Wealth Accumulation, Credit Card Borrowing, and Consumption-Income Comovement

David Laibson, Andrea Repetto, and Jeremy Tobacman

October 2001*

Preliminary and incomplete; please do not cite.

Abstract. Research in the consumption literature usually attempts to explain a single empirical regularity at a time. In this paper we ask whether a single consumption model can simultaneously explain a wide range of facts about wealth accumulation, credit card borrowing, and consumption-income comovement. Our analysis implements a lifecycle simulation framework that incorporates dependents, stochastic labor income, liquidity constraints, liquid assets, revolving credit, illiquid assets, and retirement. We use the Method of Simulated Moments to econometrically compare the choices of our simulated consumers with the choices of actual consumers as calculated from household survey data. We reject the null hypothesis of exponential discounting in favor of the hyperbolic discounting alternative. However, the hyperbolic model also fails to explain all of the micro-survey evidence. High levels of wealth accumulation and high frequencies of credit card borrowing are difficult to reconcile with any existing model.

*David Laibson is Paul Sack Associate Professor of Political Economy, Harvard University, and Faculty Research Fellow, National Bureau of Economic Research, Cambridge, Massachusetts. Andrea Repetto is Assistant Professor of Industrial Engineering, University of Chile, Santiago, Chile. Jeremy Tobacman is a graduate student in economics, Harvard University, Cambridge, Massachusetts. We received helpful advice from Alberto Alesina, Robert Barro, Christopher Carroll, Greg Mankiw, Ariel Pakes, Sam Thompson, Nicholas Souleles, and seminar audiences at Harvard University, Staford University, and the NBER. Laibson acknowledges financial support from the National Science Foundation (SBR-9510985), the National Institute on Aging (R01-AG-16605), the MacArthur Foundation and the Olin Foundation; Repetto, from DID-Universidad de Chile, FONDECYT (1990004), Fundación Andes and an institutional grant to CEA from the Hewlett Foundation; and Tobacman from a National Science Foundation Graduate Research Fellowship.

JEL classification: D91 (Intertemporal Consumer Choice; Lifecycle Models and Saving), E21 (Consumption; Saving)

1. INTRODUCTION

American households voluntarily accumulate large stocks of wealth as they approach retirement. The wealth-to-income ratio of the median household aged 50-59 is over three (Survey of Consumer Finances). Consumers *also* borrow frequently and extensively on credit cards over the lifecycle. At any point in time, seventy percent of households with credit cards carry interest-accruing balances. Aggregating over *all* households with a card, credit card balances not including the float average more than \$6000 per household (Survey of Consumer Finances, Federal Reserve Board). The average debt-weighted credit card interest rate is 16% (Federal Reserve Board). This high cost borrowing suggests that most households are effectively liquidity constrained. Perhaps as a consequence of this, consumption exhibits excess sensitivity to predictable changes in income. Consumption-income comovement regressions typically yield estimates between 0 and $\frac{1}{2}$ for the marginal propensity to consume out of predictable changes in income, with "consensus" estimates around $\frac{1}{4}$ (Hall, 1978; Hall and Mishkin, 1982; Campbell and Mankiw, 1989; Shea, 1995). One distinctive manifestation of this consumption-income comovement is the anomalous drop in consumption in the years around retirement (Banks et al. 1998, Bernheim et al. 1997).

Research in the consumption literature usually attempts to explain one or two empirical regularities at a time. In this paper we ask whether a *single* model can simultaneously explain *all* of the regularities summarized above. To do this, we build an institutionally rich model of lifecycle decision-making. Our model includes income uncertainty, retirement, liquidity constraints, age-dependent mortality, children and adult dependents, revolving credit, and illiquid investment instruments. Our model also incorporates hyperbolic discounting, in which discount rates are higher in the short run than in the long run. In particular, we assume that time preferences are characterized by the discount function, $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, ...\}$, with $\beta \leq 1$ and $\delta \leq 1$ (Phelps and Pollak 1968, Laibson 1997a). This formulation reduces to the standard exponential discount function when $\beta = 1$.

We use the Method of Simulated Moments to estimate the discounting parameters in our model. The exponential discounting model is overwhelmingly rejected by the data, but neither the exponential model nor the hyperbolic model can match *both* the wealth accumulation *and* credit card borrowing facts. Consumers appear to be of two minds, acting both patiently and impatiently. Even a dynamically inconsistent hyperbolic discounting model cannot completely explain these dichotomous behaviors.

The paper is organized as follows. Section 2 summarizes the key features of the model and our calibration choices. The empirical data used for the estimation are presented in Section 3. We explain the MSM procedure in Section 4. Section 5 presents our results, and Section 6 concludes.

2. MODEL

We adopt a variant of the model developed in Laibson, Repetto, and Tobacman (2001), hereafter LRT. This model is based on the simulation literature pioneered by Carroll (1992, 1997), Deaton (1991), and Zeldes (1989b) and extended by Hubbard, Skinner, and Zeldes (1994, 1995), Engen, Gale, and Scholz (1994), Gourinchas and Parker (2001), and Laibson, Repetto, and Tobacman (1998). LRT (2001) incorporates most of the features of earlier lifecycle simulation models and adds new ones, including credit cards, age-dependent household size, and access to an illiquid asset. We summarize the key features of the model below. Interested readers are referred to LRT (2001) for details and for extensions including options to declare bankruptcy and to borrow against illiquid collateral (as in mortgages against housing).

We divide the presentation of the model into eight parts: demographics, income from transfers and wages, liquid assets and noncollateralized debt, illiquid assets, dynamic and static budget constraints, preferences, equlibrium, simulation.

2.1. Demographics. We consider households in which the the household head has a high school diploma but no college degree. These households constitute 59% of the population (US Census Bureau, 1995). In previous work (LRT 1998, 2001) we have found qualitatively similar results when examining households whose head does not have high school degree, and when analyzing households whose head has both a high school degree and a college degree.

We assume that independent economic life begins at age 20. Households face a time-varying, exogenous hazard rate of survival, s_t, calibrated with data from the U.S. National Center for Health Statistics (1993). To simplify our computational analysis, we assume that no household lives past age 90. Household composition varies deterministically and exogenously with age (calibrated from the PSID) as children and adult dependents enter and leave the household. Following Blundell et al. (1994) and Attanasio and Weber (1995), we define effective household size as the number of adults plus .4 times the number of children under 18. We assume households always have both a head and a spouse.

2.2. Income from transfers and wages. Let Y_t represent all period t after-tax income from transfers and wages, including labor income, asset income, inheritances, private defined benefit pensions, and all government transfers including social security. We assume labor is supplied inelastically, so Y_t is exogenous. We model $y_t = \ln(Y_t)$ during working life as the sum of a cubic polynomial in age, a Markov process u_t that approximates an underlying AR(1) process, and an iid normally distributed error term. During retirment, we model

y_t as the sum of a linear polynomial in age and an iid normally distributed error term. Retirement occurs exogenously at age 63. The income process and the retirement age are calibrated from the PSID.

2.3. Liquid assets and noncollateralized debt. Let X_t represent liquid asset holdings at the beginning of period t excluding income from wages and transfers. If $X_t < 0$ then credit card debt was held between $t-1$ and t. We introduce a credit limit at age t of 30% of average income at age t. This implies that $X_t \geq -(0.30) \cdot \bar{Y}_t$. This credit limit is calibrated from the 1995 SCF. Our model precludes consumers from simultaneously holding liquid assets and credit card debt, though such behavior has been documented among a small fraction of consumers by Gross and Souleles (2000) and Bertaut and Haliassos (2001).

We assume the gross real after-tax interest rate received from liquid assets is $R = 1.0375$, corresponding to the historical return on a portfolio of two-thirds stocks and one-third bonds, assuming an average tax rate of 25%. The gross real interest rate on credit-card loans is $R^{CC} = 1.1175$, two percentage points *below* the mean debt-weighted real interest rate reported by the Federal Reserve Board. This low value is chosen to implicitly capture the impact of bankruptcy and default, which lower consumers' effective interest payments.

2.4. Illiquid assets. Let Z_t represent illiquid asset holdings at the beginning of period t, and assume that $Z_t \geq 0, \forall t$. Illiquid assets in our model generate two types of returns: capital gains and consumption flows. We set the rate of capital gains equal to $R^z = 0$ and the annual consumption flows equal to $\gamma = .05$ times the value of the asset. Hence, the return on illiquid assets is considerably higher than the return on the liquid asset. For computational tractability, we adopt the assumption here of complete illiquidity of Z; transaction costs are sufficiently large that the asset can never be sold.[1]

Three issues arise when interpreting this illiquid asset. First, the specification is obviously quite stylized, but Z shares some similarities with home equity. Consider a consumer who owns a house of fixed real value H and derives annual consumption flows from the house of γH. Suppose the consumer has a mortgage of size M, and hence home equity of $H - M$. The real cost of the mortgage is ηM, where $\eta = i \cdot (1 - \tau) - \pi$ is the nominal mortgage interest rate corrected for the tax deductibility of mortgage interest payments and inflation. If we assume $\eta \approx \gamma$, the net benefit to the homeowner from the house is $\gamma H - \eta M. \approx \gamma(H - M) = \gamma Z$.

Second, the assumption of total illiquidity increases the motive for credit card borrowing in both the exponential and hyperbolic versions of our model. When illiquid assets have favorable returns but are highly illiquid, consumers will want to hold those illiquid assets and use credit card borrowing to smooth consumption volatility due to high frequency shocks in the income process.

Finally, for hyperbolic consumers even small delays between requests for liquidity (i.e., applications for home equity lines) and actual access to liquidity (i.e., approval of applications and release of funds) have a behavioral effect equivalent to total illiquidity (Laibson 1997a). In other words, even fairly small actual or perceived transactions costs involved in extracting liquidity from home equity will deter hyperbolic consumers from applying for home equity loans. Hence, the extreme illiquidity assumption is particularly appropriate for the hyperbolic model.

2.5. Dynamic and static budget constraints. Let I_t^X represent net investment into the liquid asset X during period t, and let I_t^Z represent net investment into the illiquid asset

[1] Angeletos et al. (2001) and LRT (2000) find that assuming partial illiquidity, ie. fixed and proportional costs of withdrawal from the Z asset, has similar implications to the full illiquidity assumption we adopt here.

Z during period t. Then the dynamic budget constraints are given by,

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

$$Z_{t+1} = R^Z \cdot (Z_t + I_t^Z) \tag{2}$$

where R^X and R^Z are the real interest rates, respectively, on liquid wealth and illiquid wealth. 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:

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

The static budget constraint is:

$$C_t = Y_t - I_t^X - I_t^Z$$

The state variables at the beginning of period t are liquid wealth $(X_t + Y_t)$, illiquid wealth (Z_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.

2.6. Preferences. We adopt iso-elastic instantaneous utility functions (i.e., constant relative risk aversion) and hyperbolic discount functions. 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

$\alpha, \gamma > 0$.[2]

Figure 1 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 $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots\}$. This discount function was first used to study intergenerational preferences (Phelps and Pollak 1968). Laibson (1997a) used Phelps and Pollak's discount function to capture the properties of intra-personal time preferences. Figure 1 plots the quasi-hyperbolic discount function for the case of $\beta = .7$ and $\delta = .956$.[3] 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 t and $t+1$ represents the weight placed on utils at time $t+1$ relative to the weight placed on utils at time t. From the perspective of self t, the discount factor between periods t and $t+1$ is $\beta\delta$, but the discount factor that applies between any two later periods is δ. Since we take β to be less than one, this implies a short-term discount 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

[2]See 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.

[3]The particular parameter values used in this example correspond to the calibration used in this paper for households with a high school educated head.

Figure 1 Discount Functions

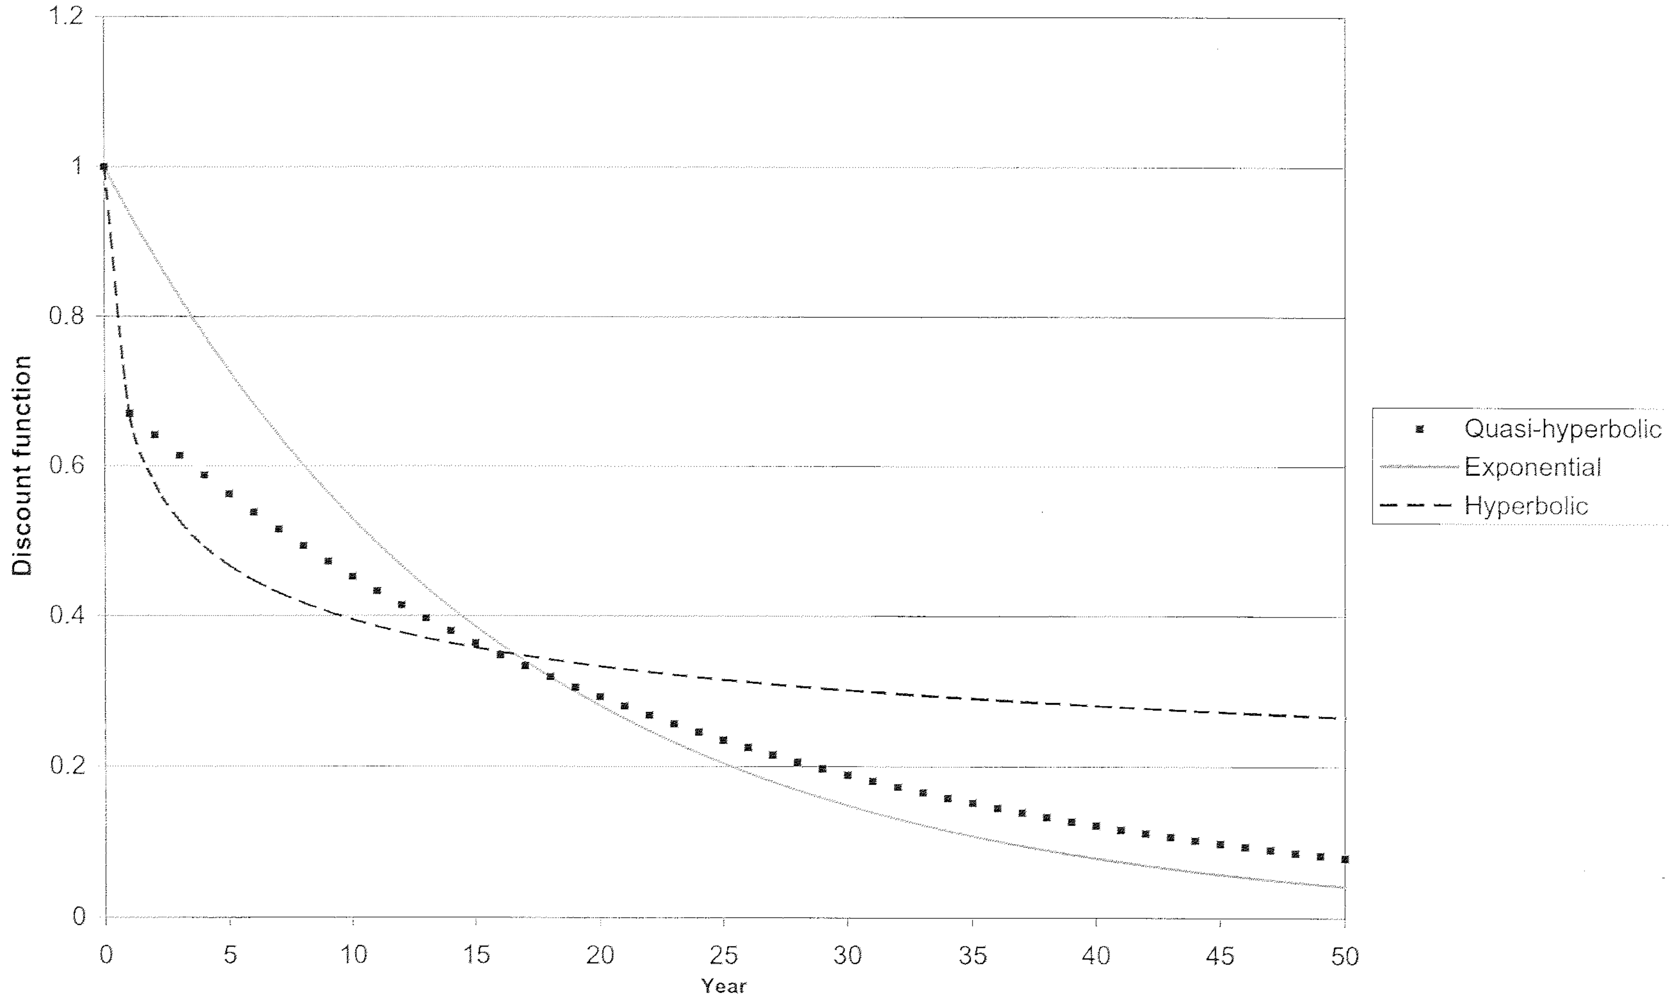

Source: Authors' calculations. Exponential: δ^t, with δ=0.939; hyperbolic: $(1+\alpha t)^{-\gamma/\alpha}$, with α=4 and γ=1; and quasi-hyperbolic: $\{1,\beta\delta,\beta\delta^2,\beta\delta^3,...\}$, with β=0.7 and δ=0.957.

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.

This potential for dynamic inconsistency in preferences implies that we must specify the preferences of all of the temporally distinct selves over the life-cycle: $t \in \{20, 21, ..., 90\}$. 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:

$$\beta \sum_{i=1}^{90-t} \delta^i \left(\Pi_{j=1}^{i-1} s_{t+i}\right) \left[s_{t+i} \cdot u(C_{t+i}, Z_{t+i}, n_{t+i}) + (1 - s_{t+i}) \cdot B(X_{t+i}, Z_{t+i}, D_{t+i})\right].$$

Note that n_t is the effective household size, ρ is the coefficient of relative risk aversion, γZ_t represents the consumption flow generated by 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 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 $t+i$.

2.7. Equilibrium. When $\beta < 1$ the household has dynamically inconsistent preferences, and hence the consumption problem cannot 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 $\{20, 21, ..., 90\}$ 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, $\{t,\ X+Y,\ Z,\ u\}$, 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 $V_{t,t+1}(X_{t+1}+Y_{t+1}, Z_{t+1}, u_{t+1})$ represent the time $t+1$ continuation payoff function of self t. Then the objective function of self t is:

$$u(C_t, Z_t, n_t) + \beta\delta E_t V_{t,t+1}(\Lambda_{t+1}) \tag{3}$$

where Λ_{t+1} represents the vector of state variables: $\{X_{t+1}+Y_{t+1}, Z_{t+1}, u_{t+1}\}$. Self t chooses C_t to maximize this expression. The sequence of continuation payoff functions is defined recursively:

$$V_{t-1,t}(\Lambda_t) = s_t[u(C_t, Z_t, n_t) + \delta E_t V_{t,t+1}(\Lambda_{t+1})] + (1-s_t)E_t B(\Lambda_t) \tag{4}$$

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 3 — reflecting self t's discount factor between periods t and $t+1$ — but does not appear in Equation 4, since self $t-1$ does not use the β factor to discount between periods t and $t+1$.

Equations 3 and 4 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.[4] 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.[5]

2.8. Simulation. Our simulations represent the behavior of individual households. We generate 5000 independent streams of income realizations according to the process described in Subsection 2.2. Households make equilibrium decisions at each age using the information available at that age (current and lagged income realizations). From the resulting simulated profiles of C, X, Z, and Y, we calculate the summary moments used in the MSM estimation procedure described below.

[4]See Laibson 1997b.

[5]While economists and psychologists have a great deal of evidence that consumers are not perfectly rational, that does not imply that we know what alternative to rationality should be adopted. There are no well-studied, generally applicable bounded rationality models. We are keenly interested in the recent developments in the reinforcement learning literature (e.g., Erev and Roth 1997, and Camerer and Ho 1996). However, reinforcement models like these are difficult to apply to the analysis of savings decisions, since it is not clear why and even if savings decisions are rewarding in the short-run. Perhaps the lack of short- and medium-run reinforcement provides another explanation for undersaving. We are also sympathetic to the model of "naif" behavior first proposed by Robert Strotz (1956) and more recently studied by Akerlof (1991), and O'Donoghue and Rabin (1997, 1998). 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.

3. DATA

Our MSM implementation formally compares the simulated values of five moments to their empirical analogues. Table 1 summarizes our empirical findings. The Table reports each empirical moment, its shorthand name, and its estimated standard error. We present a complete description of our data procedures in Appendix 1.

The first statistic, $wealth$, is the median wealth-to-income ratio for households with heads aged 50-59. The value of $wealth = 3.88$ implies that American households accumulate a substantial amount of wealth before retirement. Our wealth measure does not include involuntary wealth accumulation like social security and other types of defined benefit pensions.

The second statistic, $\%\ Visa$, is the fraction of households who borrow on credit cards. Our analysis indicates that 68% of US households pay interest on credit card debt each month. This number excludes households who pay their Visa bill in full at the end of the month, and excludes households who pay their Visa bill late.

We construct the third statistic in the table, $mean\ Visa$, by dividing age-specific credit card borrowing by mean age-specific income. We then average this fraction over the lifecycle. The average household has oustanding credit card debt equal to 13% of the mean income of its age cohort.

The fourth statistic, CY, represents the degree of anomalous comovement of consumption and income. We find that the excess sensitivity of consumption in response to predictable changes in income is 23% of the income innovation. This figure is consistent with other analyses. Hall and Mishkin (1982) report a statistically significant coefficient of .200, Hayashi (1985) reports a significant coefficient of .158, Altonji and Siow (1987) report an insignificant coefficient of .091, Attanasio and Weber (1993) report an insignificant coefficient of .119, Attanasio and Weber (1995) report an insignificant coefficient of .100, Shea (1995) reports a marginally significant coefficient of .888, Lusardi (1996) reports a significant coefficient of .368, and Souleles (1999) reports a significant coefficient of .344. See Deaton (1992) and Browning and Lusardi (1996) for discussion of the excess sensitivity literature.

Banks et al. (1998) and Bernheim et al. (1997) document a decline in consumption in the years around retirement that cannot be explained by standard effects like a reduction in work-related expenditures. Using the PSID we find an anomalous retirement consumption drop of $Cdrop = 9\%$.

Table 1

STATISTIC	$\bar{m}_d$	$se_{\bar{m}_d}$
median 50-59 $\frac{wealth}{income}$ ($wealth$)	3.880	.245
% borrowing on 'Visa'? ($\% \ Visa$)	.681	.015
mean (borrowing$_t$ / mean(income$_t$)) ($mean \ Visa$)	.126	.010
Consumption-Income comovement (CY)	.231	.112
retirement C drop ($Cdrop$)	.091	.067

These five moments reflect important empirical characteristics of lifecycle behavior including wealth accumulation, borrowing, and consumption-income comovement. Essential to the results is simultaneous inclusion in the set of moments of $wealth$ (or some other measure of long-run patience) and $\% \ Visa$, $mean \ Visa$, CY, $Cdrop$ or some other measure of short-run impatience. The analysis below demonstrates that the data reflects a simultaneous preference for patience and impatience. This tension ultimately drives econometric identification of the model's parameters.

4. METHOD OF SIMULATED MOMENTS PROCEDURE

The method of simulated moments (MSM) was developed by Ariel Pakes and David Pollard (1989) and Darrell Duffie and Kenneth Singleton (1993). We use MSM to formally compare the predictions of our model to the empirical evidence on lifecycle behavior described in Section 3. In addition, MSM enables us to formally test the nested null hypothesis of exponential discounting $\beta = 1$. The current Section, which describes the MSM procedure, is self-contained and may be skipped on a first reading.

In order to describe the MSM procedure we adopt the following notation. Let θ be a vector of preference parameters. Suppose θ^* is the true vector. Let $\bar{m}_d$ be a vector of N (fixed) empirical moments. Suppose $\bar{m}_d$ has asymptotic variance-covariance matrix Ω. Let $m_s = m_s(\theta)$ be a vector of N simulated moments which correspond to the empirical moments $\bar{m}_d$. Note that the simulated moments exhibit a functional dependence on the vector of preference parameters. Let $e(\theta) \equiv \bar{m}_d - m_s(\theta)$ be the deviation of the simulated moments from the empirical moments given θ. Let W be an NxN weighting matrix. Let $q(\theta) \equiv e(\theta) \cdot W^{-1} \cdot e(\theta)'$ be a scalar-valued loss function, equal to the weighted sum of squared deviations of simulated moments from their corresponding empirical values.

Our procedure is to minimize the loss function $q(\theta)$,[6] and define the MSM estimator as,

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}}\, q(\theta). \tag{5}$$

[6]We perform this minimization with Matlab's built-in Nelder-Mead simplex algorithm. This algorithm is slower but more robust than derivative-based methods, and here it is preferred because of the nonconvexities in hyperbolic consumption functions discussed in Subsection 2.7.

Pakes and Pollard (1989) demonstrate that $\hat{\theta}$ is a consistent estimator of θ^*, and $\hat{\theta}$ is asymptotically normally distributed. For $W = \Omega$, $\hat{\theta} \longrightarrow N(\theta^*, \Sigma)$ asymptotically,[7] where

$$\Sigma = \left[\left(\frac{\partial e_i(\hat{\theta})}{\partial \theta_j} \right)' \Omega^{-1} \left(\frac{\partial e_i(\hat{\theta})}{\partial \theta_j} \right) \right]^{-1} \tag{6}$$

Note that the derivatives in this expression are evaluated at $\hat{\theta}$. The intuition for this equation is most easily seen by analogy to the case of estimation of one parameter by one moment with the familiar method of moments. First, observe that the standard error of the parameter estimate is increasing in the standard error of the empirical moment. If the moment is imprecisely estimated, we can attach little certainty to the parameter estimate. Also, suppose in this simple one parameter, one moment case that $m_s(\theta)$ is very flat near the optimum. Then large changes in the parameter would have only a small effect on the simulated moment. Consequently, the location of the true minimum of the loss function will be relatively uncertain. Conversely, if $m_s(\theta)$ is steeply sloped very close to the optimum, so small changes in the parameter have a dramatic effect on the moment, one can be relatively confident about the parameter estimate: the standard error of $\hat{\theta}$ will be small.

Returning to the general case of MSM, the expression for Σ above likewise indicates that large derivatives of the moments as functions of the parameters result in small standard errors for the estimated parameters. In other words, if the moments are very sensitive to the parameters, the parameters are more likely to be precisely estimated. The Ω^{-1} term in the center captures the notion that redundant or imprecisely estimated empirical moments in general do not tightly constrain the MSM parameter estimates.[8]

[7] Technically the *wealth* moment is not differentiable. This has no bearing on the consistency of the estimator, but we are currently evaluating procedures, like those in French and Jones (2001) and Newey and McFadden (1994), that properly implement MSM and ensure asymptotic normality when using quantile moments.

[8] Subsequent versions of this paper will contain an appendix with detailed, technical notes about our

5. RESULTS

In this section we estimate the discounting parameters β and δ. We also consider constrained estimates in which we restrict the value of β. In our benchmark estimates we use all of the moments described above, but we also consider alternative identification strategies that use subsets of our moments. We check robustness and compare our results to those of LRT (2001). Throughout this analysis we fix the coefficient of relative risk aversion at $\rho = 2$. Finally, we focus on the case $W = \Omega$, which yields a more efficient estimator than the alternatives $W = I_N$ and $W = diag(\Omega)$.

5.1. Exponential Case. We begin our analysis by estimating δ while constraining $\beta = 1$ (the exponential discounting case). We obtain

$$[\hat{\delta} : \beta = 1] = .833$$
$$(.003)$$

The credit card borrowing moments, *% Visa* and *mean Visa*, drive this estimate of the exponential discount function. This exponential model matches *wealth* very poorly, but *wealth* is relatively less tightly estimated than the credit card moments. Recall that the empirical value is *wealth* = 3.88, with a standard error of 0.25. The loss function minimized by the MSM procedure weights the moments by the inverse of the variance matrix, so less tightly estimated moments have less of an effect on the parameter estimates. Intuitively, the low variances of *% Visa* and *mean Visa* cause those moments to identify $\hat{\delta}$. Column 1 of Table 2 reports the values of the moments at the estimated $\hat{\delta}$. The simulated credit

procedure, including formal integration of the "first stage" estimation of income process parameters and discussion of the (small) additional errors introduced by substituting simulated moments derived from finite samples of simulated data.

card moments are close in absolute terms to their empirical analogues, which are reproduced from Table 1 in Column 2 of Table 2 (Column 3 reproduces the empirical standard errors). The consumption-income comovement statistics, CY and $Cdrop$, are well-matched in this exponential simulation as well. However, *wealth* evaluated at $\hat{\delta}$ is approximately zero. This is a serious shortcoming of the model, given that the actual value of *wealth* is 3.88.[9]

Table 2

STATISTIC	Column 1 $m_s(\hat{\delta}; \beta = 1)$	Column 2 $\bar{m}_d$	Column 3 $se_{\bar{m}_d}$
wealth	-0.03	3.88	0.25
$\%\ Visa$	0.66	0.68	0.015
$Visa/\bar{Y}$	0.15	0.13	0.01
CY	0.27	0.23	0.11
$C\ drop$	0.16	0.09	0.07
$q(\hat{\theta})$	278		

Intuition for the particular value of $\hat{\delta}$ comes from the exponential Euler Equation (LRT 2001). Consider a stripped-down version of the exponential model. Specifically, 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:

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

[9]We'll formally characterize the degree of seriousness through a specification test, described at the end of the results section.

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

$$\begin{aligned} \text{discount rate} &= -\ln(\delta) \\ &= -\rho E_t \Delta \ln (C_{t+1}) + r + \tfrac{\rho^2}{2} V_t [\Delta \ln (C_{t+1})] \end{aligned} \tag{7}$$

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 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.[10] 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.[11] Hence, for our analysis, the appropriate real interest rate is the real credit card borrowing rate, $r = r^{cc} \approx 0.1175$. Recall that this four percentage points below the debt-weighted credit card interest rate. The difference reflects adjustments for inflation and default.

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 vari-

[10] E.g., Gourinchas and Parker (2001).
[11] SCF, 1995 cross-section.

ance 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 0.025. 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 borrowing in the credit card market, a portfolio decision that suggests low levels of liquid wealth accumulation and hence necessarily imperfect consumption smoothing.

We are now in a position to evaluate $-\ln(\delta)$.

$$\begin{aligned} -\ln(\delta) &= -2 \cdot 0 + 0.1175 + \frac{2^2}{2}(0.025) \\ &= 0.1675 \end{aligned}$$

This back-of-the-envelope Euler Equation calculation compares well with the estimate implied by our MSM procedure:

$$-\ln \delta = -\ln 0.833 = 0.1827$$

Such high long-term discount rates are problematic. As we will show below, observed household wealth accumulation profiles can only be explained with much *lower* discount rates. 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% ($\rho = 3$).[12] Hence, these observations suggest a puzzle. Consumers act impatiently in the credit market but act patiently when accumulating for retirement. LRT (2001) call this the Debt Puzzle.

5.2. Hyperbolic Case. Next we explore whether relaxing the $\beta = 1$ constraint improves the model's ability to match the data. Theory suggests that it may. The Euler Equation analysis above only applies to exponential consumers ($\beta = 1$). As Harris and Laibson (2001) have shown, making the discount function hyperbolic generates a critical modification of the Euler Equation. They find the following form for the Hyperbolic Euler Equation:

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

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 $\beta\delta$.

[12] All 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.

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 ($\simeq -\ln \beta\delta$), 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

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

Hence, illiquid wealth accumulation is primarily driven by δ, not β, implying that the consumer accumulates illiquid wealth as if she had a discount rate of $-\ln(\delta) = .05$.

With the potential for effective discount rates of $-\ln \beta\delta$ 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.

With these theoretical observations in mind, we now ask whether relaxing the $\beta = 1$ constraint can resolve the Debt Puzzle and explain the available evidence on consumption-income comovement. First, we fix $\beta = .7$ and estimate δ. In this case we find

$$\boxed{\begin{array}{r}[\hat{\delta} : \beta = .7] = .907 \\ (.001)\end{array}}$$

This $\hat{\delta}$ corresponds to a discount rate of $-\log(\hat{\delta}) = .097$, which is still problematically large. The values of the simulated moment functions evaluted at $\hat{\delta}$ are reported in Column 2 of Table 3 (Column 1 repeats column 1 of Table 2). As in the $\beta = 1$ constrained case, the small standard errors on the credit card moments mean that they are matched quite closely in absolute terms. Again, the simulated comovement measures also resemble the empirical values. However, the wealth measure is still just $wealth(\hat{\delta} : \beta = .7) = 0.21$, fourteen standard deviations from the empirical value.

Table 3

STATISTIC	Column 1 $m_s(\hat{\delta}; \beta = 1)$	Column 2 $m_s(\hat{\delta}; \beta = .7)$	Column 3 $m_s(\hat{\delta}, \hat{\beta})$	Column 4 $\bar{m}_d$	Column 5 $se_{\bar{m}_d}$
wealth	-0.03	0.21	3.05	3.88	0.25
% Visa	0.66	0.64	0.62	0.68	0.015
$Visa/\bar{Y}$	0.15	0.17	0.17	0.13	0.01
CY	0.27	0.37	0.48	0.23	0.11
C drop	0.16	0.12	0.26	0.09	0.07
$q(\hat{\theta})$	278	271	86		

Next, we report the two parameter estimates, when β is not fixed. For this general case, our benchmark estimate is

$\hat{\delta} = .969$	$\hat{\beta} = .195$
$(.003)$	$(.009)$

The long run discount factor δ is estimated to be 0.969. Consumers with such a high degree of patience over long horizons accumulate large stocks of wealth, for example in preparation for retirment. The model drives $\hat{\delta}$ to this high value in order to match the *wealth* moment. To generate this level of wealth accumulation the long-run discount rate, $-\ln\delta$, must be approximately equal to the rate of return on savings (i.e., 3.75%-5.00%). In fact, as reported in Column 3 of Table 3, this high value of $\hat{\delta}$ results in substantial simulated wealth accumulation: $wealth(\hat{\delta}, \hat{\beta}) = 3.05$, less than four standard errors from the empirical value.

Most of the wealth these simulated consumers accumulate is in illiquid form. Illiquid assets are attractive to them for two reasons. First, illiquid assets have a higher rate of return in our model than liquid assets, so these very patient consumers wish to take advantage of it over long horizons. Second, as described in the next paragraph, these consumers have a very low value of β. Illiquidity helps them overcome their dynamic inconsistency problems.

The estimate of β is close to $\hat{\beta} = 0.2$ and very precise, with a standard error of 0.01. We conclude that the MSM procedure rejects the $\beta = 1$ null hypothesis of exponential discounting with very high levels of confidence; the t-statistic is $\frac{1-\hat{\beta}}{s.e._{\hat{\beta}}} = 89$.

The credit card borrowing and consumption-income comovement moments drive β below unity. Such low values of β allow the model to better replicate this short term impatient behavior implied by the costly credit card borrowing and consumption-income comovement. However, tension remains because of the high value of $\hat{\delta}$: $\%\ Visa$ is poorly matched because

$\hat{\delta}$ is so high, and the simulated value of *wealth* is still far too low because $\hat{\beta}$ is so small. Even the relatively flexible model of hyperbolic discounting cannot reconcile such large amounts of simultaneous wealth accumulation, credit card borrowing, and consumption-income comovement.

Formally we know that if the model is correct then $q(\hat{\theta})$ has a χ^2 distribution. Recall that $q(\hat{\theta})$ is the sum of squared deviations of the simulated moments (evaluated at the parameter estimates) from the empirical moments, weighted by the inverse of the empirical variances, The number of degrees of freedom of the χ^2 equals the number of moments used for the estimation minus the number of parameters estimated. Thus in the one-parameter constrained cases, when we fixed β first at 1 and then at .7, $q(\hat{\theta}) \sim \chi^2(4)$, while in the unconstrained case when we estimated both β and δ, $q(\hat{\theta}) \sim \chi^2(4)$. The values of $q(\hat{\theta})$ for these different specifications are reported in the bottom row of Table 3. By comparison, the 99% critical values of the $\chi^2(3)$ and the $\chi^2(4)$ are 11.3 and 13.3, respectively, which are *much* smaller than the actual $q(\hat{\theta})$'s. Thus we overwhelmingly reject both the exponential and hyperbolic models; neither can adequately reconcile the high empirical levels of wealth accumulation, credit card borrowing, and consumption-income comovement.

These results are very robust to changes in the specification. For example, we have replicated the estimation with different sets of moments. When we try to match only *wealth* and *% Visa* with the data from LRT 2001 we obtain almost identical parameter estimates:

$\hat{\delta} = .970$	$\hat{\beta} = .172$
(.002)	(.014)

However, we fail to match the moments. See Column 1 of Table 4. Failure to perfectly identify this two-moment, two-parameter case is additional evidence that a Debt Puzzle is present (LRT 2001).

Modifying the model's interest rates caused changes in the parameter estimates, but did not lead to acceptance of the specification test. Specifically, we explored the case where credit card borrowing and illiquid asset accumulation were both more attractive. Setting $R = 1.02, R^{CC} = 1.09$, and $\gamma = .07$ resulted in estimates of $[\hat{\delta}, \hat{\beta}] = [.946, .481]$, but the same tension persisted between the wealth accumulation and credit card borrowing moments. We summarize this run in Column 2 of Table 4.

Table 4

STATISTIC	Column 1 $m_s(\hat{\delta}, \hat{\beta})$	Column 2 $m_s(\hat{\delta}, \hat{\beta})$	Column 4 $\bar{m}_d$	Column 5 $se_{\bar{m}_d}$
wealth	2.666	2.56	3.88	0.25
% *Visa*	0.681	0.63	0.68	0.015
$Visa/\bar{Y}$	-	0.17	0.13	0.01
CY	-	0.37	0.23	0.11
C *drop*	-	0.15	0.09	0.07
$q(\hat{\theta})$	6.16	89		

Column 1: Use only the moments $wealth = 3.24$ and % $Visa = 0.70$.

Column2: All five moments, but $R = 1.02, R^{CC} = 1.09$, and $\gamma = .07$

6. CONCLUSION

To be written.

7. APPENDIX 1: DATA

Our discussion here of the data procedures first addresses the moments derived from the Survey of Consumer Finances, *wealth*, % *Visa*, and *mean* $Visa_t/\bar{Y}_t$, and then turns to the moments derived from the Panel Study of Income Dynamics, *CY* and *Cdrop*.

SCF Moments. We use the 1989, 1992, 1995, and 1998 SCF's to compute the *wealth* moment. We derive % *Visa* and *mean* $Visa_t/\bar{Y}_t$ from the 1995 and 1998 SCF's. For all three moments, we control for cohort effects, household demographic characteristics, and business cycle effects; we undertake this task to make the characteristics of the populations underlying the real-world data and the simulated data fully analogous. We assign to households in our model – the simulated households – the mean cohort, demographic, and business cycle effects, so by controlling for these effects in the SCF data we align the characteristics of the empirical and simulated populations. The exact procedure is this:

For each variable of interest x first use weighted least squares, where the weights are the SCF population weights, to estimate the regression

$$x_i = FE_i + BCE_t + CE_i + AE_i + \xi_i \quad (8)$$

Here FE_i is a family size effect that consists of three variables, the number of heads, the number of children, and the number of dependent adults in the household. BCE_i is a business cycle effect proxied by the unemployment rate in the household's region of residence. In 1992, 1995, and 1998, the unemployment rate is the rate in the household's Census Division. In 1989 the nationwide unemployment rate was used because information on household location is not available in that year's public use data set. CE_i is a cohort effect that consists of a full set of five-year cohort dummies, AE_i is an age effect that consists of a full set of age dummies, and ξ_i is an error term.

Next, with these effects, perform the correction. Define the "typical" household to be identical to the household modelled and used in the simulations (i.e. possessing two heads, exogenous age-varying numbers of children and adult dependents, an average cohort effect, and an average unemployment effect, [using the same means as in the estimation of the income process]). Then for each variable create a new variable $\widehat{x_i}$ (corrected x_i) that captures what x_i would have been had household i been typical. For example, if household i is identical to the "typical" household except for having more children, we set $\widehat{x_i} = x_i + \beta(\overline{nkids} - nkids_i)$, where β is the coefficient for number of kids in the regression above and $\overline{nkids}$ is the average number of children in the household as a function of the head's age. All moments were subsequently estimated using $\widehat{x_i}$.

For the *wealth* moment, we restrict the sample to households with heads aged between 50 and 59, and we calculate $\widehat{x_i}$ = wealth/income corrected for cohort, family and business cycle effects for each household in the restricted sample. *wealth* is the median of $\widehat{x_i}$ in the sample, using the SCF population weights. We use a bootstrap technique to find the standard error of *wealth*. Specifically, from the actual restricted, corrected sample of $\widehat{x_i}$ we draw with replacement to generate 5000 random samples of a size comparable to the actual sample. Each of the random bootstrap samples is formed by first taking age-by-age random samples, of sizes such that aggregation replicates the distribution of ages in the original SCF sample. The standard error of *wealth* is the standard deviation of the medians of the 5000 random bootstrap samples.

We begin constructing $\%$ *Visa* by creating a dummy variable *hasdebt* that is equal to one for household i if i has a positive outstanding credit card balance in the SCF. We correct *hasdebt* for cohort, demographic, and business cycle effects, generating $\widehat{x_i}$. We then regress

$\widehat{x_i}$ on a full set of age dummies. *% Visa* is a linear combination of the estimated coefficients on the age dummies, where the weights are those implicit in the SCF mortality weights. The standard error is computed directly from the weights and the standard errors on the age dummy estimates.

Construction of $mean\ Visa_t/\bar{Y}_t$ is confounded by the fact that aggregate credit card borrowing data from the Federal Reserve Board indicate that 1995 and 1998 SCF borrowing magnitudes are biased downward by a factor of three. We correct for this bias as follows. First we compute average outstanding interest-bearing balances. According to the Fed, aggregate debt outstanding at year-end 1995 and 1998 were \$443 billion and \$561 billion, respectively. From these figures we subtract an upper bound on the float (the balances that are still in their grace period which do not accrue interest), which we obtain by dividing total purchase volume by 12. We then divide by the number of US households with credit cards, using Census Bureau data on total households and SCF data on the percentage of households with cards. We obtain average household borrowing conditional on having a card of \$5115 in 1995 and \$6411 in 1998. These figures are consistent with those from a proprietary account-level data set analyzed by David Gross and Nicholas Souleles (1999a, 1999b, 2000).

In our simulations we focus on households headed by people with high school degrees, so next we use the SCF data on borrowing to scale the Fed average borrowing figure for just the high school educated group (ie, we try to remove any bias caused by heterogeneity between households without high school degrees, with just high school degrees, and with college degrees . In particular, define α such that

$$debt_{all}^{Fed} = \alpha \cdot (w_{nhs} debt_{nhs}^{SCF} + w_{hs} debt_{hs}^{SCF} + w_{coll} debt_{coll}^{SCF})$$

with weights w_{nhs}, w_{hs}, and w_{coll} defined by the proportion of educational categories in the population (0.25, 0.5, 0.25, respectively) and $debt_{educ}^{source}$ equal to the average debt reported by *source* for educational group *educ*. Focussing now exclusively on the high school educational group, let $debt_i^{SCF}$ be the level of credit card debt reported in the SCF for household i. Let $debt_i = \alpha \cdot debt_i^{SCF}$ be the corrected credit card debt. Calculate age specific income means $(\overline{y}_t)$ and create $debtinc_i$ as $debt_i/\overline{y}$. (When calculating the age-specific income means we group together ages 20-21, 70-74, 75-79, and 80+ because we have very few observations at those ages.) Then, as with *wealth* and $\%\ Visa$, we correct $debtinc_i$ for cohort, demographic, and business cycle effects, creating $\widehat{x_i}$, and regress $\widehat{x_i}$ on a full set of age dummies. The moment $mean\ Visa_t/\bar{Y}_t$ is a linear combination of the estimated coefficients on the age dummies, using the weights implicit in the SCF mortality rates. As for $\%\ Visa$, the standard error is computed directly from the weights and the standard errors on the age dummy estimates.

Covariances between the SCF moments were constructed by bootstrapping. We generated 5000 age-by-age bootstrap samples from the SCF (i.e. 5000 samples of 20 year old households, 5000 samples of 21 year old households, and so on), and the 3 SCF moments were each estimated using each of the 5000 replications. The covariance matrix we use is the covariance of the 3 moments across the 5000 bootstrap samples.

7.1. PSID Moments. We use PSID data from 1978 to 1992 to estimate the CY and $Cdrop$ moments. In the data, we define consumption to include food, rent, and utilities (the most general definition available in the PSID). The rental value of an owner-occupied home is assumed to be 5% of the value of the home. If the household neither owns nor rents, rent is the self-reported rental value of the home if it were rented. Total consumption was projected from the PSID's partial measure using the CEX: in the CEX we regress total consumption on food, rent and utilities consumption, and then we use the coefficient estimates to infer total consumption from the available PSID measure. Since wealth is observed in the PSID only in 1984 and 1989, in the other years we estimate wealth using the intertemporal budget constraint. The after-tax real interest rate is assumed to be 3.75%.

We construct the CY moment by using 2SLS to estimate

$$\Delta \ln(C_{it}) = \alpha E_{t-1} \Delta \ln(Y_{it}) + X_{it}\beta + \varepsilon_{it}.$$

We assume an MA(1) process for the error term and instrument for $E_{t-1}\Delta\ln(Y_{it})$ with $\ln Y_{it-3}$ and $\ln Y_{it-4}$. The overidentification test does not reject this specification. The vector X_{it} includes age, cohort, and business cycle effects, the change in effective family size, the mortality rate, and lagged wealth.

Cdrop is the sum of the elements of $\hat{\gamma}$ derived from estimation of

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

where I_{it}^{RETIRE} is a set of dummy variables that take the value of one in periods $t-1$, t, $t+1$ and $t+2$ if period t is the age of retirement. The vector of control variables X_{it} includes mortality rates, the change in effective household size, cohort effects, and business cycle effects.

We obtained the covariance between CY and *Cdrop* by jointly estimating the above regressions.

8. REFERENCES

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

328.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 2001. "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation" *Journal of Economic Perspectives*. Vol. 15, No. 3, Summer 2001, pp. 47-68.

———, ———, ———, ———, and ———. 2001. "The Hyperbolic Buffer Stock Model: Calibration, Simulation, and Empirical Evaluation" NBER working paper.

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

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

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

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

Bernheim, B. Douglas. 1995. "Do Households Appreciate Their Financial Vulnerabilities? An Analysis of Actions, Perceptions, and Public Policy." *Tax Policy for Economic Growth in the 1990s*. Washington, D.C.: American Council for Capital Formation, pp.1-30.

———, Jonathan Skinner and Steven Weinberg. 1997. "What Accounts for the Variation in

Retirement Wealth Among US Households?" Working Paper 6227. Cambridge, Mass: National Bureau of Economic Research.

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

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

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

Campbell, John Y. and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence," in Olivier J. Blanchard and Stanley Fischer (eds.) *NBER Macroeconomics Annual* 1989. Cambridge, Mass: MIT Press, pp. 185-216.

Carrillo, Juan and Thomas Mariotti. 2000. "Strategic Ignorance as a Self-Disciplining Device," forthcoming in *Review of Economic Studies*.

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

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

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

Coase, Ronald H. 1972. "Durability and Monopoly." *Journal of Law & Economics.* Vol 15 (1), p. 143-49.

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

———. 1992. *Understanding Consumption.* Oxford University Press: Oxford.

——— and John Muellbauer 1986. "On Measuring Child Costs: With Applications to Poor Countries." *Journal of Political Economy.* 94:4, pp. 720-744.

DellaVigna, Stefano, and Paserman, M. Daniele. 2000. "Job Search and Hyperbolic Discounting." Maurice Falk Institute for Economic Research (Jerusalem, Israel), Discussion Paper No. 00.15.

——— and Ulrike Malmendier. 2001. "Self-control in the market: evidence from the health club industry." Harvard University mimeo.

Diamond, Peter and Botond Koszegi. 1998. "Hyperbolic Discounting and Retirement." MIT mimeo.

Duffie, Darrell, and Kenneth J. Singleton. 1993. "Simulated Moments Estimation of Markov Models of Asset Prices." *Econometrica.* 61:4, pp. 929-952.

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

Farkas, Steve and Jean Johnson. 1997. "Miles to Go: A Status Report on Americans' Plans for Retirement." Public Agenda.

Frederick, Shane, George Loewenstein, and Edward O'Donoghue (2001) "Time Discounting: A Critical Review," forthcoming in *Choice Over Time II*, New York: Russell Sage Foundation.

French, Eric and John Jones. 2001. "The Effects of Health Insurance and Self-Insurance on Retirement Behavior." Mimeo.

Gourinchas, Pierre-Olivier and Jonathan Parker. 2001. "Consumption Over the Life-Cycle," forthcoming in *Econometrica*.

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.

Hall, Robert E. 1978. "Stochastic Implications of the Life-Cycle/Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*. 96, pp. 971-87.

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

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

______ and ______. 2001 "Instantaneous Gratification," mimeo Harvard University.

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

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

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

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

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

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

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

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

———, Andrea Repetto and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement." *Brookings Papers on Economic Activity.* 1, pp. 91-196.

———, ——— and ———. 2001."A Debt Puzzle," in eds. Philippe Aghion, Roman Frydman, Joseph Stiglitz, Michael Woodford, *Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps*, forthcoming.

Loewenstein, George, and Drazen Prelec. 1992. "Anomalies in Intertemporal Choice: Evidence and an Interpretation." *Quarterly Journal of Economics*. 107, pp. 573-98.

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

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

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

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

Newey, Whitney, and Daniel McFadden. 1994. in R Engle and D. McFadden, *Handbook of Econometrics*.

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

——— and ——— 1999b. "Incentives for Procrastinators." *Quarterly Journal of Economics*. 114:3, pp. 769-816.

——— and ——— 2001. "Choice and Procrastination." *Quarterly Journal of Economics.* 116:1, pp.121-160..

Pakes, Ariel and David Pollard. 1989. "Simulation and the Asymptotics of Optimization Estimators." *Econometrica.* 57:5, pp 1027-57.

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

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

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

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

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

Stern, Steven, 1997. "Simulation-Based Estimation." *Journal of Economic Literature.* 35:4, pp. 2006-2039.

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

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

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