parents about 30-40 percent of what they spend on themselves.

²⁸We believe that reported income of less than \$1000 is more likely to reflect a coding or reporting error than to reflect a true report. Recall that our income definition includes all government transfers.

by weighted least squares, using the PSID population weights. This equation is estimated twice, once for households in the labor force and once for retired households. Income of household i in period t is determined by a household size effect (HS_{it}), a polynomial in age, a time effect (TE_t), and a cohort effect (CE_i). The household size effect integrates the effects of three variables: the number of heads in the household (head only or head and spouse), the number of children, and the number of dependent adults. We specify the age polynomial as third degree for our pre-retirement regression and linear for our post-retirement regression. Following Gourinchas and Parker (1997), and to circumvent the problem that age, time, and birth year are perfectly correlated, we assume that the time effect is related to the business cycle and that it can be proxied by the unemployment rate. We use the unemployment rate in the household's state of residence, taken from the Bureau of Labor Statistics. Our cohort effects control for birth year to account for permanent differences in productivity across cohorts.²⁹ We use five-year age-cohorts, the oldest born in 1910-14 and the youngest born in 1970-74. Table 4 reports the income regressions for each education group.

We calculate f^W and f^R — the polynomials in the model of the previous section — by setting the cohort and unemployment effects equal to the sample means, setting the number of heads equal to two, and the number of dependents — children and adults — equal to the age varying smoothed profiles estimated in the previous subsection. This allows us to recover variation in expected income over the lifecycle for a household that has a typical lifecycle evolution in household size, experiences no business cycle effects, and has a typical cohort effect. Figure 5 plots the exponentiated values of f^W and f^R for the three education categories.

To study the stochastic component of pre-retirement non-asset household income we exploit the panel dimension of the PSID. We model the unexplained part of measured non-asset income (ξ_{it}) as the sum of an individual fixed effect, an AR(1) process, and a purely transitory shock:

$$\xi_{it} = \vartheta_i + u_{it} + v_{it}^W = \vartheta_i + \alpha u_{it-1} + \epsilon_{it} + v_{it}^W$$

²⁹See Ameriks and Zeldes (2000) and Attanasio and Weber (1993) for a discussion of cohort effects.

The individual fixed effect is included to account for permanent differences in income that are not completely captured by the educational categories, particularly differences in human capital and earning ability.

Let $\sigma_{v,W}^2$ be the variance of the transitory shock v^W , and σ_ϵ^2 , the variance of ϵ . Also, let $C_k \equiv E(\Delta\xi_t \Delta\xi_{t-k})$ represent the theoretical autocovariances of $\Delta\xi$. Then

$$\begin{aligned} C_0 &= \frac{2\sigma_\epsilon^2}{1+\alpha} + 2\sigma_v^2 \\ C_1 &= \frac{-\sigma_\epsilon^2 \cdot (1-\alpha)}{1+\alpha} - \sigma_v^2 \\ &\vdots \\ C_d &= \frac{-\alpha^{d-1} \sigma_\epsilon^2 \cdot (1-\alpha)}{1+\alpha} \end{aligned}$$

We estimate the parameters σ_ϵ^2 , $\sigma_{v,W}^2$ and α using weighted GMM by minimizing the distance between the theoretical and the empirical first seven autocovariances. The estimated parameters are presented in Table 5. These parameter values are almost identical to the values reported by Hubbard et al (1994), who estimate an identical after-tax income process.

The transitory noise in retirement income is inferred by estimating

$$\xi_{it} = \vartheta_i + v_{it}^R$$

on retired households, where ϑ_i is a household fixed effect, and v_{it}^R has variance $\sigma_{v,R}^2$.

In the numerical simulations, we set the individual effect equal to zero, and we represent u_t (an AR1 process) with a two-state Markov process. The latter is done to save computational time. The Markov process is symmetric, taking on two states $\{-\theta, +\theta\}$, with symmetric transition probability p . To make this two-state Markov process match the variance and autocovariance of u_t , we set $\theta = \sqrt{\frac{\sigma_\epsilon^2}{1-\alpha^2}}$ and $p = \frac{\alpha+1}{2}$.

To calculate the typical retirement age by education group we look at households that expe-

rienced a transition into retirement over the observed period (using the Bernheim et al (1997) definition of retirement). We find that the mean age at which households without a high school diploma (with a high school diploma, with a college degree) begin retirement is 61 (63, 65).

5.3. Liquid assets and non-collateralized debt. We calibrate the credit limit $\lambda \cdot \bar{Y}_t$ using the 1995 SCF. Specifically, for each education group we identify the households with credit cards and calculate for each age t

$$\lambda_t = \sum_h \frac{\theta_{ht}(\text{credit limit})_{ht}}{\bar{Y}_t}.$$

where h indexes households, and θ_{ht} is the population weight of household h who is t years old. The age profiles of λ_t are virtually flat, while the levels are quite similar across education groups, with an overall weighted average of almost 24%. We selected $\lambda = .30$, a number larger than the observed mean, to take into account the fact that the SCF reports the credit limit associated with Visa, Mastercard, Discover, and Optima cards only, and does not include credit limit information on store and other charge cards. It is worth noting that the four listed cards accounted, on average, for about 80% of total credit card debt according to the 1995 SCF.

5.4. Illiquid assets and collateralized debt. For our benchmark simulation we assume an extreme form of transaction costs:

$$\psi(I^Z) = \begin{cases} 0 & \text{if } I^Z > 0 \\ \infty & \text{if } I^Z \leq 0 \end{cases}.$$

In other words, purchases of the illiquid asset generate no transaction costs, but sales are infinitely costly. Alternatively, one could simply assume that sales costs are sufficiently large to make sales of the illiquid asset unappealing. By making the illiquid asset extremely illiquid we heighten the need for credit card borrowing, since the illiquid asset cannot be used to buffer transitory income shocks. Our simulation code is sufficiently flexible to consider other less extreme assumptions, which we do in Section 7 on robustness.

In our benchmark simulations we allow no collateralized debt, and therefore set the downpayment fraction $\mu = 1$. We explore the parameterization $\mu = .10$ in our robustness checks.

5.5. Dynamic and static budget constraints. We set the value of the after-tax real interest rate on liquid savings equal to 3.75 percentage points. This assumes that liquid assets are invested in a diversified portfolio of stocks and bonds ($\frac{2}{3}$ stocks and $\frac{1}{3}$ bonds), and that the effective tax rate on real returns is 25%.

In our benchmark simulation, we do not allow the household to declare bankruptcy. In this case, we set the real interest rate on credit card loans to 10.75 percentage points, three percentage points *below* the mean debt-weighted real interest rate measured by the Federal Reserve Board. We do this to bias up our credit card borrowing, and to implicitly capture the effect of bankruptcy. Actual annual bankruptcy rates of one percent per year, imply that the effective interest rate is one percentage point below the observed interest rate.

Later in the paper we explicitly model bankruptcy, allowing consumers to escape their uncollateralized debt obligations with some penalty. When bankruptcy is explicitly modelled, we set the real interest rate on credit card loans to 13.75%, equal to the rate measured by the Federal Reserve Board.

We set the real return on illiquid assets to 0, but assume that illiquid assets generate a consumption flow equal to 5.00 percent of the value of the illiquid asset (i.e., $\gamma = .05$). Hence, illiquid assets have the same pre-tax gross return as liquid assets, but illiquid assets generate consumption flows that are by-and-large not taxed (e.g., housing). Hence, the “after-tax” return on illiquid assets is considerably higher than the after-tax return on other assets. We explore an even higher rate of return in our robustness checks.

Finally, we set the after-tax real interest rate on collateralized debt to 5.00 percentage points. Hence, the pre-tax real interest rate is 6.67 percentage points, assuming that interest payments on collateralized debt are tax deductible (e.g., housing).

5.6. Bankruptcy. In our benchmark simulations we do not allow bankruptcy and instead lower the credit card interest rate three percentage points to reflect the probability that the debt will not all be repaid. In Section 7, we consider a simulation that explicitly allows households to enter bankruptcy. We describe the assumptions for this case here.

If bankruptcy is declared in period t , we assume the following consequences: consumption drops permanently to a proportion $\alpha_{\text{Bankruptcy}}$ of the expected value of permanent income (where permanent income is evaluated at the date at which bankruptcy is declared), X drops permanently to zero, Z drops permanently to $\min\{Z^{\text{Bankruptcy}}, Z_t - D_t\}$, and D_t drops permanently to zero. We set $Z^{\text{Bankruptcy}} = \$100,000$ to reflect state laws that allow bankrupt households to retain partial or full ownership of their primary residence.³⁰ We found that setting $\alpha_{\text{Bankruptcy}} = 1$ generates simulated bankruptcy rates that approximately match observed bankruptcy rates (on average .7% of our simulated households enter bankruptcy each year). This match arises because consumers value the flexibility of choosing the timing of consumption. Recall that early-life child rearing and high rates of time preference make it optimal to consume more when young. In our simulations, declaring bankruptcy forces the households to give up this flexibility (i.e., they are forced to consume the annuity value of their human and physical wealth). Naturally, this annuity assumption is unrealistic. It simply serves as a calibrated “punishment” for declaring bankruptcy. We know that our assumed punishment has realistic utility consequences because of the associated frequency with which bankruptcy is endogenously chosen by our simulated consumers. In other words, the utility consequence is roughly realistic since our simulated consumers choose bankruptcy as often as real-world consumers.

5.7. Preferences.

Coefficient of relative risk aversion: ρ . We adopt a utility function with a constant coefficient of relative risk aversion. In our benchmark calibration we set the coefficient of relative

³⁰See Repetto (1998).

risk aversion, ρ , equal to two, a value which lies in the middle of the range of values that are commonly used in the consumption literature (i.e., $\rho \in [.5, 5]$).³¹

Time preferences: β . In Section 6 we simulate exponential economies and hyperbolic economies. In these simulations we assume that the economy is either populated exclusively by exponential households (i.e., $\beta = 1$) or exclusively by hyperbolic households, which we model by setting $\beta = .7$. Most of the experimental evidence suggests that the one-year discount rate is at least 30%-40%.³² We experiment with β values below .7 in Section 7.

Bequests: We parameterize the bequest payoff function as

$$B(X_t, Z_t, D_t) = (R - 1) \cdot \max\{0, X_t + \frac{2}{3}(Z_t - D_t)\} \cdot \frac{\alpha^{\text{Bequest}} \cdot u_1(\bar{y}, 0, \bar{n})}{1 - \delta} \quad (13)$$

where \bar{n} is average effective household size over the life-cycle, and \bar{y} is average labor income over the life-cycle (calculated separately for each educational group). We arbitrarily set $\alpha^{\text{Bequest}} = 1$, but test other values in our section on robustness. We multiply bequeathed illiquid wealth by two-thirds to capture the idea that much of that wealth can only be liquidated with substantial transactions costs (e.g., furniture, automobiles, and to a more limited extent housing). Note that $B(X_t, Z_t, D_t)$ is weakly increasing in X_t and $Z_t - D_t$.

To motivate our specific functional form assumptions, recall that

$$u(C, Z, n) = n \cdot \frac{\left(\frac{C+\gamma Z}{n}\right)^{1-\rho} - 1}{1 - \rho},$$

implying that,

$$u_1(\bar{y}, 0, \bar{n}) = \left(\frac{\bar{y}}{\bar{n}}\right)^{-\rho}.$$

³¹See Laibson, Repetto, and Tobacman (1998) for a detailed discussion of calibration of ρ , and an argument that ρ is closer to .5 than to 5.

³²See Ainslie (1992) for a review.

Equation 13 follows from assuming that the bequest recipient's total consumption is approximately equal to \bar{y} , the bequest recipient's effective household size is \bar{n} , and the bequest recipient consumes bequeathed wealth as an annuity.

Time preferences : δ . Having fixed all of the other parameters, we are left with three free parameters in our hyperbolic simulations — $\delta_{\text{hyperbolic}}^{\text{NHS}}$, $\delta_{\text{hyperbolic}}^{\text{HS}}$, $\delta_{\text{hyperbolic}}^{\text{COLL}}$ — and three free parameters in our exponential simulations — $\delta_{\text{exponential}}^{\text{NHS}}$, $\delta_{\text{exponential}}^{\text{HS}}$, $\delta_{\text{exponential}}^{\text{COLL}}$. The superscripts NHS, HS, and COLL represent our three educational groups. In our simulations we pick the various δ values so that our simulations replicate the actual level of pre-retirement wealth holdings. Specifically, we pick δ such that the *simulated* median ratio of total wealth to income for individuals between ages 50 and 59 matches the *actual* median in the data (SCF). When we construct total wealth from the SCF, we include liquid assets (checking accounts, savings accounts, money market accounts, call accounts, CD's, bonds, stocks, mutual funds, cash, less credit card debt), and illiquid assets (IRA's, defined contribution pension plans, life insurance, trusts, annuities, vehicles, home equity, real estate, business equity, jewelry/furniture/antiques, home durables, less education loans). We do not include defined benefit pension wealth, such as claims on the Social Security System. When we measure total wealth in our simulations, we add: $X + Z + \frac{Y}{24}$, where X represents liquid assets (excluding current labor income), Z represents illiquid assets, and Y represents annual after-tax labor income. The last term is included to reflect average cash-inventories used for (continuous) consumption. 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 $\frac{Y}{24}$.

The SCF data is taken from the 1983, 1989, 1992, and 1995 surveys. We match the mean of the medians across those four years of surveys. The empirical medians and their means are reported in Table 6. The (mean) median ratio of net wealth to income for individuals between ages 50 and 59 is 2.5 for households whose head has no high school degree, 3.2 for households whose head's highest educational attainment is a high school education, and 4.3 for households whose head has a college

degree.

The discount rates ($1-\delta$) that replicate these wealth to income ratios are reported in Table 7. Three properties stand out. First, the discount rates generally fall with educational attainment. Since the shape of the labor income profile is roughly similar across educational groups, a relatively high discount rate is needed to replicate the relatively low wealth to income ratio of the least educated households. Second, the discount rates for the hyperbolic consumers are lower than the discount rates for the exponential consumers. Since hyperbolic consumers have two sources of discounting — β and δ — the hyperbolic δ 's must be higher than 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 we “equalize” the underlying willingness to save between the exponential and hyperbolic consumers. Third, all of our calibrated long-term discount rates are sensible, falling between .04 and .09. Note that these discount rates do not include mortality effects which add roughly another .01 to the discount rates discussed above.

5.8. Equilibrium. To numerically solve for our backwards induction solution, we have developed an algorithm based on local grid searches that iterates our functional operators (Equations 8 and 9).³³

6. SIMULATION RESULTS

We begin by presenting our results on the exponential households ($\beta = 1$). Throughout this section, we focus on households in the HS group and on aggregates, since results for households in the NHS and COLL groups are qualitatively similar to the results for the HS group.

6.1. Exponential Simulation Results. Figure 6 plots the average consumption profile for households whose heads have a high school education (HS group). The average labor income profile is plotted for comparison. Low frequency consumption-income comovement is evident in

³³A description of the algorithm is available from the authors.

this figure. Figure 7 plots the realized consumption and income path for a single household in the HS group. This Figure demonstrates both high and low frequency consumption-income comovement.

Figure 8 plots the simulated median and mean amount of credit card borrowing, along with the age-dependent credit limit.³⁴ Since X_t is liquid wealth, credit card borrowing is defined as $\max\{0, -X_t\}$. Hence, when X_t is negative, credit card borrowing equals $-X_t$. When X_t is positive, credit card borrowing equals zero.

Figure 9 plots the mean level of illiquid wealth (Z_t), liquid wealth (X_t), and illiquid plus liquid wealth (Z_t+X_t) for our simulated households in the HS group. Liquid wealth incorporates the effects of credit card borrowing, and borrowing is sufficiently large to make average liquid wealth negative before age 25. The precautionary motive generates buffer stock saving which eventually overtakes credit card borrowing in the 30's, pushing average liquid wealth above zero. In mid-life the buffer stock vanishes because the consumer can now buffer transitory income shocks by cutting back her substantial investment flow into illiquid assets.

To evaluate the accuracy of the model, we focus on the proportion of households who are borrowing on their credit cards. We focus on this variable since there does not exist a reliable public-use data source for household level credit card borrowing magnitudes (see Section 2). Figure 10 plots the simulated proportion of households in the HS group who are borrowing on their credit card. On average 20.5% of the simulated exponential households borrow on their credit card. This proportion is well below 70%, the observed fraction of HS households that report that they are credit card borrowers in the SCF (1995 cross-section. See Table 1). Naturally, one would like to control for cohort effects when making such comparisons between the simulated data and the SCF data. Hence, in Figure 10 we plot both the simulated profile and the cohort-adjusted empirical estimate of the fraction of households in the HS group who are borrowing on their credit cards.³⁵

³⁴Simulated levels (like means and medians) are not directly comparable to the aggregate averages, since the simulated values correspond to a single representative cohort. At any given point in time many different cohorts of consumers coexist, each with a specific proportional shift of the expected income profile. Consequently, each cohort has its own proportionally shifted profile for level variables like consumption, credit card debt and wealth.

³⁵See Figure 1 and the appendix.

The estimated empirical profile lies everywhere above the simulated profile.

Similar results arise for the simulated exponential households in the NHS and COLL groups. In the NHS (COLL) group, the borrowing frequency is 22% (28%). These results are particularly puzzling because they reverse the empirical ranking of the educational groups. In the 1995 SCF, the reported frequency is 68% for the NHS group and 53% for the COLL group.

We calculate population aggregates by taking weighted averages across our three groups of households: NHS, HS, COLL. These groups respectively represent roughly 25%, 50%, and 25% of the household population, but since we are focusing on households with credit cards, we assume that the percentages are actually 22.6%, 48.3%, and 29.2%. These proportions are consistent with the 1995 SCF which reports that 72% of households in the NHS group have credit cards, 77% of households in the HS group have credit cards, and 93% of households in the COLL group have credit cards.³⁶ Figure 11 plots the simulated aggregate median and mean amount of credit card borrowing, along with the mean of the credit limit. Figure 12, plots the aggregate percentage of households that are borrowing on their credit cards. It is immediately apparent that these aggregate plots do not match the observed data. In the simulated aggregate data, 23% of households borrow on their credit cards at any point in time. In the observed data at least 63% of all households with credit cards borrow on their credit cards at any point in time. Figure 12 also plots the aggregate empirically estimated fraction of households who are borrowing on their credit cards, removing cohort effects. This estimated empirical profile lies uniformly above the simulated profile.

We also compare the simulated borrowing frequencies across wealth categories. Table 8 reports the simulated borrowing frequencies across age-contingent wealth quartiles (for both exponential and hyperbolic simulations). These values can be compared to the empirical frequencies in Table 2. It is immediately apparent that the exponential simulations do not match the empirical data. Two tensions arise. First, as already pointed out, the exponential borrowing frequencies are too low. Second, the exponential borrowing frequencies drop off too sharply as wealth rises. For example

³⁶E.g., $\frac{(.25)(.72)}{(.25)(.72)+(.50)(.77)+(.25)(.93)} = .226$

for the 40-49 year olds in the HS group, the quartile-based simulated borrowing frequencies take values: 54%, 20%, 9% and 2%. By contrast, the empirical frequencies take values: 86%, 79%, 74%, 50%. Similar contrasts arise for other age categories and educational groups. Simulated exponential borrowing is too infrequent, and this empirical failure is particularly dramatic among the high wealth households. Contrary to the data, high wealth simulated exponential households practically do not borrow at all. This mismatch is most striking at the youngest ages. Simulated exponential consumers between ages 20-29 and 30-39 in their respective top wealth quartiles borrow at an average frequency below 1%. This contrasts with empirical borrowing frequencies of 68% (ages 20-29, top wealth quartile) and 59% (ages 30-39, top wealth quartile).

6.2. Hyperbolic Simulation Results. We now turn to our benchmark hyperbolic simulations. Figure 13 compares the exponential and hyperbolic consumption paths for the HS group. These paths are almost identical, except for a small hyperbolic consumption boom at the beginning of life, and the relatively steeper decline in hyperbolic consumption during the retirement period.³⁷

Figure 14 compares total wealth accumulation of exponential and hyperbolic consumers. Two properties distinguish the hyperbolic households. First, the hyperbolic households borrow more when young, depressing total wealth and even driving it below zero for a substantial portion of the lifecycle. Second, hyperbolic households hold more illiquid wealth, which cannot be dissaved and hence elevates total wealth when old. These comparisons are shown in Figure 15, which plots illiquid wealth for exponentials and hyperbolics, and Figure 16, which plots liquid wealth for exponentials and hyperbolics. Similar exponential-hyperbolic contrasts arise for the simulated households in the NHS and COLL groups.

The relative scarcity of liquid wealth is associated with high levels of credit card borrowing for simulated hyperbolic households. Households in the NHS, HS, and COLL groups borrow at

³⁷Like the exponential simulations, the hyperbolic simulations also exhibit low and high frequency comovement between consumption and income (see Laibson et al, 1998). Hence the hyperbolic model is consistent with the empirical regularities documented by Carroll (1992, 1997a), Gourinchas and Parker (1999) and others.

respective frequencies of 60%, 58%, and 49%. These percentages are similar to those in the SCF data: 68%, 70%, 53%. In both the simulations and the data, the NHS and HS frequencies are approximately equal, and the COLL frequency is noticeably lower.

We now turn to comparisons of population aggregates (aggregating across the three educational groups). Figure 17 plots the median amount of simulated credit card borrowing for exponential and hyperbolics, along with the simulated age-dependent credit limit. Figure 18 plots the proportion of households who are borrowing on their credit cards. For simulated hyperbolic households the aggregate borrowing frequency is 55%, compared to 23% of the simulated exponential households. Recall that at least 63% of households are currently borrowing on their credit cards. Figure 18 also plots the estimated cohort-adjusted lifecycle profile of borrowing frequencies. This profile lies everywhere above the simulated exponential profile, but either intersects or nearly intersects the hyperbolic profile at ages 21, 66, and 90.

Finally, we compare the simulated borrowing frequencies across wealth categories. Reconsider Tables 2 and 8, which report the empirical and simulated borrowing frequencies across age-contingent wealth quartiles. Like the exponential simulations, the hyperbolic simulations also predict too little borrowing of high wealth households. For example for the 40-49 year olds in the HS group, the quartile-based hyperbolic borrowing frequencies take values: 84%, 60%, 42% and 24%. The exponential borrowing frequencies take values: 54%, 20%, 9% and 2%. The SCF empirical frequencies take values: 86%, 79%, 74%, 50%. Hence, both the hyperbolic and exponential borrowing frequencies drop off too quickly as wealth rises. Similar patterns arise for other age categories and educational groups.

In summary, the hyperbolic model seems broadly consistent with the empirical data. Hyperbolic consumers borrow at approximately the right average frequency. Moreover, hyperbolics with college educations borrow less frequently than hyperbolics without a college degree. The principal failure of the hyperbolic model is the prediction that high wealth households will borrow at relatively low frequencies. High wealth households in the SCF borrow too frequently to match the predictions of

either the hyperbolic or the exponential model.

7. ROBUSTNESS CHECKS

The results reported in the previous section are robust to substantial variation in all of the calibration assumptions. In every variant that we have considered (a fraction of which are reported here), exponential households continue to hold credit card debt far too infrequently.

Table 9 summarizes these results. The first row of the Table reports our benchmark simulations (see previous two subsections) for the exponential and hyperbolic households in the HS group. Rows 2-15 report perturbations to these benchmark cases. In each of these rows, the benchmark simulation is perturbed by changing the calibration values of important parameters in the model.

Those perturbed parameters are identified in the first column of Table 9. To simulate behavior with the perturbed parameter values, we replicate the calibration procedure described in Section 5. Specifically, we numerically find the values of $\delta_{\text{exponential}}^{\text{HS}}$ and $\delta_{\text{hyperbolic}}^{\text{HS}}$ that generate simulated wealth accumulation that matches the SCF mid-life median wealth-to-income ratio. Hence, each row of Table 9 uses a new pair of values of $\delta_{\text{exponential}}^{\text{HS}}$ and $\delta_{\text{hyperbolic}}^{\text{HS}}$. Column two reports $\delta_{\text{exponential}}^{\text{HS}}$. Column three reports the simulated percentage of exponential consumers who borrow on their credit card at any point in time. Column four reports the average amount of credit card debt held by exponential consumers. Likewise, column five reports $\delta_{\text{hyperbolic}}^{\text{HS}}$, column six reports the simulated percentage of hyperbolic consumers who borrow on their credit card at any point in time, and column seven reports the average amount of credit card debt held by hyperbolic consumers.

All of the simulations in rows 2-15 have been implemented with partition jumps of \$2000 for the liquid asset and jumps of \$50,000 for the illiquid asset. By contrast, in the benchmark cases (row 1), we use a partition with jumps of \$500 for the liquid asset and jumps of \$10,000 for the illiquid asset.³⁸ We adopt a relatively coarse partition in rows 2-15, because many of these simulations are far more complex than the benchmark simulations (e.g., some of the state spaces and action spaces

³⁸The large partition jumps for the illiquid asset reflect the fact that illiquid assets tend to be lumpier than liquid assets.

are relatively large in these new runs). Even with the coarse partition, some of these robustness simulations take nearly two weeks to execute.

Row 2 matches the benchmark simulation, but adopts the relatively coarse partition. These results provide a check that changing the coarseness of the partition does not significantly change the original benchmark simulation results. The other reported robustness checks are summarized below:

Row 3 In the benchmark formula for effective household size, children are weighted with a factor of .4 relative to adults. The simulations reported in row 3 change the weighting on children from .4 to .6.

Row 4 In the benchmark simulations, disinvestment from the illiquid asset is not permitted. The simulations reported in row 4 allow such disinvestment and assume disinvestment transaction costs: a fixed cost of \$10,000 and a .1 proportional cost.

Row 5 In the benchmark simulations the required downpayment for the illiquid asset is 100 percent. The simulations reported in row 5 assume a downpayment of 10 percent.

Row 6 In the benchmark simulations the real interest rate on credit card debt is 11.75 percent. The simulations reported in row 6 assume a credit card interest rate of 9.75 percent.

Row 7 The simulations reported in row 7 assume a credit card interest rate of 13.75 percent.

Row 8 In the benchmark simulations, bankruptcy is not allowed. The simulations reported in row 8 allow households to declare bankruptcy and set the credit card interest rate to 13.75 percent.³⁹

Row 9 In the benchmark simulations the coefficient of relative risk aversion, ρ , is set to two. The simulations reported in row 9 assume $\rho = 1$.

³⁹Since households can declare bankruptcy, we no longer need to set a lower credit card interest rate to account for non-payment.

Row 10 The simulations reported in row 10 assume $\rho = 3$.

Row 11 In the benchmark hyperbolic simulations, the hyperbolic discount parameter, β , is set to .7. The hyperbolic simulation reported in row 11 assumes $\beta = .6$.

Row 12 The hyperbolic simulation reported in row 12 assumes $\beta = .8$.

Row 13 In the benchmark simulations the altruism parameter, α^{Bequest} , is set to one. The simulations reported in row 13 assume $\alpha^{\text{Bequest}} = .5$.

Row 14 In the benchmark simulations the total consumption flow from the illiquid asset, γ , is 5 percent per year. The simulations reported in row 14 assume a flow of 6 percent.

Table 9 demonstrates two points. First, the simulation results are not sensitive to our model and calibration assumptions. No reasonable variation in the modeling assumptions drives the simulated exponential borrowing far from the levels in our benchmark simulation. Second, calibrated hyperbolic households always borrow between two and four times as often as their exponential counterparts. This difference arises, even though hyperbolic and exponential consumers accumulate identical levels of pre-retirement wealth.

8. CONCLUSION

Consumers appear to be of two minds. Relatively large voluntary retirement accumulations imply exponential discount rates of only five percent. However, frequent credit card borrowing implies exponential discount rates of eighteen percent. It does not appear to be possible to calibrate realistic lifecycle models to match both observed levels of voluntary retirement savings and the observed frequency of credit card borrowing. We call this apparent paradox, The Debt Puzzle.

We have also suggested a resolution to this puzzle. If consumers have hyperbolic discount functions, then they may act both patiently *and* impatiently. Our calibrated simulations show that hyperbolic consumers will save aggressively for retirement, primarily in illiquid form, *and* borrow frequently in the credit card market.

9. REFERENCES

- Ameriks, John and Stephen P. Zeldes 2000. "How Do Household Portfolio Shares Vary With Age?" mimeo.
- Ausubel, Lawrence M. 1991. "The Failure of Competition in the Credit Card Market," *American Economic Review*, 81 (1), 50-81.
- Ainslie, George. 1992 *Picoeconomics*, Cambridge: Cambridge University Press.
- Akerlof, George A. 1991. "Procrastination and Obedience." *American Economic Review, Papers and Proceedings*, 1-19.
- Attanasio, Orazio. 1999. "Consumption," in (eds.) John Taylor and Michael Woodford, "Handbook of Macroeconomics", North Holland.
- , and Guglielmo Weber. 1993. "Consumption Growth, the Interest Rate, and Aggregation," *Review of Economic Studies* 60(3), 631-49.
- Barro, Robert. 1997. "Myopia and inconsistency in the neoclassical growth model," Harvard 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, 57-80.
- Brito, Dagobert L. and Peter R. Hartley 1995. "Consumer Rationality and Credit Cards," *Journal of Political Economy*, 103:2, 400-433.
- Carroll, Christopher D. 1992. "The Buffer Stock Theory of Saving: Some Macroeconomic Evidence." *BPEA*, 2:1992, 61-156.

- . 1997a. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112: 1-57.
- . 1997b. "Death to the log-linearized consumption Euler equation! (And very poor health to the second-order approximation)," NBER working paper 6298.
- 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: 267-72.
- Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica* 59: 1221-48.
- Deaton, Angus, and John Muellbauer 1986. "On Measuring Child Costs: With Applications to Poor Countries," *Journal of Political Economy*, 94:4, 720-744.
- Diamond, Peter and Botond Koszegi. 1998. "Hyperbolic discounting and retirement," MIT mimeo.
- Engen, Eric M., and William G. Gale. 1993. "IRAs and Saving: A Life-Cycle Consumption Simulation." UCLA mimeo.
- , ———, and John Karl Scholz. 1994. "Do Saving Incentives Work?" *BPEA* 1:1994, 85-180.
- , ———, and Cori Uccello. 1999. "The Adequacy of Retirement Saving" *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*.

- Gourinchas, Pierre-Olivier and Jonathan Parker. 1999. "Consumption Over the Life-Cycle." mimeo.
- Harris, Christopher, and David Laibson. 2000. "Dynamics Choices of Hyperbolic Consumers." forthcoming *Econometrica*.
- 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): 59-125.
- , ———, ———. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103: 360-99.
- 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." mimeo, Harvard University.
- . 1998. Comments on "Personal Retirement Saving Programs and Asset Accumulation," by James M. Poterba, Steven F. Venti, and David A. Wise, published in Ed. David A. Wise, "Studies in the Economics of Aging," NBER and the University of Chicago Press, 106-24.
- , Andrea Repetto, and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement," *Brookings Papers on Economic Activity* 1, pp.91-196.

- , and Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics* 57: 573-98.
- Morrison, Anne Klazien. (1998), "An Anomaly in Household Consumption and Savings Behavior: The Simultaneous Borrowing and Lending of Liquid Assets," mimeo, University of Chicago.
- Mulligan, Casey. 1997. "A Logical Economist's Argument Against Hyperbolic Discounting." University of Chicago mimeo.
- O'Donoghue, Ted, and Matthew Rabin. 1997a. "Doing It Now or Later." *American Economic Review*.
- , and ———. 1997b. "Incentives for Procrastinators." Working Paper, Northwestern University.
- Phelps, E. S., and R. A. Pollak. 1968 "On Second-best National Saving and Game-equilibrium Growth." *Review of Economic Studies* 35: 185-199.
- Repetto, Andrea. 1998. "Personal Bankruptcies and Individual Wealth Accumulation." mimeo, CEA, Universidad de Chile.
- Strotz, Robert H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies* 23: 165-180.
- Thaler, Richard H., and H. M. Shefrin. 1981. "An Economic Theory of Self-Control." *Journal of Political Economy* 89: 392-406.
- Zeldes, Stephen P. 1989. "Consumption and liquidity constraints: an empirical investigation," *Journal of Political Economy*, 97 (2), pp. 305-346.
- . 1989. "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence." *Quarterly Journal of Economics*. Vol. 104 (2). p 275-98. May 1989.

10. APPENDIX: ESTIMATING THE LIFE-CYCLE PROFILE OF THE FRACTION OF HOUSEHOLDS
BORROWING ON THEIR CREDIT CARDS.

This appendix describes the construction of the profiles of the proportion of households borrowing on their credit cards (see Figure 1). We construct these profiles so they can be compared directly to the profiles generated by our simulations.

The consumers in our model have the following characteristics: (1) they have always had a credit or charge card, (2) they have an exogenous level of education, (3) they experience no business cycle effects, and (4) they live in households with a head and a spouse. Using data from the SCF, we construct age profiles that condition on these four characteristics. We build the profiles as follows.

First, we exclude all households with heads younger than 20 or older than 90, and drop households who do not have credit or charge cards.

We then assign all households to one of the three educational groups in the model. We also assign households to a cohort group, in order to account for the fact that different generations of consumers have had differential access to revolving credit. Older cohorts may not have developed the habit of using credit cards because they did not have credit cards when they were young adults. We use 3 year cohorts, matching the frequency of the SCF, the youngest born in 1973-75, and the oldest born in 1898-1900. To account for business cycle effects in the data that are not in the model, we use the unemployment rate in the household's census division. Finally, we create a dummy variable to account for the head's marital status.

We then construct a dummy variable d_{it} equal to 1 if household i held credit card debt at date t , and 0 otherwise, and estimate the equation

$$d_{it} = f(\text{age}_{it}) + \beta_1 ce_i + \beta_2 \mu_{it} + \beta_3 ms_{it} + \beta_4 educ_{it} + \varepsilon_{it} \quad (14)$$

where $f(\text{age}_{it})$ is a function of the head's age, ce_i is a complete set of cohort dummies (excluding the last one), μ_{it} is the rate of unemployment in the household's census division, ms_{it} is a marital

status dummy equal to one if the head is married, $educ_{it}$ is a set of education dummies (excluding the HS group dummy), and ε_{it} represents classical measurement error. We estimate the equation by weighted least squares – using the SCF population weights – as well as by probit and logit models. Since these methods generated almost identical results, we only report the weighted least squares results.

To model the function $f(\cdot)$, we experimented with several alternatives, including polynomials in age, age dummies and linear splines, all generating similar results. Figure 1 plots the estimated profile of the proportion of HS households borrowing on their cards, using a spline with knots at ages 35, 50, 65, and 80. To match the characteristics of the households in the model, we set the youngest cohort dummy equal to one, and set all the other cohort dummies to zero. We also evaluate the unemployment rate at the average rate in the sample, and set the marital status dummy equal to one. Thus, the figure represents the fraction of married households born in the youngest cohort, facing the average unemployment rate.

Table 1. Credit Card Debt^{a,b}

	Conditional on Having a Credit Card			
	% with Card	% with Debt	Balance	
			Mean	Median
All categories				
20-29	0.72	0.77	1668	746
30-39	0.77	0.76	2114	772
40-49	0.85	0.72	2487	760
50-59	0.84	0.60	1603	343
60-69	0.83	0.43	980	0
70+	0.80	0.27	250	0
All ages	0.80	0.63	1715	343
No high school diploma				
20-29	0.68	0.83	1823	849
30-39	0.66	0.77	2559	943
40-49	0.77	0.84	2988	815
50-59	0.73	0.71	1910	549
60-69	0.71	0.55	1115	129
70+	0.76	0.35	285	0
All ages	0.72	0.68	1832	429
High school graduates				
20-29	0.60	0.84	1885	935
30-39	0.74	0.86	1673	858
40-49	0.81	0.73	2274	772
50-59	0.84	0.72	1424	515
60-69	0.85	0.44	722	0
70+	0.75	0.28	265	0
All ages	0.77	0.70	1537	472
College graduates				
20-29	0.89	0.65	1364	600
30-39	0.92	0.65	2213	532
40-49	0.93	0.64	2340	497
50-59	0.96	0.40	1545	0
60-69	1.00	0.26	1143	0
70+	0.93	0.13	180	0
All ages	0.93	0.53	1767	94

Source: Authors' calculations based on the 1995 SCF.

^a Includes traditional cards such as Visa, Mastercard, Discover and Optima, and other credit or charge cards such as Diners Club, American Express, store cards, airline cards, car rental cards, and gasoline cards. Excludes business and company cards.

^b The total credit card debt is constructed on the basis of the responses to the following SCF question:
 "After the last payments were made on this (these) account(s), roughly what was the balance still owed on this (these) account(s)?"

Table 2. Fraction of Households Borrowing on Credit Cards Across the Distribution of Wealth^{a,b}

Age group	Wealth Distribution Percentile			
	Less than 25	25-50	50-75	Over 75
All categories				
20-29	0.87	0.77	0.70	0.65
30-39	0.86	0.80	0.69	0.51
40-49	0.79	0.76	0.56	0.41
50-59	0.75	0.65	0.40	0.27
60-69	0.55	0.40	0.25	0.18
70+	0.48	0.26	0.11	0.05
Incomplete High School				
20-29	0.91	0.83	0.67	0.82
30-39	0.73	0.82	0.78	0.70
40-49	0.84	0.85	0.80	0.60
50-59	0.83	0.67	0.75	0.45
60-69	0.60	0.51	0.39	0.25
70+	0.57	0.30	0.24	0.10
High School Graduates				
20-29	0.89	0.78	0.82	0.73
30-39	0.90	0.83	0.83	0.66
40-49	0.86	0.79	0.74	0.50
50-59	0.79	0.72	0.55	0.40
60-69	0.60	0.42	0.31	0.24
70+	0.47	0.29	0.09	0.14
College Graduates				
20-29	0.81	0.65	0.51	0.56
30-39	0.82	0.61	0.55	0.39
40-49	0.71	0.53	0.44	0.20
50-59	0.63	0.38	0.24	0.22
60-69	0.41	0.20	0.09	0.10
70+	0.28	0.07	0.06	0.03

Source: Authors' calculations based on the 1983-1995 SCFs.

^a Conditional on having a credit card.

^b We calculated the fraction of households who are borrowing in each quartile of the wealth distribution contingent on age and education group, for every SCF year. The table reports the weighted average across the 4 SCF years, using the proportion of households with credit cards in a given year/category as weights.

Table 3. Estimated Age-Number of Children and Age-Number of Dependent Adults Profiles

	Children	Dependent Adults
High School Dropouts		
Constant	0.12143 (0.0111)	0.00002 (0.0000)
Age	0.16690 (0.0048)	0.41411 (0.0125)
Age ²	0.00238 (0.0001)	0.00396 (0.0001)
High School Graduates		
Constant	0.00613 (0.0006)	8E-09 (0.0000)
Age	0.32402 (0.0054)	0.72718 (0.0160)
Age ²	0.00450 (0.0001)	0.00713 (0.0002)
College Graduates		
Constant	0.00005 (0.0000)	4E-12 (0.0000)
Age	0.55628 (0.0139)	1.00347 (0.0413)
Age ²	0.00729 (0.0002)	0.00965 (0.0004)

Source: Author's calculations based on data from the PSID. Standard errors in parenthesis.

We estimated the following model by NLLS:

$$x_{it} = \beta_0 \exp(\beta_1 \text{age}_{it} - \beta_2 \text{age}_{it}^2) + \varepsilon_{it}$$

where x_{it} is either the number of dependent children or the number of dependent adults in the household, and ε_{it} represents iid noise.

Table 4. Estimated Age-Non Asset Income Profiles ^a

	Less than High School	High School Graduates	College Graduates
In the labor force ^b			
Age	0.077 (0.039)	0.118 (0.021)	0.223 (0.038)
Age ² /100	-0.172 (0.074)	-0.201 (0.050)	-0.390 (0.086)
Age ³ /10000	0.092 (0.045)	0.081 (0.035)	0.204 (0.059)
N head and spouse	0.668 (0.035)	0.548 (0.019)	0.462 (0.032)
N kids	0.012 (0.010)	-0.033 (0.005)	-0.023 (0.008)
N dep. adults	0.167 (0.011)	0.170 (0.008)	0.022 (0.021)
Other effects ^d	7.958	7.439	6.029
Retired ^c			
Age	-0.039 (0.024)	-0.002 (0.013)	-0.009 (0.008)
N head and spouse	0.656 (0.316)	0.554 (0.084)	0.327 (0.140)
N kids	0.042 (0.096)	0.199 (0.172)	-0.560 (0.102)
N dep. adults	0.421 (0.092)	0.204 (0.102)	0.162 (0.081)
Other effects ^d	9.927	8.433	10.172

Source: Authors' calculations based on data from the PSID 1983-90.

^a The dependent variable is the natural logarithm of non-asset after tax household income. It includes lump sum payments such as inheritances. Standard errors are in parenthesis.

^b A household is in the labor force if anyone in the household worked 1500 hours or more in that year or in any subsequent year.

^c A household is retired if no member works more than 500 hours per year in the current year or in any year in the future.

^d Includes the effects of a constant, cohort dummies, and the unemployment rate in the household's state of residence, evaluated with each regressor set equal to its sample mean.

Table 5. Estimated Age-Income Processes

	Less than High School	Completed High School	Completed College
In the Labor Force^{a,b}			
α	0.881 (0.022)	0.782 (0.006)	0.967 (0.007)
Variance of ε	0.024 (0.006)	0.029 (0.003)	0.019 (0.002)
Variance of υ	0.041 (0.005)	0.026 (0.003)	0.014 (0.002)
Retired^{c,d}			
Variance of υ	0.077 (0.019)	0.051 (0.013)	0.042 (0.013)

Source: Authors' calculations based on data from the PSID, 1983-90.

Standard errors in parenthesis.

The coefficient α and the variances of ε and υ were estimated using GMM.

^a A household is in the labor force if anyone in the household worked 1500 hours or more in that year or in any subsequent year.

^b Model estimated, using the residuals of the regressions reported in Table 4:

$$\xi_{it} = \zeta_i + u_{it} + v_{it} = \zeta_i + \alpha u_{it-1} + \varepsilon_{it} + v_{it}$$

^c A household is retired if no member works more than 500 hours per year in the current year or in any year in the future.

^d Model estimated, using the residuals of the regressions reported in Table 4:

$$\xi_{it} = \zeta_i + v_{it}$$

Table 6. Wealth-Income Ratios

Age Group	Means					Medians				
	1983 ^a	1989	1992	1995	Average	1983 ^a	1989	1992	1995	Average
All categories										
20-29	1.26	3.29	1.07	1.42	1.76	0.45	0.41	0.42	0.52	0.45
30-39	2.97	2.70	2.59	2.38	2.66	1.32	1.27	1.03	1.14	1.19
40-49	5.16	6.69	4.78	4.98	5.40	2.07	2.45	1.87	1.84	2.06
50-59	8.00	8.06	8.82	8.03	8.23	2.91	3.90	3.87	3.34	3.50
60-69	11.82	19.56	15.30	14.43	15.28	4.07	5.73	5.14	5.13	5.02
70+	13.06	24.08	21.35	24.91	20.85	4.67	7.02	10.13	8.30	7.53
Incomplete High School										
20-29	0.54	1.49	0.78	0.93	0.94	0.22	0.32	0.31	0.42	0.32
30-39	1.87	2.26	1.71	1.65	1.87	0.52	1.27	0.58	0.76	0.78
40-49	3.13	6.64	3.43	4.22	4.35	1.07	2.02	1.53	1.30	1.48
50-59	3.67	6.21	4.44	5.82	5.03	2.29	3.41	2.19	2.16	2.51
60-69	7.19	14.25	9.59	9.73	10.19	2.98	5.00	3.73	3.30	3.75
70+	9.67	24.81	16.56	18.42	17.37	3.75	5.97	9.05	6.95	6.43
High School Graduate										
20-29	1.40	2.63	1.10	1.44	1.64	0.46	0.40	0.37	0.47	0.42
30-39	3.08	1.97	2.59	2.22	2.47	1.22	0.86	0.94	1.17	1.05
40-49	3.72	4.11	2.32	3.94	3.52	2.20	2.33	1.22	1.69	1.86
50-59	11.39	7.53	9.18	6.51	8.65	2.78	3.69	3.75	2.74	3.24
60-69	13.10	18.06	15.80	15.35	15.57	4.31	6.53	5.44	6.55	5.71
70+	18.55	21.74	21.79	23.46	21.39	6.08	7.85	10.90	9.25	8.52
College Graduate										
20-29	1.31	5.91	1.31	1.97	2.63	0.63	0.82	0.46	0.92	0.71
30-39	3.20	3.72	3.23	3.23	3.34	1.75	1.58	1.44	1.35	1.53
40-49	9.49	8.85	7.34	6.22	7.97	2.33	3.28	2.69	2.42	2.68
50-59	7.90	11.19	12.39	12.12	10.90	3.57	4.78	4.71	4.32	4.34
60-69	21.89	34.40	23.15	21.73	25.29	7.98	8.38	8.49	9.05	8.48
70+	18.08	24.34	32.09	39.35	28.47	11.03	9.85	12.89	14.09	11.97

Sources: SCF, Social Security Administration, Congressional Budget Office and Pechman (1989).

Income is after tax non-asset income, plus bequests. Taxes include Social Security deductions, and Federal income taxes. Social Security deductions were imputed using OASDI-HI tax rates and maximum taxable earnings. Federal income taxes were imputed using effective tax rates as reported by the CBO and Pechman.

^a Bequests are imputed using Laibson, Repetto and Tobacman (1998) calculations.

Table 7. Calibrated Long-term Discount Rates^a

	Exponential Consumers	Hyperbolic Consumers
High school dropouts	0.0880	0.0700
High school graduates	0.0560	0.0440
College graduates	0.0550	0.0440

Source: Authors' calculations.

^a The table reports the long term discount rates (1- δ) that replicate the average wealth-income ratios for households with heads between ages 50 and 59, as reported in Table 6.

Table 8. Simulated Share Borrowing Across the Wealth Distribution^a

Age group	Wealth quartile			
	0-25	25-50	50-75	75+
Incomplete High School - exponential				
20-29	1.00	0.34	0.00	0.00
30-39	1.00	0.06	0.01	0.00
40-49	0.68	0.19	0.05	0.00
50-59	0.45	0.20	0.05	0.00
60-69	0.09	0.06	0.05	0.00
70+	0.26	0.30	0.31	0.37
Incomplete High School - hyperbolic				
20-29	1.00	0.89	0.15	0.13
30-39	1.00	0.68	0.41	0.24
40-49	0.91	0.65	0.49	0.37
50-59	0.75	0.55	0.43	0.31
60-69	0.46	0.43	0.40	0.39
70+	0.72	0.84	0.96	0.98
High School Graduates - exponential				
20-29	1.00	0.25	0.00	0.00
30-39	0.79	0.07	0.02	0.00
40-49	0.54	0.20	0.09	0.02
50-59	0.33	0.17	0.09	0.04
60-69	0.07	0.05	0.04	0.03
70+	0.41	0.33	0.32	0.14
High School Graduates - hyperbolic				
20-29	1.00	0.74	0.17	0.10
30-39	1.00	0.56	0.36	0.19
40-49	0.84	0.60	0.42	0.24
50-59	0.73	0.54	0.44	0.27
60-69	0.56	0.57	0.70	0.45
70+	0.93	0.97	0.98	0.32
College Graduates - exponential				
20-29	1.00	0.98	0.01	0.00
30-39	1.00	0.32	0.01	0.00
40-49	0.70	0.06	0.02	0.03
50-59	0.64	0.14	0.11	0.01
60-69	0.90	0.26	0.02	0.00
70+	0.59	0.10	0.00	0.00
College Graduates - hyperbolic				
20-29	1.00	1.00	0.38	0.03
30-39	1.00	0.90	0.13	0.06
40-49	1.00	0.85	0.24	0.11
50-59	1.00	0.73	0.22	0.00
60-69	1.00	0.57	0.01	0.00
70+	1.00	0.52	0.00	0.00

Source: Author's simulations.

^a Fraction of simulated households who borrow in each wealth quartile of an age-education group.


Table 9. Robustness Checks^a

	Exponential Simulations			Hyperbolic Simulations		
	Calibrated Discount Rate	Proportion Borrowing	Average Debt	Calibrated Discount Rate	Proportion Borrowing	Average Debt
1 Benchmarks (with fine partition)	0.056	0.21	\$907.54	0.044	0.57	\$3,748.28
2 Benchmarks (with coarse partition)	0.056	0.18	\$904.99	0.042	0.48	\$3,234.01
3 Heavier weight on children ($\kappa = .6$)	0.052	0.16	\$798.71	0.040	0.43	\$3,060.63
4 Reversible investment in Z	0.056	0.17	\$890.78	0.042	0.44	\$3,109.76
5 Debt-financed purchase of Z	0.059	0.21	\$1,075.90	0.049	0.59	\$4,036.43
6 Credit card interest rate 9.75	0.056	0.21	\$1,114.58	0.042	0.51	\$3,582.20
7 Credit card interest rate 13.75	0.056	0.15	\$716.75	0.042	0.44	\$2,923.80
8 Bankruptcy allowed (interest rate 13.75)	0.056	0.16	\$969.90	0.042	0.42	\$3,228.59
9 CRRRA = 1	0.049	0.15	\$553.73	0.036	0.55	\$4,038.07
10 CRRRA = 3	0.063	0.17	\$880.32	0.049	0.38	\$2,460.16
11 $\beta = .6$	N/A	N/A	N/A	0.038	0.54	\$3,891.19
12 $\beta = .8$	N/A	N/A	N/A	0.046	0.38	\$2,347.36
13 Altruism parameter = .5	0.052	0.18	\$876.40	0.036	0.52	\$3,435.13
14 Illiquid rate of return = .06	0.063	0.27	\$1,340.07	0.047	0.56	\$3,893.96

Source: Authors' simulations.


^a The table shows the average amount borrowed and the fraction of households borrowing for different calibration assumptions. All education groups are included.

Figure 1. Fraction of Households Borrowing on their Credit Cards




Source: SCF. Calculated from a regression on a linear spline in age, cohort dummies, the unemployment rate, a marital status dummy, and a set of education dummies.

Figure 2. Discount functions




Source: Authors' calculations. Exponential: δ^t , with $\delta=0.939$; hyperbolic: $(1+\alpha)^{-\beta t}$, with $\alpha=4$ and $\beta=1$; and quasi-hyperbolic: $(1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots)$, with $\beta=0.7$ and $\delta=0.957$.

Figure 3. Implied discount rates




Source: Authors' calculations. Figure plots the exponential discount rates and coefficients of relative risk aversion that are consistent with a second order approximation of the Euler equation, under alternative calibration values for the real interest rate (.14 and .13), the expected one year rate of consumption growth (0 and .01), and the conditional variance of consumption growth (.025 and .015).

Figure 4. Number of Children and of Dependent Adults over the Life-Cycle




Source: PSID. Values were estimated using nonlinear least squares on the model $x_{it} = \beta_0 + \beta_1 \text{age}_{it} + \beta_2 \text{age}_{it}^2 + \epsilon_{it}$, where x_{it} is either the number of dependent children or the number of dependent adults in the household, and ϵ_{it} represents iid noise.

Figure 5. Estimated Age-Income Profiles



Source: PSID. Figure plots estimated non-asset after-tax household income, by age and education group. Values are calculated from a regression of the log of income on a cubic polynomial in age, cohort dummies, family size, and the unemployment rate. The figure plots the age effects, with other regressors set equal to their means.

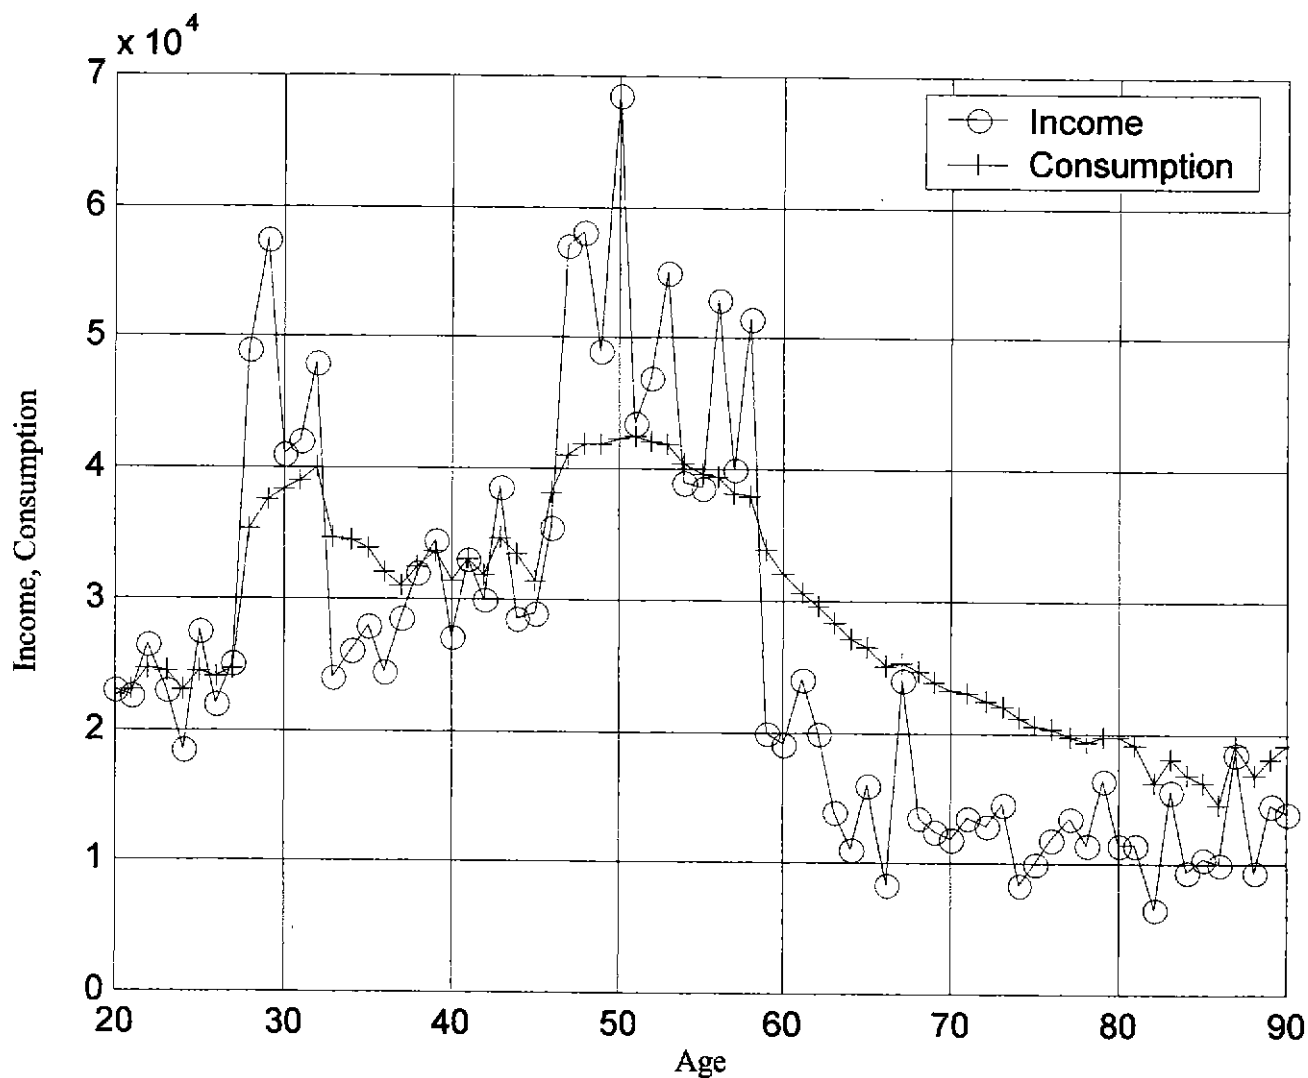
Figure 6: Simulated Mean Income and Consumption of Exponential Households



Source: Authors' simulations.

The figure plots the simulated average values of consumption and income for households with high school graduate heads.

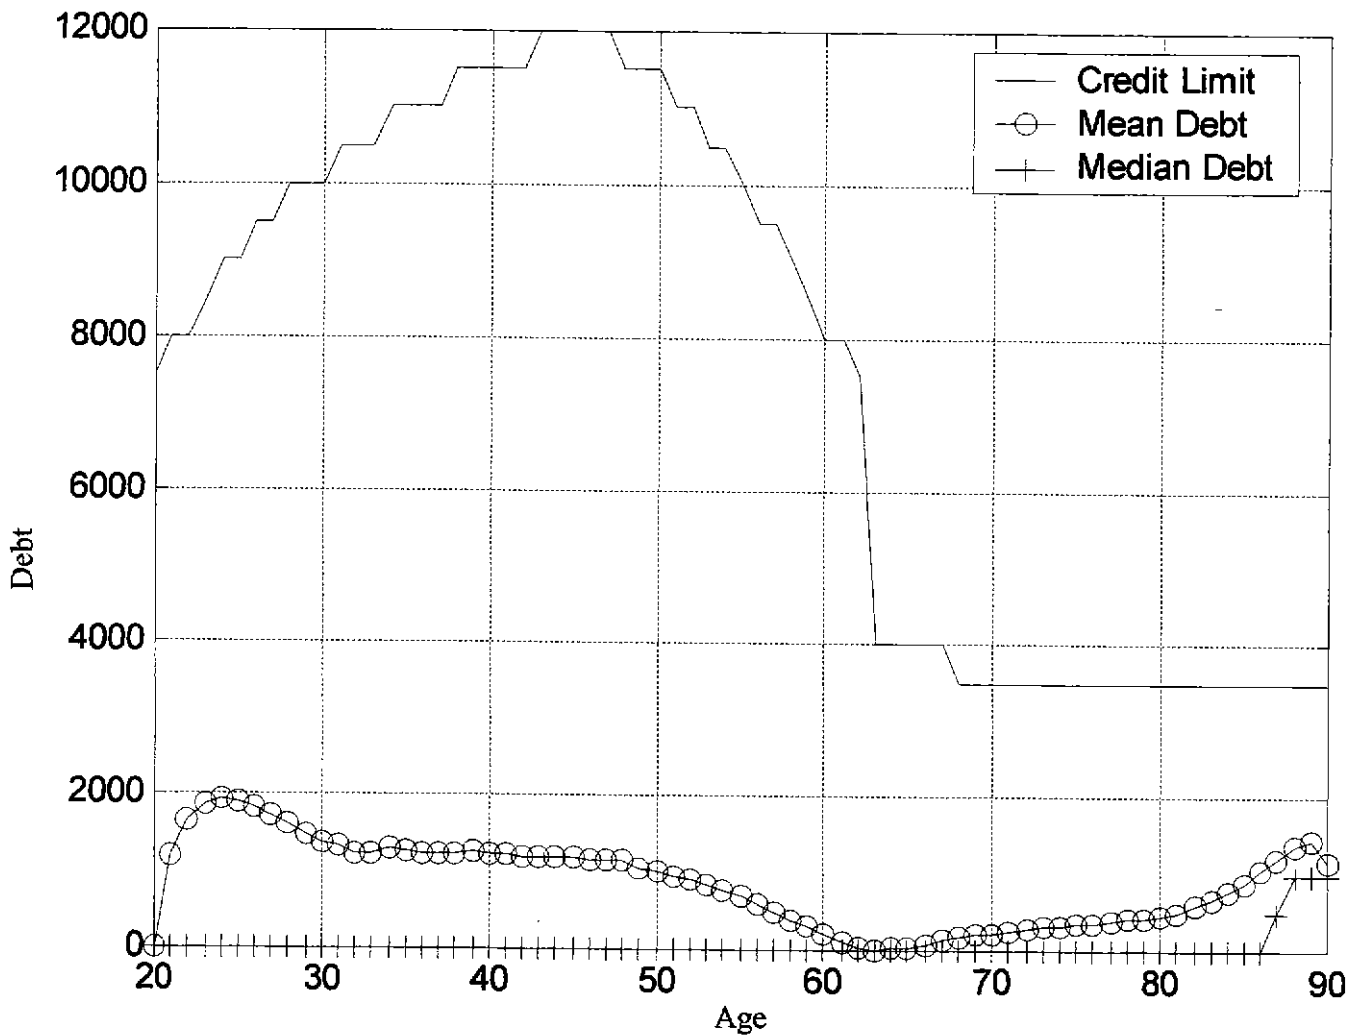
Figure 7 : Simulated Income and Consumption of a Typical Exponential Household



Source: Authors' simulations.

The figure plots the simulated life cycle profiles of consumption and income for a typical household with a high school graduate head.

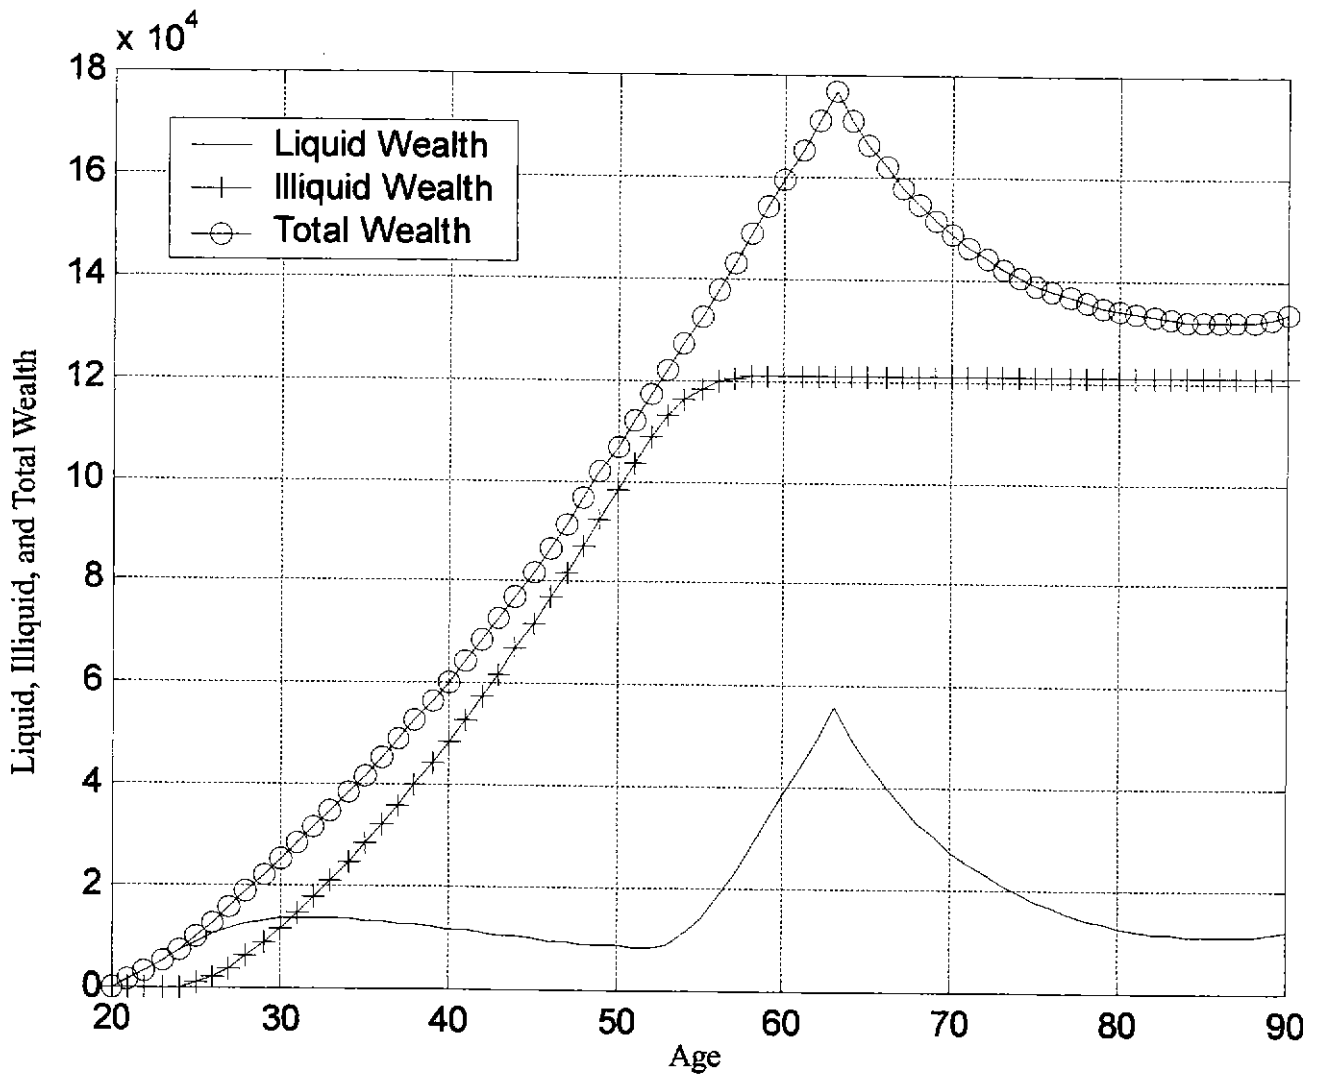
Figure 8: Simulated Mean & Median Debt, and Credit Limit, for Exponential Households



Source: Authors' simulations.

The figure plots the simulated median and mean amount of credit card debt, along with the age-dependent credit limit, for households with high school graduate heads.

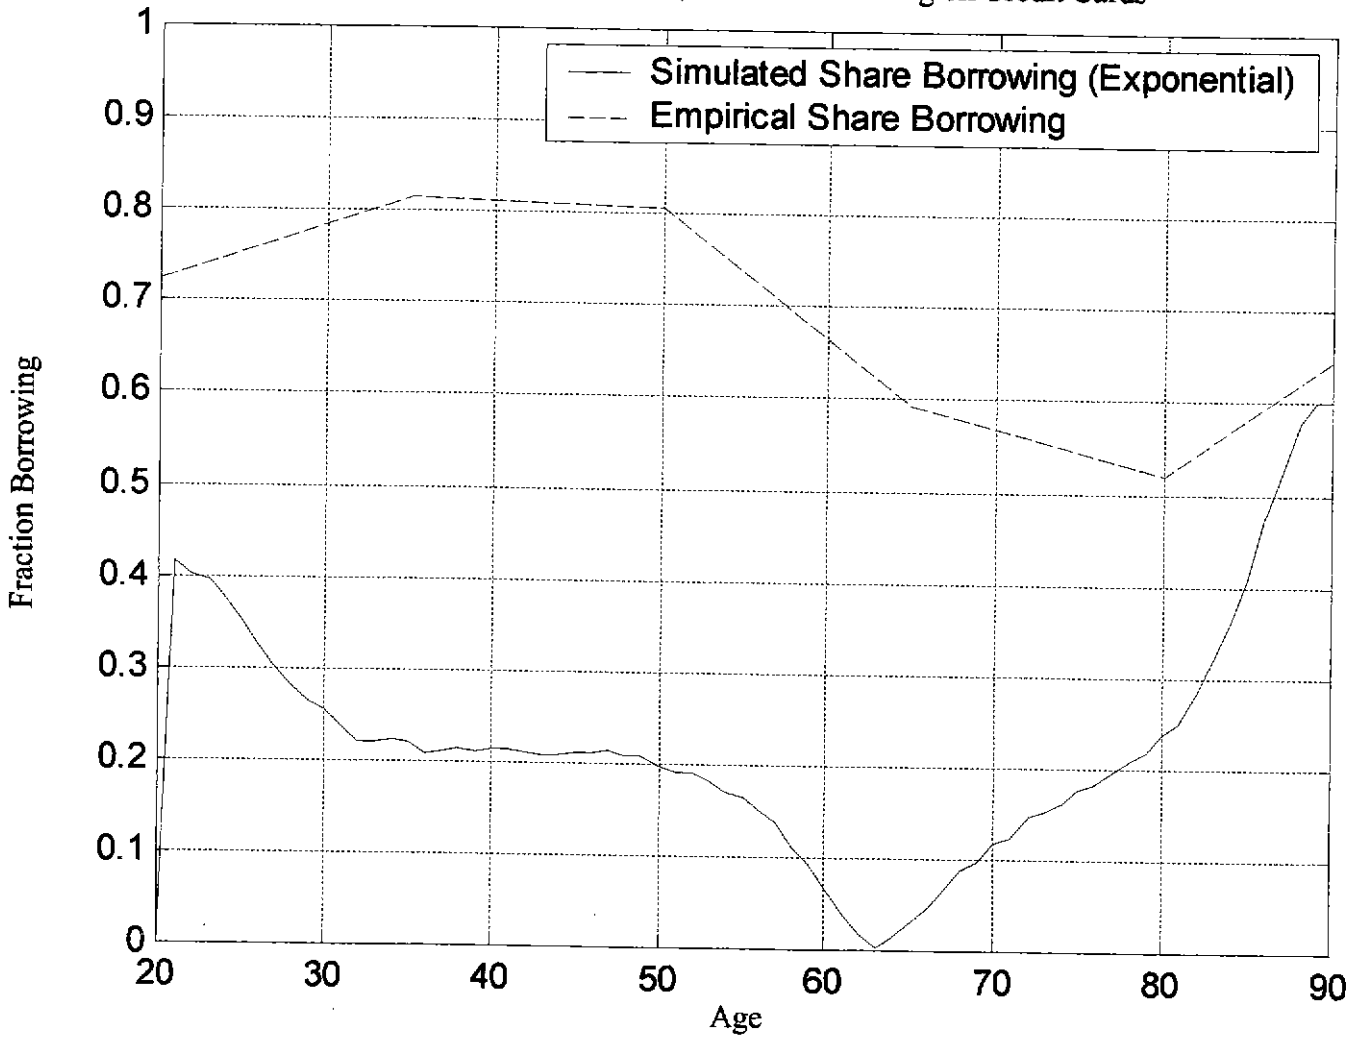
Figure 9: Simulated Mean Liquid, Illiquid, and Total Wealth for Exponential Households



Source: Authors' simulations.

The figure plots the simulated mean level of liquid, illiquid and total wealth for households with high school graduate heads.


Figure 10: Fraction of Households Borrowing on Credit Cards



Source: Authors' simulations, and Survey of Consumer Finances.

The figure plots the simulated fraction of households with a high school graduate head who are borrowing on their credit cards, along with the estimated life-cycle profile from Figure 1.

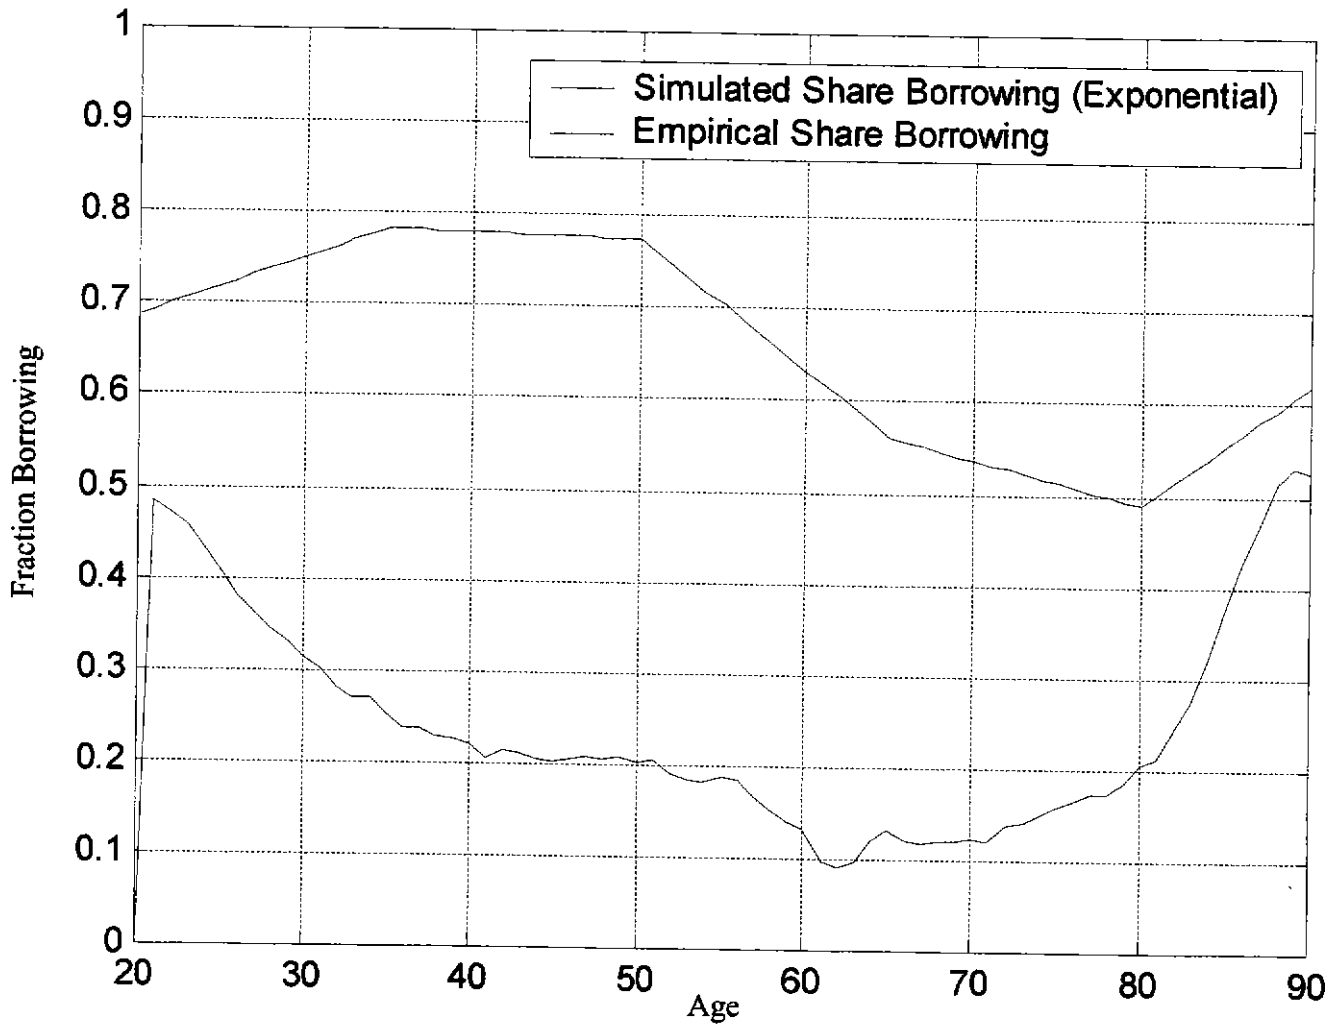
Figure 11: Simulated Mean & Median Debt, and Credit Limit, for Exponential Households



Source: Authors' simulations.

The figure plots the simulated median and mean amount of credit card debt, along with the age-dependent mean credit limit, for all educational groups.


Figure 12: Fraction of Households Borrowing on Credit Cards



Source: Authors' simulations, and Survey of Consumer Finances.

The figure plots the simulated fraction of households who are borrowing on their credit cards, along with the estimated life-cycle profile from Figure 1, for all educational groups.


Figure 13: Mean Consumption of Exponential and Hyperbolic Households



Source: Authors' simulations.

The figure plots average consumption over the life-cycle for simulated exponential and hyperbolic households with high school graduate heads.


Figure 14: Mean Total Wealth of Exponential and Hyperbolic Households



Source: Authors' simulations.

The figure plots average wealth over the life-cycle for simulated exponential and hyperbolic households with high school graduate heads.


Figure 15: Mean Illiquid Wealth of Exponential and Hyperbolic Households



Source: Authors' simulations.

The figure plots average illiquid wealth over the life-cycle for simulated exponential and hyperbolic households with high school graduate heads.

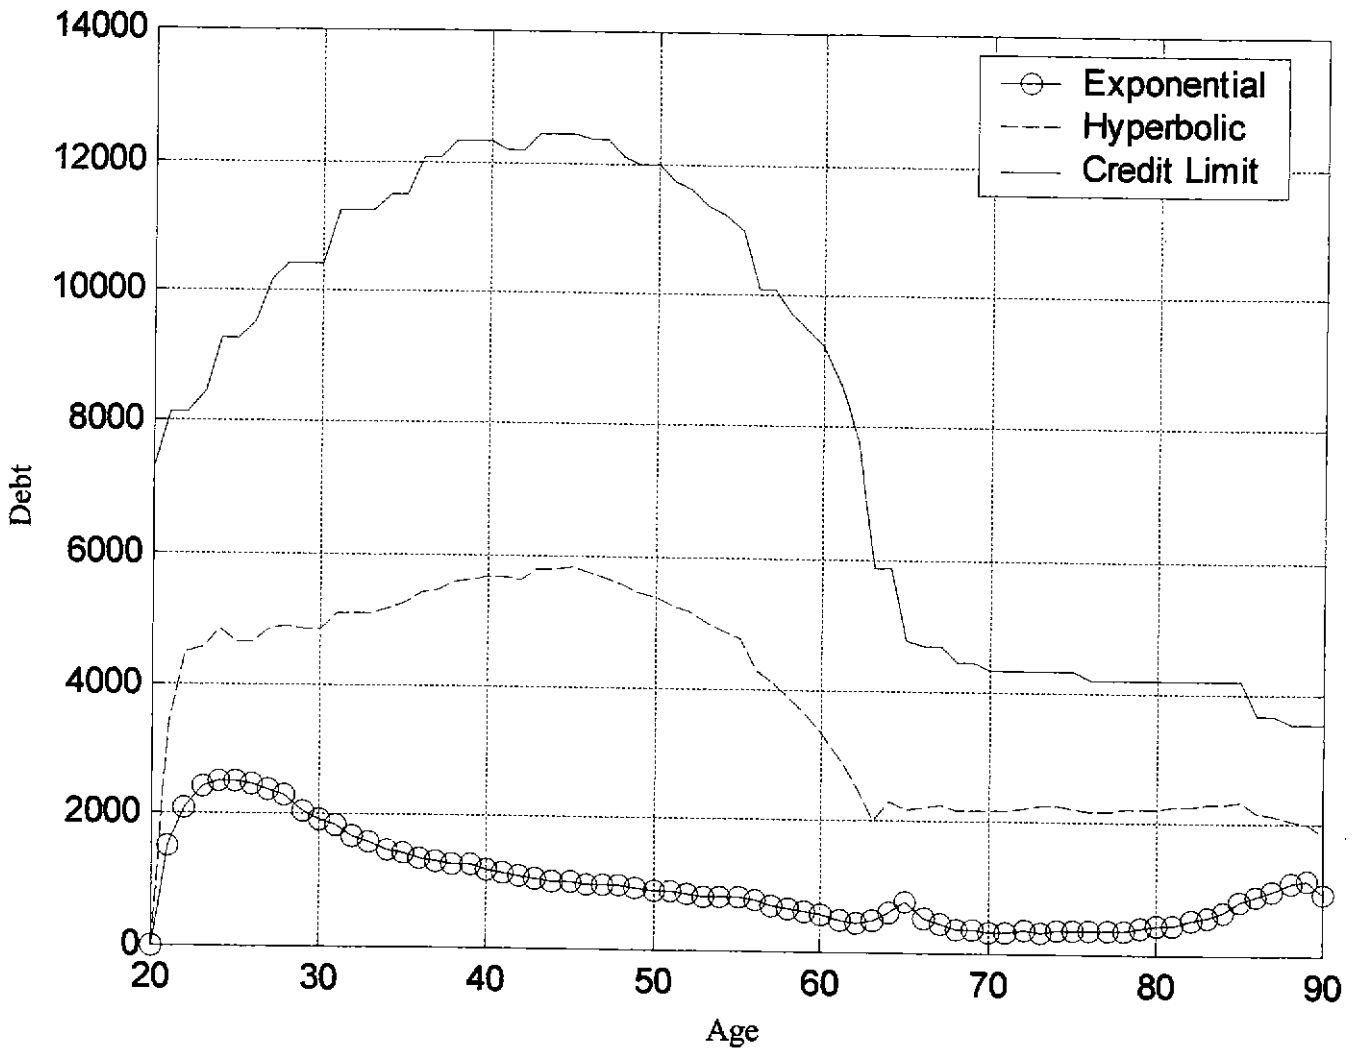
Figure 16: Mean Liquid Wealth of Exponential and Hyperbolic Households



Source: Authors' simulations.

The figure plots average liquid wealth over the life-cycle for simulated exponential and hyperbolic households with high school graduate heads.

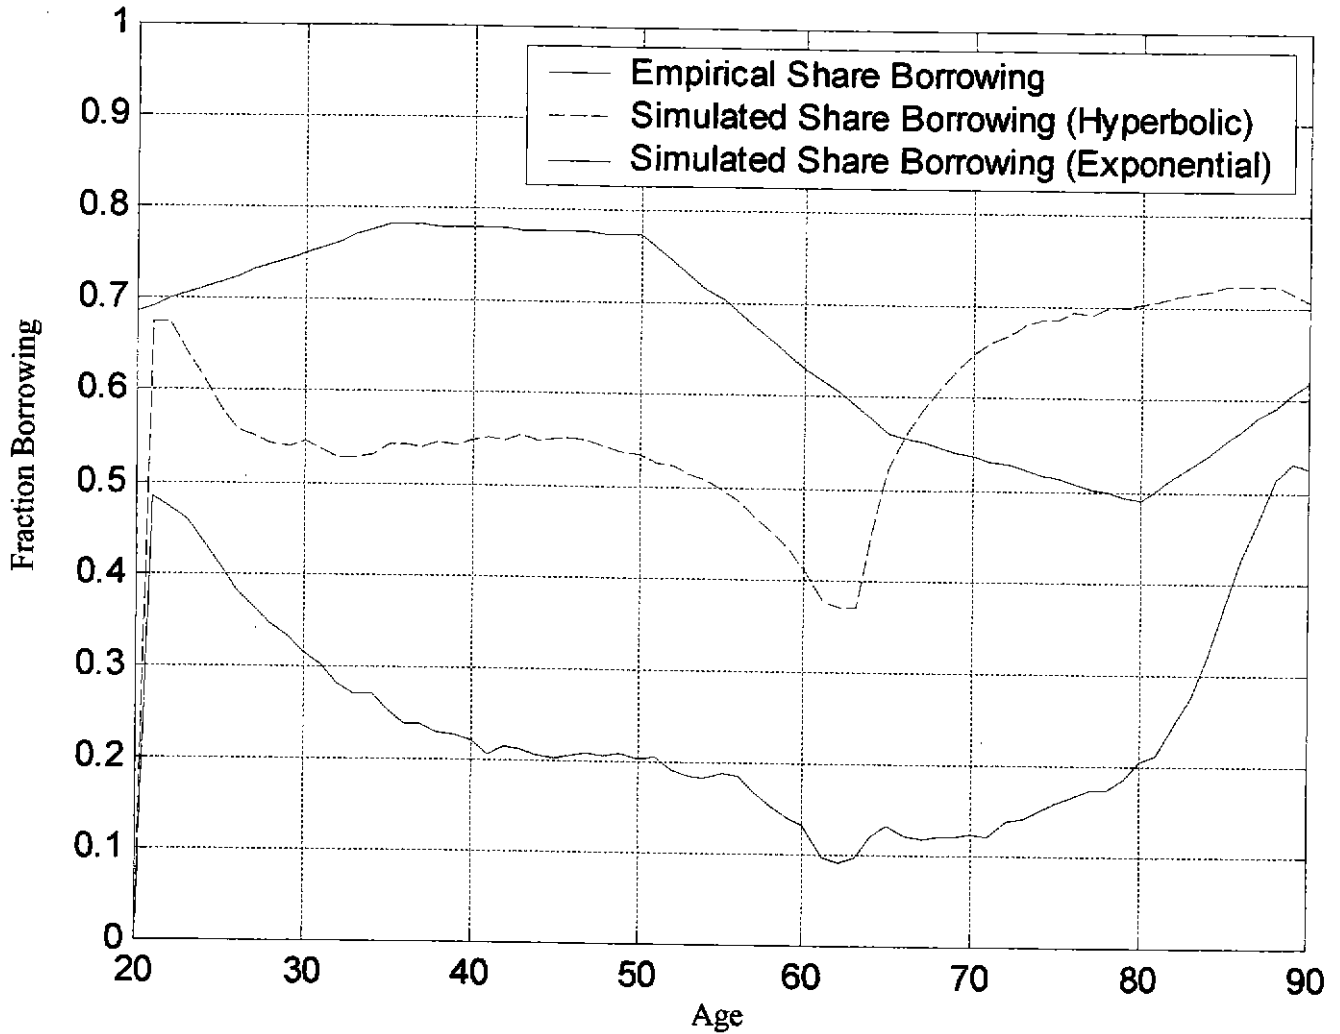
Figure 17: Mean Credit Limit and Simulated Mean Debt for Exponential and Hyperbolic Households



Source: Authors' simulations.

The figure plots the mean credit limit and mean credit card debt over the life-cycle for simulated exponential and hyperbolic households, for all educational groups.

Figure 18: Fraction of Households Borrowing on Credit Cards



Source: Authors' simulations, and Survey of Consumer Finances.

The figure plots the simulated fraction of exponential and hyperbolic households who are borrowing on their credit cards, along with the estimated life-cycle profile from Figure 1, for all educational groups.