

THE HYPERBOLIC CONSUMPTION MODEL:
CALIBRATION, SIMULATION, AND EMPIRICAL EVALUATION.

GEORGE-MARIOS ANGELETOS, DAVID LAIBSON,
ANDREA REPETTO, JEREMY TOBACMAN, STEPHEN WEINBERG

Forthcoming *Journal of Economic Perspectives* (August 2001)*

ABSTRACT. Laboratory and field studies of time preference find that discount rates are much greater in the short-run than in the long-run. Hyperbolic discount functions capture this property. This paper presents simulations of the savings and asset allocation choices of households with hyperbolic preferences. The behavior of the hyperbolic households is compared to the behavior of exponential households. The hyperbolic households hold relatively more illiquid wealth and relatively less liquid wealth. The hyperbolic households borrow much more frequently in the revolving credit market. The hyperbolic households exhibit greater consumption-income comovement and experience a greater drop in consumption around retirement. The hyperbolic simulations match observed consumption and balance sheet data much better than the exponential simulations.

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

*George-Marios Angeletos is Assistant Professor of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts. David Laibson is Paul Sack Associate Professor of Political Economy, Harvard University, 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. Stephen Weinberg is a graduate student in economics, Harvard University, Cambridge, Massachusetts. We received helpful advice from David Altig, Orazio Attanasio, Chris Carroll, Christopher Harris, Greg Mankiw, Jonathan Parker and seminar audiences at Harvard University, University of California San Diego, University of Chicago, and the National Bureau Economic Research. We are particularly grateful for the thoughtful and detailed editorial guidance of Brad De Long, Timothy Taylor, and Michael Waldman. Eddie Nikolova, and Laura Serban provided excellent research assistance. 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; and Repetto, from DID-Universidad de Chile, FONDECYT (1990004), Fundación Andes and an institutional grant to CEA from the Hewlett Foundation.

1. INTRODUCTION

Our preferences for the long-run tend to conflict with our short-run behavior. When planning for the long-run, we *intend* to meet our deadlines, exercise regularly, and eat healthfully. But in the short-run we have little interest in revising manuscripts, jogging on the stairmaster, and skipping the chocolate soufflé à la mode. Delay of gratification is a nice long-term goal, but instant gratification is disconcertingly tempting.

This gap between long-run intentions and short-run actions is apparent across a wide range of behaviors, including saving choices. A 1997 survey found that three quarters of respondents (76 percent) believe they should be saving more for retirement. Looking only at respondents who believed they were at an age where “you should be seriously saving already,” the survey found that 55 percent reported being “Behind” in their savings and only 6 percent reported being “Ahead” (Farkas and Johnson, 1997). The report on the survey concluded: “[T]he gaps between people’s attitudes, intentions, and behavior are troubling and threaten increased insecurity and dissatisfaction for people when they retire. Americans are simply not doing what logic — and their own reasoning — suggests they should be doing.”

A 1993 Luntz Webber/Merrill Lynch survey found similar answers when baby boomers were asked “What percentage of your annual household income do you think you should save for retirement?” and “What percentage of your annual household income are you now saving for retirement?” The median reported shortfall was 10 percent of income and the mean gap was 11.1 percent (Bernheim, 1995).

Such survey evidence resonates with popular and professional financial planning advice. Financial planners seem to recognize self-control limits when they advise consumers to “use whatever means possible to remove a set amount of money from your bank account each month before you have a chance to spend it” (Rankin, 1993).

These observations suggest that households have self-control problems, but this conclusion is tempered by doubts about the quality of the evidence. Most of the data described above is either

anecdotal or based on attitudinal survey questions. By contrast, almost all mainstream economic analysis focuses on consumer *behavior*, not consumer *attitudes* about ambiguous concepts like what people believe they should be doing.

In this paper, we report new behavioral evidence that self-control problems importantly influence savings choices (see also Thaler and Shefrin, 1981, and Thaler, 1992). We adopt a conceptual framework that integrates a standard economic theory of life-cycle planning with a psychological model of self-control. This integrated model provides a parsimonious formal framework in which to evaluate the quantitative effects of self-control problems.

We build our modelling framework on three principles. First, our model adopts the standard assumptions of modern consumption models, including the ideas that consumers have uncertain future labor income and face liquidity constraints in the sense that they have limited ability to borrow against this future labor income (Carroll, 1992, 1997, this issue; Deaton, 1991; Zeldes, 1989b). Second, we extend the consumption literature by allowing our simulated consumers to borrow on credit cards and by including a partially illiquid asset in the consumers' menu of investment options. Third, we assume that consumers have both a short-run preference for instantaneous gratification and a long-run preference to act patiently. This combination of time preferences is usually called "hyperbolic" discounting (Ainslie, 1992), since generalized hyperbolas were first used to capture such intertemporal preferences (Chung and Herrnstein 1961).

Hyperbolic discounting generates the self-control problem that motivates our analysis. When two rewards are both far away in time, decision-makers act relatively patiently. For example, a worker prefers a twenty minute break in 101 days, rather than a fifteen minute break in 100 days. But when both rewards are brought forward in time, preferences exhibit a reversal, reflecting more impatience; the same person prefers a fifteen minute break right now, rather than a twenty minute break tomorrow. Far in advance, the consumer prefers to be patient between t and $t + 1$, waiting the extra day for the extra five minutes of break time. But, when date t actually arrives, the consumer's preferences have switched, and the consumer now prefers to act impatiently, taking the

shorter break immediately.

This type of preference ‘change’, or dynamic inconsistency, is reflected in many common experiences. Such reversals should be well understood by everyone who willfully sets the alarm clock the night before, only to repeatedly press the snooze button the morning after — and even better understood by anyone who goes out of their way to put their alarm clock on the other side of the room. Hyperbolic consumers will report a gap between their long-run goals and their short-run behavior. They will not achieve their desired level of “target savings,” since short-run preferences for instantaneous gratification undermine the consumers’ efforts to implement patient long-run plans.

The hyperbolic discounting model generates numerous empirical predictions that distinguish the model from the standard model with exponential discounting. First, households with hyperbolic discount functions will hold their wealth in an illiquid form, since such illiquid assets are protected from consumption splurges. Second, households with hyperbolic discount functions are very likely to borrow on their credit cards to fund instant gratification. Thus households with hyperbolic discount functions are likely to have a high level of revolving debt, despite the high cost of credit card borrowing. Third, since hyperbolic households have little liquid wealth, they are unable to smooth consumption, generating a high level of comovement between income and consumption. Indeed, hyperbolic households will even exhibit a high level of comovement between predictable changes in income and changes in consumption. Fourth, this comovement between income and consumption will stand out around retirement, when labor income falls and the lack of liquid wealth generates a necessary fall in consumption.

These rich empirical predictions enable us to distinguish between the exponential discounting model and the hyperbolic discounting model. In this paper we calibrate the exponential and hyperbolic models so that both models match available evidence on the level of retirement savings. We then show that the hyperbolic discounting model better matches available consumption and asset allocation data from the Panel Survey on Income Dynamics and the Survey of Consumer Finances.

2. HYPERBOLIC DISCOUNTING

Robert Strotz (1956) first suggested that people are more impatient when they make short-run tradeoffs than when they make long-run tradeoffs. Since then dozens of formal experiments have shown that decision makers are more impatient in the short-run than in the long-run. Time preference 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.¹ Most of these experiments elicit the values of rewards at different time horizons. Then, experimenters deduce the shape of the “discount function,” which measures the value of utility, as perceived from the present, at each future time period.

Figure 1 plots examples of discount functions, which are normalized to take a value of one at zero delay. The horizontal axis represents the duration of delay, τ , and the vertical axis represents the value of a util delayed by τ periods. The downward slope of the discount functions implies that delaying a reward reduces its value.

Economists usually assume that discount functions are exponential. Specifically, a util delayed τ periods is worth δ^τ as much as a util enjoyed immediately ($\tau = 0$). Typically, economists assume that δ is less than one, capturing the fact that future utils are worth less than current utils. In Figure 1, δ is set at 0.944, which is the annual discount factor used in our exponential simulations below. In this exponential case, the discount function declines at a constant rate over time — 5.6 percent per year.²

However, the experimental evidence implies that the actual discount function declines at a greater rate in the short-run than in the long-run. In other words, delaying a short-run reward by a few days reduces the value of the reward more in percentage terms than delaying a long-run reward by a few days. When researchers estimate the shape of the discount function based on choices by

¹See Solnick et al 1980, Thaler 1981, Navarick 1982, Millar and Navarick 1984, King and Logue 1987, Kirby and Herrnstein 1995, Kirby and Marakovic 1995, 1996, Kirby 1997, Read et al 1996. See Ainslie 1992 and Frederick et al 2001 for partial reviews of this literature. See Mulligan 1997 for a critique.

²The rate of decline of the discount function is given by $-\frac{f'(\tau)}{f(\tau)}$, where $f(\tau)$ is the discount function. For the exponential case, $f(\tau) = \delta^\tau$, the rate of decline is $-\ln(\delta) \simeq 1 - \delta$.

experimental subjects, the estimates are better approximated by generalized hyperbolic functions than by exponential functions. In the original psychology literature, researchers used hyperbolic discount functions like $\frac{1}{\tau}$ and $\frac{1}{1+\alpha\tau}$, with $\alpha > 0$ (Chung and Herrnstein 1961, Ainslie 1992). The most general hyperbolic discount function (Loewenstein and Prelec 1992) weights events τ periods away with factor $\frac{1}{(1+\alpha\tau)^{\gamma/\alpha}}$, with $\alpha, \gamma > 0$. Figure 1 plots such a generalized hyperbolic discount function. Note that the generalized hyperbolic discount function declines at a faster rate in the short-run than in the long-run, matching the key feature of the experimental data.

To capture this qualitative property, Laibson (1997a) adopted a discrete-time discount function, $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots\}$, which Phelps and Pollak (1968) had previously used to model intergenerational time preferences. This “quasi-hyperbolic function” reflects the sharp short-run drop in valuation measured in the experimental time preference data and has been adopted as a research tool because of its analytical tractability. The quasi-hyperbolic discount function is only “hyperbolic” in the sense that it captures the key qualitative property of the hyperbolic functions: a faster rate of decline in the short-run than in the long-run. Laibson (1997a) adopted the phrase “quasi-hyperbolic” to emphasize the connection to the “hyperbolic discounting” literature in psychology (Ainslie, 1992). O’Donoghue and Rabin (1999a) call the $\{1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots\}$ preferences “present biased.” Krusell and Smith (2000) call these preferences “quasi-geometric.” Akerlof (1991) used a similar discount function to model the salience of the present: $\{1, \beta, \beta, \beta, \dots\}$.

Figure 1 plots the particular parameterization of the quasi-hyperbolic discount function used in our simulations, $\beta = .7$ and $\delta = .957$. Using annual periods, these parameter values roughly match experimentally measured discounting patterns. Delaying an immediate reward by a year reduces the value of that reward by approximately $\frac{1}{3} \approx 1 - \beta\delta$. By contrast, delaying a distant reward by an additional year reduces the value of that reward by a much smaller percentage: $1 - \delta$.

All forms of hyperbolic preferences induce dynamic inconsistency.³ Consider the discrete-time

³Dynamic inconsistency refers to preferences which contradict the decision-maker’s own preferences at a later date. For example, imagine that on Monday I prefer to quit smoking on Tuesday but that on Tuesday I change my mind (with no new information) and now prefer to quit smoking on Wednesday. This agent holds preferences that

quasi-hyperbolic function. From the perspective of time 0, the value of a util at time 11 relative to the value at time 10 is $\frac{\beta\delta^{11}}{\beta\delta^{10}} = \delta$. However, *from the perspective of time 10*, the value of a util at time 11 (1 period in the future) relative to the value at time 10 (right now) is $\frac{\beta\delta}{1} = \beta\delta$. Since δ is close to one, the period 0 perspective implies that utils in periods 10 and 11 are close in value. So self 0 wishes to be *patient* when considering tradeoffs between periods 10 and 11. Since β is much less than 1, the time 10 perspective implies that a util in period 10 is worth a lot more than a util in period 11. So self 10 wishes to be *impatient* between periods 10 and 11.

This dynamic inconsistency forces the hyperbolic consumer to grapple with intra-personal strategic conflict. Early selves wish to force their preferences on later selves. Later selves do their best to maximize their own interests. Economists have modeled this situation as an intra-personal game played among the consumer's temporally situated selves (Strotz, 1956). Recently, hyperbolic discount functions have been used to explain a wide range of anomalous economic choices, including procrastination, contract design, drug addiction, self-deception, retirement timing, and undersaving (e.g., Akerlof 1991, Barro 1999, Benabou and Tirole 2000, Carrillo and Marriotti 2000, Diamond and Koszegi 1998, Laibson 1994, 1996, 1997a, and O'Donoghue and Rabin 1999a, 1999b, 2000). We focus here on the implications for life-cycle savings decisions.

In the analysis that follows, we analyze the hyperbolic model with "sophisticated" consumers. These consumers correctly predict that later selves will not honor the preferences of early selves. Thus, the early selves take actions that seek to constrain their future selves.

An appealing alternative is to assume that consumers make current choices under the false belief that later selves will act in the interests of the current self. Such "naifs" have optimistic forecasts in the sense that they believe that future selves will carry out the wishes of the current self. Under

are dynamically inconsistent. A distinct kind of dynamic inconsistency sometimes arises in strategic interactions between distinct agents. For example, a durable goods manufacturer may wish to commit to permanently charge a high price for a good, thereby encouraging customers to buy the good immediately instead of waiting for subsequent price declines. But if this early buying were to take place, the manufacturer would then wish to subsequently lower the price to attract buyers who weren't willing to buy at the original price (Coase 1972; see also Kydland and Prescott 1977 and Barro and Gordon 1983).

this belief, the current self constructs the sequence of actions that maximizes the preferences of the current self. The current self then implements the first action in that sequence, expecting future selves to implement the remaining actions. Of course, those future selves conduct their own optimization and therefore implement actions that potentially do not maximize earlier selves' preferences. This "naif" assumption was first proposed by Strotz (1956) and has been studied by Akerlof (1991) and O'Donoghue and Rabin (1999a, 1999b, 2000), who show that naifs and sophisticates sometimes behave in radically different ways. However, in the particular economic setting that we simulate, the choices of naifs and sophisticates are nearly identical, so we focus only on the sophisticates. Details of analogous naif simulations are available in a longer working paper, Angeletos et al. (2001).

3. A MODEL OF CONSUMPTION OVER THE LIFECYCLE

We analyze a special case of a model of hyperbolic discounting developed in Laibson, Repetto and Tobacman (2000). 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 (1999), and Laibson, Repetto and Tobacman (1998). The model of Laibson, Repetto and Tobacman (2000) incorporates most of the features of previous life-cycle simulation models and adds new features, including credit cards, variation in household size over time, and the ability to invest in illiquid assets. We summarize the key features of the model below and refer interested readers to Laibson, Repetto and Tobacman (2000) for details and for extensions including an option to declare bankruptcy and an ability to borrow against illiquid collateral (for example, mortgages on housing).

In the simulations presented here, households live for a maximum of 90 years, beginning economic life at age 20, and retiring at age 63. Household composition — number of adults and non-adults — varies over the life-cycle. Labor income is autocorrelated over time, but can be affected by stochastic shocks. The level of labor income and the size of the shocks are calibrated

to match empirical data. In the current paper, we will focus only on households whose head has only a high school degree, which accounts for roughly half of U.S. households. Laibson, Repetto and Tobacman (2000) analyze households across three different levels of educational attainment.

Households in our simulations may hold both liquid assets and illiquid assets, and they may borrow up to a credit limit equal to 30 percent of average labor income for their age-group. The real after-tax interest rate received from liquid assets is 3.75 percent. We chose this return because it corresponds to the return realized by a household with two-thirds of its assets in stocks and one-third in risk free bonds, assuming an average tax rate of 25%. The real interest rate on credit-card loans is 11.75 percent, 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. The credit limit — 30 percent of income — is calibrated from data on credit limits in the Survey of Consumer Finances. The illiquid asset, which can be thought of as housing, generates annual consumption flows equal to 5 percent of the value of the asset. Hence, the return on illiquid assets is considerably higher than the return on other assets. However, the illiquid asset can only be sold with a transaction cost equal to \$10,000 and 10 percent of the value of the asset.

These asset market assumptions imply that the household has an implicit crude “commitment” technology. Because sales of the illiquid asset generate transaction costs, wealth invested in the illiquid asset is partially protected from future splurges. This transaction cost enables the consumer to constrain herself *imperfectly* by investing wealth in the illiquid asset. Had we assumed that *perfect* commitment technologies existed, then the sophisticated consumer with a hyperbolic discount function would behave much like a consumer with an exponential discount function. With a perfect commitment technology the sophisticated hyperbolic consumer would be able to commit to future actions that maximize the young self's preferences for the future, which are exponential.

We assume that households have preferences toward risk characterized by a coefficient of rela-

tive risk aversion of two.⁴ Households have either an exponential discount function ($\delta_{\text{exponential}}^t$) or a quasi-hyperbolic discount function ($\beta\delta_{\text{hyperbolic}}^t$). For the hyperbolic case, we fix $\beta = .7$, corresponding to the one-year discount factor typically measured in laboratory experiments.

We assume that the economy is either populated exclusively by exponential households or exclusively by hyperbolic households. We pick $\delta_{\text{exponential}}$ and $\delta_{\text{hyperbolic}}$ to match empirical levels of retirement saving. Specifically, $\delta_{\text{exponential}}$ is picked so that the exponential simulations generate a median wealth to income ratio of 3.2, for individuals between ages 50 and 59. The median of 3.2 is calibrated from the Survey of Consumer Finances.⁵ The hyperbolic discount factor, $\delta_{\text{hyperbolic}}$, is also picked to match the empirical median of 3.2.

The discount factors, $\delta_{\text{exponential}}$ and $\delta_{\text{hyperbolic}}$, that replicate the SCF wealth to income ratio are .944 for the exponential model and .957 for the hyperbolic model. Since hyperbolic consumers have two sources of discounting — β and δ — the hyperbolic δ 's lie above the exponential δ 's. Recall that the hyperbolic and exponential discount functions are calibrated to generate the same amount of pre-retirement wealth accumulation. In this manner the calibrations “equalize” the underlying willingness to save between the exponential and hyperbolic consumers. The calibrated long-term discount factors are sensible when compared to discount factors that have been used in similar exercises by other authors. Finally, note that these discount factors do not include age-dependent mortality effects, which reduce the respective discount factors by an additional one percent on average per year. The underlying question is how this generally similar willingness to save is translated into actual patterns of consumption over a lifetime.

When a household has a hyperbolic discount function the household will have dynamically inconsistent preferences. So the problem of allocating consumption over time cannot be treated as a straightforward optimization problem. A *sophisticated* hyperbolic consumer realizes that selves

⁴Typically, the coefficient of relative risk aversion is set between one and five. If a household has a coefficient of relative risk aversion of two, then a household is indifferent between sure consumption of \$66,667 and a 50/50 gamble between \$50,000 and \$100,000 of consumption.

⁵Wealth does not include social security wealth and other defined benefit pensions, which are already built into the model in the form of post-retirement “labor income.”

at later points in time will not implement the policies that are optimal from the perspective of selves at earlier points in time.

Following the work of Strotz (1956), we model a sophisticated consumer as a sequence of rational players in an *intra*-personal game. Selves $\{20, 21, \dots, 90\}$ 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. At an equilibrium, all selves choose optimal strategies given the strategies of all other selves. We numerically compute the equilibrium strategies using a backwards induction algorithm.

4. RESULTS OF THE EXPONENTIAL AND HYPERBOLIC SIMULATIONS

Figure 2 plots the mean consumption profile for households with an exponential discount function and households with a hyperbolic discount function. Figure 2 also plots the calibrated mean labor income profile which is fixed by the data and hence is the same for both types of consumers. The exponential and hyperbolic consumption profiles roughly track the labor income profile. This comovement is driven by two factors. First, low income early in life holds down consumption, since consumers cannot borrow much against future income. After all, in this model the credit card borrowing limit is .30 of one year's income, which is not enough to smooth consumption over the life cycle. Second, consumption needs peak in mid-life, as the number of children increase. The number of children in the household reaches a peak of 2.09 when the household head is 36 years old. As the number of children declines, the household begins to support more and more adult dependents (i.e., the grandparents of the children), reaching a peak of .91 adult dependents at age 51.

Figure 2 compares the consumption profile of exponential households with the profile of hyperbolic households. These two consumption profiles are almost indistinguishable, with small differences arising at the very beginning of life, around retirement, and at the very end of life. At the beginning of life, consumers with hyperbolic discount rates go on a spending spree financed with credit cards, leading to higher consumption than the households with the exponential dis-

count function.⁶ Around retirement, hyperbolic consumption falls more steeply than exponential consumption, since hyperbolic households have most of their wealth in illiquid assets, which they cannot cost-effectively sell to smooth consumption. At the end of life, hyperbolic consumers have more illiquid assets to sell, supporting a higher level of late-life consumption.

The top panel of Figure 3 plots the mean levels of liquid assets, illiquid assets, and total assets for our simulated exponential households. Liquid assets include year-end liquid financial assets and $\frac{1}{24}$ of annual labor income. The latter term adjusts for the fact that our annualized discrete time model doesn't contain any motive to hold cash. If labor income is paid in equal monthly installments, $\frac{Y}{12}$, and consumption is smoothly spread over time, then average cash inventories resulting from monthly income will be $\frac{Y}{24}$.

Liquid wealth is held as a buffer against income shocks and against the fall in income at retirement. As a result, liquid wealth has a local peak just before retirement. Illiquid wealth is accumulated up until retirement and is then sold off after age 70. Most households sell all of their illiquid wealth in one transaction, thereby minimizing transaction costs, which include proportional and fixed components. However a small proportion of wealthy households continue to hold illiquid wealth through the end of life because of a bequest motive. The sell-off in illiquid wealth generates a late-life jump in liquid wealth as assets are shifted from one account to the other.

Liquid financial assets accumulate until they reach a temporary plateau at age 30. This buffer stock of liquid wealth is used to ride out transitory shocks during working life. More liquid wealth is accumulated in the decade before retirement (ages 53-63) to smooth out the drop in labor income at retirement. Illiquid accumulation begins at age 30 and peaks at age 63. Late in life, illiquid wealth is sold, transformed into liquid wealth, and then consumed.

The bottom panel of Figure 3 plots the mean level of liquid liabilities — i.e., credit card debt on which interest is paid — for our simulated households with an exponential discount rate. The bottom panel shows that credit card borrowing grows quickly early in life. It then remains fairly

⁶See Gourinchas and Parker (1999) for empirical evidence on an early life consumption boom.

steady between ages 30 and 40, and then is paid off between ages 40 and 60.

Before comparing the wealth accumulation of simulated exponential and simulated hyperbolic households, note that the total assets accumulated — including liquid and illiquid assets — are almost identical. This similarity is built into the simulation framework, since after all, the time preference parameters for the exponential and hyperbolic discount functions were calibrated to match observed levels of pre-retirement wealth holding. The only noticeable difference between the simulated total asset profiles occurs after retirement. Households with exponential discount functions spend down their retirement savings much more quickly than those with hyperbolic discount functions, since those with exponential discount functions hold less of their assets in illiquid form.

Figure 4 illustrates this point by plotting the average illiquid asset holdings of exponential and hyperbolic households. Households with hyperbolic discount functions begin accumulating the illiquid asset earlier and continue actively accumulating longer. Households with hyperbolic discount functions are more willing to hold illiquid wealth for two reasons. First, they view illiquid assets as a commitment device, which they value since it prevents later selves from splurging saved wealth. Second, illiquid assets have the same property as the goose that laid golden eggs (Laibson 1997a). The asset promises to generate substantial benefits, but these benefits can only be realized by holding the asset for a long time. Trying to extract value quickly — by slaughtering the goose or selling the illiquid asset — reduces the value of these assets. Such illiquid assets are particularly valuable to hyperbolics, since hyperbolics have a relatively low *long-run* discount rate.⁷ Hyperbolics and exponentials dislike illiquidity for the standard reason that illiquid assets can't be used to buffer income shocks, but this cost of illiquidity is partially offset for hyperbolics since they value commitment and they more highly value the long-run dividends of illiquid assets. Hence, on net, illiquidity is more costly for a household with exponential discount functions than a household

⁷The long-run discount rate of a hyperbolic consumer, $-\ln(\delta_{\text{hyperbolic}}) = -\ln(.957) = .044$, is calibrated to lie below the long-run discount rate of an exponential consumer, $-\ln(\delta_{\text{exponential}}) = -\ln(.944) = .058$. To see these effects graphically, note that in Figure 1 both hyperbolic curves eventually rise above the exponential curve.

with hyperbolic discount functions, explaining why hyperbolics hold a higher share of their wealth in illiquid form.

Conversely, households with hyperbolic discount functions tend to hold relatively little liquid wealth. Figure 5 plots the liquid financial assets and credit card debt for the two types of simulated households. Households with hyperbolic discount functions hold far more credit card debt and lower levels of liquid assets than households with exponential discount functions. The hyperbolic households end up holding relatively little liquid wealth and high levels of liquid debt because liquidity tends to be used to satisfy the hyperbolic taste for instant gratification. Households with hyperbolic discount functions view credit cards as a mixed blessing. Credit cards enable future selves to splurge (a cost) but credit cards also provide liquidity when income shocks hit the household.

5. EMPIRICAL EVALUATION OF THE SIMULATION RESULTS

In this section we evaluate empirically the predictions of the simulation models discussed above. We focus on simulation predictions about liquid and illiquid wealth accumulation, credit card borrowing, and consumption-income comovement. Relative to households with exponential discount functions, households with hyperbolic discount functions hold less liquid wealth, hold more illiquid wealth, borrow more aggressively on credit cards, and smooth consumption less successfully over the lifecycle. As we will argue, the hyperbolic model can make sense of a wide range of facts about the life-cycle choices of U.S. households.

5.1. Wealth Accumulation and Asset Allocation. We begin by considering the percentage of households with high levels of liquid wealth. Using our simulated data, we calculate the ratio of liquid assets to annual labor income for every household at every age. We then report the percentage of households which have liquid asset holdings that are at least as large as one month of labor income. The results of this analysis of simulated data and survey data are reported in Table 1. For example, from ages 40 to 49, 72 percent of simulated households with exponential discount

functions hold liquid assets greater than one month of labor income. The analogous number for households with hyperbolic discount functions is only 38 percent. For comparison, between 26 percent and 42 percent of households between ages 40 and 49 in the Survey of Consumer Finances hold liquid financial assets greater than one month of labor income. The range of empirical values reflects three different liquid asset definitions adopted in our analysis. The narrowest definition of liquid assets is just cash, checking, and savings accounts. An intermediate definition of liquid assets includes those items plus money market accounts. Our most inclusive definition of liquid assets includes the previous items plus call accounts, certificates of deposit, bonds, stocks and mutual funds. Our results change very little as we vary the liquid assets definition. The proportions of households in the Survey of Consumer Finances with liquid assets greater than one month of income, based on these three definitions, are shown in the final three columns of Table 1.

Consider our intermediate definition of liquid financial assets. Using this definition, on average 43 percent of actual households hold liquid assets greater than one month of household income. This percentage rises over the life cycle. The simulated exponential model does a poor job of approximating this survey data. The exponential profile lies *everywhere* above the empirical profile, with an average difference of 30 percentage points. The simulation of households with hyperbolic discount functions does a much better job of approximating the empirical measures of liquid wealth holding. The hyperbolic profile intersects the empirical profile, with an average difference of only 3 percentage points. However, the hyperbolic simulations do not predict the sharp empirical rise in liquid wealth holding over the lifecycle.

In Table 2 we evaluate the theoretical models by analyzing the simulated quantity of liquid assets as a share of total assets, what we call the “liquid wealth share.” In the data from the Survey of Consumer Finances, the average liquid wealth share is only 10 percent (using the intermediate definition of liquid wealth) and neither the exponential nor hyperbolic simulations come close to matching this number, though the hyperbolic simulations are a bit closer to the mark. This failure may arise because illiquid assets in the real world are less illiquid than illiquid assets in our

simulations. By lowering the transactions costs of our simulated illiquid assets we can increase our simulated consumers' willingness to hold illiquid assets. In addition, liquid assets in the real world may provide a lower rate of return than liquid assets in our model. If most liquid assets are checking account balances — instead of stocks and corporate bonds — then the liquid return may be lower than 3.75 percent. A lower *liquid* return would increase the relative appeal of the *illiquid* asset.

Revolving credit — that is, credit card borrowing — represents an important form of liquidity. Low levels of liquid net assets are naturally associated with high levels of credit card debt. At any point in time, only 19 percent of simulated consumers with an exponential discount function borrow on their credit cards, compared to 51 percent of those with a hyperbolic discount function. By comparison, in the 1995 Survey of Consumer Finances, 70 percent of households with credit cards report that they did not fully pay their credit card bill the last time that they mailed in a payment. Both of the simulated results are low relative to this empirical benchmark, but hyperbolic simulations do come closer to matching the available data.

Analogous results arise when we measure the average amount borrowed on credit cards. On average, simulated households with exponential discount functions owe \$900 of interest-paying credit card debt, *including* the households with no debt. By contrast, simulated households with hyperbolic discount functions owe \$3,400 of credit card debt. The actual amount of credit card debt owed per household with a credit card is over \$5,000 (*including* households with no debt, but *excluding* the float).⁸ Again, the hyperbolic simulations provide a better approximation than the exponential simulations.

5.2. Comovement of Consumption and Labor Income. Since Hall's (1978) pathbreaking work, the core of the empirical consumption literature has been based on tests of comovement

⁸This average balance includes households in all education categories. It is calculated on the basis of aggregate information reported by the Federal Reserve. This figure is consistent with values from a proprietary account-level data set assembled by David Gross and Nicholas Souleles (1999a, 1999b, 2000). See Laibson, Repetto, and Tobacman (2000) for a much more detailed analysis of credit card borrowing.

between consumption and income. Hall argued that if current consumption is based on asset wealth and the net present value of future income flows, then changes in consumption will only occur if news arrives which changes wealth or the expectations of future income. As a result, the standard model of consumption, in which consumers face no liquidity constraints, implies that the marginal propensity to consume out of predictable changes in income should be zero.

By contrast, empirical estimates of the marginal propensity to consume out of predictable changes in income lie generally between 0 and .5, with typical estimates between .1 and .3. For example, Altonji and Siow (1987) report a statistically insignificant coefficient of .091; Attanasio and Weber (1993) report an insignificant coefficient of .119; Attanasio and Weber (1995) report an insignificant coefficient of .100; Hall and Mishkin (1982) report a significant coefficient of .200; Hayashi (1985) reports a significant coefficient of .158; Lusardi (1996) reports a significant coefficient of .368; Shea (1995) reports a marginally significant coefficient of .888; and Souleles (1999) reports a significant coefficient of .344. See Deaton (1992) and Browning and Lusardi (1996) for discussion of the literature on the excess sensitivity of consumption to income.

We have replicated the standard comovement regressions using the 1978-1992 surveys from the Panel Survey of Income Dynamics. We use a range of definitions of consumption and several different assumptions on the measurement error in income (Angeletos et al. 2001). We predict expected income growth with a range of instruments: lagged income, lagged hours worked by the head and spouse, and race and marital status dummies. Our empirical estimates of the marginal propensity to consume out of predictable changes in income lie between .19 and .33 and most are significant at the 5 percent level.⁹

We have also estimated the standard comovement regression using our *simulated* households with exponential discount functions. For these households, the estimated marginal propensity

⁹Specifically, we estimate,

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

Here X_{it} is a vector of control variables. The standard consumption model (without liquidity constraints), predicts $\alpha = 0$; the marginal propensity to consume out of predictable changes in income should be zero.

to consume out of predictable changes in income is only 0.03. This estimate lies slightly above zero because of liquidity constraints. But liquidity constraints do not bind or come close to binding often enough in the calibrated exponential model to push the marginal propensity to consume far above zero. The calibration of the exponential discount function yields a discount factor $\delta_{\text{exponential}} = .944$, implying an exponential discount rate of .056. This discount rate is not high enough to make the liquidity constraints matter. Hence, consumption comoves only weakly with predictable changes in income. With higher discount rates the implied comovement would be stronger, but such higher discount rates are inconsistent with the observed empirical level of retirement wealth accumulation. Recall that the level of retirement wealth was used to estimate $\delta_{\text{exponential}}$.

We have also estimated the marginal propensity to consume out of predictable changes in income using the simulated households with hyperbolic discount functions. For our hyperbolic households we estimate a marginal propensity to consume of 0.166. This estimate compares well with the available empirical evidence. Households with hyperbolic discount functions hold more of their wealth in illiquid form than households with exponential discount functions. So households with hyperbolic discount functions are more likely to hit liquidity constraints, raising their marginal propensity to consume out of predictable changes in income.

We also use our simulations to investigate income-consumption comovement around the time of retirement. Banks et al. (1998) and Bernheim et al. (1997) argue that consumption falls particularly steeply as consumers enter retirement. From the standpoint of the standard theory of consumption, this decline may be anomalous; rational consumers should anticipate retirement and should smooth their consumption even though their income is falling.

Using data from the Panel Survey of Income Dynamics, we confirm the results of Bernheim et al. and estimate a statistically significant excess drop in consumption of 11.6 percent in the

four-year window around retirement (Angeletos et al. 2001).¹⁰ We then ran analogous regressions on our simulated data. For households with an exponential discount function, we estimate an excess drop in consumption of 3.0 percent around retirement. For households with a hyperbolic discount function, the analogous drop is 14.5 percent, close to the corresponding PSID estimate. The underlying reason for the difference is that hyperbolic consumers hold relatively little liquid wealth. Therefore, a drop in income at retirement translates into a substantial drop in consumption, even though retirement is an exogenous, completely predictable event. Naturally, this is just one explanation for the empirical drop in consumption at retirement. Consumption may also fall because of leisure-consumption substitution and because retirement correlates with negative health or labor supply shocks.¹¹ Neither of these potentially important effects is present in our model.

6. CONCLUSIONS

All in all, a model of consumption based on a hyperbolic discount function consistently better approximates the data than does a model based on an exponential discount function. The hyperbolic discount function turns out to be useful in helping to explain why households hold relatively low levels of liquid wealth, measured either as a fraction of labor income or as a share of total wealth. The model also helps to explain why households borrow so aggressively with their credit cards. In addition, because the households with hyperbolic discount functions have relatively low levels of liquid assets and relatively high levels of credit card debt, they are less able to smooth their consumption paths in the presence of predictable changes in income. As a result, their consumption responds to predictable changes in income, matching well documented empirical patterns of consumption-income comovement. For similar reasons, the hyperbolic discount function helps to

¹⁰Specifically, we estimate,

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

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, and X_{it} is a vector of control variables, including mortality and household composition. By summing the coefficients on the four dummy variables (and switching signs), we get an estimate of the “excess” drop in consumption around retirement.

¹¹Banks et al. (1998) and Bernheim et al. (1997) argue that leisure-consumption substitutability effects and “bad news” effects can not explain the drop in consumption within a rational expectations framework.

explain why consumption drops around retirement.

Future work should enrich the lifecycle modelling framework. Expanding the number of assets and the sources of uncertainty pose important challenges. Such extensions would overwhelm the current generation of computers. But as computers get faster, richer simulations will become possible. We are also intrigued by analysis that examines the high frequency behavior of decision-makers with hyperbolic discount functions. To economize on computer resources, the model in the current paper analyzes a model with annual decisions. But self-control problems arise from moment to moment — think of the chocolate soufflé à la mode. A complete understanding of self-regulation will require high frequency analysis of intertemporal choices, including models embedded in continuous time (Barro, 1999, Harris and Laibson 2001, Luttmer and Mariotti 2000). Other interesting extensions include analysis of “mixed” economies populated by different types of consumers — exponential and hyperbolic (Laibson, Repetto and Tobacman 1998). It will also be important to understand the role of sophistication in the savings and asset allocation decisions of hyperbolic consumers (O’Donoghue and Rabin, 2000). Researchers should also endogenize the menu of contracts provided by profit-maximizing firms to hyperbolic consumers (O’Donoghue and Rabin, 1999b, Della Vigna and Malmendier 2001).

Future work should also attempt to estimate time preference parameters using field data. In preliminary work along these lines, we have empirically evaluated the model described in this paper using the method of simulated moments (Pakes and Pollard 1989). Following this approach we have tightly estimated the hyperbolic preference parameters: $\beta = .55$ (standard error .05) and $\delta = .96$ (standard error .01). This inference joins a nascent literature in which structural hyperbolic models are estimated with field data (e.g., Della Vigna and Paserman 2001).

We believe that the model of consumption based on a hyperbolic discount function will prove useful because it provides a parsimonious framework that makes sharp empirical predictions. Naturally, we expect such empirical predictions to identify both the strengths and the failings of the hyperbolic approach. We predict that the failings will reflect the fact that the hyperbolic model

is a reduced form that captures a more complicated underlying preference structure. For now, the hyperbolic model represents an empirically useful parsimonious representation of our self-control problems. Future authors will undoubtedly uncover the deeper cognitive subcomponents of the struggle for self-command.

7. REFERENCES

- Ainslie, George. 1992. *Picoeconomics*. Cambridge: Cambridge University Press.
- Akerlof, George A. 1991. "Procrastination and Obedience." *American Economic Review. (Papers and Proceedings)*, pp. 1-19.
- Altonji, Joseph and Aloysius Siow. 1987. "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data." *Quarterly Journal of Economics*. 102:2, pp. 293-328.
- Angeletos, George-Marios, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg. 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 from 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.
- Barro, Robert J. 1999. "Laibson Meets Ramsey in the Neoclassical Growth Model." *Quarterly Journal of Economics*. 114(4), pp. 1125-52.
- Barro, Robert J. and David B Gordon. 1983 "A Positive Theory of Monetary Policy in a Natural Rate Mechanism." *Journal of Political Economy* vol 91 no. 4 pp. 589-610.

- Benabou, Roland and Jean Tirole. 2000. "Willpower and personal rules." Mimeo.
- 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.
- 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.
- Deaton, Angus. 1992. *Understanding Consumption*. Oxford University Press: Oxford.
- Deaton, Angus 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.
- DellaVigna, Stefano 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.
- Engen, Eric, William Gale and John Karl Scholz. 1994. "Do Saving Incentives Work?" *BPEA*. 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.
- Gourinchas, Pierre-Olivier and Jonathan Parker. 1999. "Consumption Over the Life-Cycle." Mimeo.
- Gross, David and Nicholas Souleles. 1999a. "An Empirical Analysis of Personal Bankruptcy and Delinquency." Mimeo.

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

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

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. 2000. "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.

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

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

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

- _____ and Nino N. Marakovic. 1995. "Modeling Myopic Decisions: Evidence for Hyperbolic Delay-Discounting within Subjects and Amounts." *Organizational Behavior and Human Decision Processes*. 64:1, pp. 22-30.
- Krusell, Per and Anthony A Smith, Jr. 1999. "Consumption-Savings Decisions with Quasi-Geometric Discounting." Mimeo, University of Rochester.
- _____ and _____. 1996. "Delayed-Discounting Probabilistic Rewards Rates Decrease as Amounts Increase." *Psychonomic Bulletin and Review*. 3:1, pp. 100-104.
- Kydland, Finn E. and Edward C. Prescott. 1977 "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* vol 85 no. 3 pp. 473-491.
- Laibson, David I. 1994. "Self-Control and Savings." Ph.D. diss., Massachusetts Institute of Technology.
- . 1996. "Hyperbolic Discounting, Undersaving, and Savings Policy." Working Paper 5635. Cambridge, Mass: National Bureau of Economic Research.
- . 1997a. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics*. 62:2, pp. 443-478.
- . 1997b. "Hyperbolic Discount Functions and Time Preference Heterogeneity." Harvard mimeo.
- . 1998. Comments on "Personal Retirement Saving Programs and Asset Accumulation," by James M. Poterba, Steven F. Venti, and David A. Wise, in *Studies in the Economics of Aging*. David A. Wise, ed. Cambridge, Mass: NBER and the University of Chicago Press, pp. 106-24.
- , Andrea Repetto and Jeremy Tobacman. 1998. "Self-Control and Saving for Retirement." *BPEA*. 1, pp. 91-196.

———, ——— and ———. 2000. “A Debt Puzzle.” Mimeo.

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

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

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

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

Mulligan, Casey. 1997. “A Logical Economist’s Argument Against Hyperbolic Discounting.” University of Chicago mimeo.

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

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

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

——— and ——— 2000. “Choice and Procrastination.” Working paper.

Pakes, Ariel and David Pollard. 1989 “Simulation and the Asymptotics of Optimization Estimators.” *Econometrica*. Vol. 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.
- Rankin, Deborah M. 1993. "How to Get Ready for Retirement: Save, Save, Save." *New York Times*. March 13, pp. 33.
- Read, Daniel, George Loewenstein, Shobana Kalyanaraman and Adrian Bivolaru. 1996. "Mixing Virtue and Vice: The Combined Effects of Hyperbolic Discounting and Diversification." Working paper, Carnegie Mellon University.
- Runkle, David. 1991. "Liquidity Constraints and the Permanent-Income Hypothesis: Evidence From Panel Data." *Journal of Monetary Economics*. 27:1, pp.73-98.
- Shapiro, Matthew D. and Joel Slemrod. 1995. "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding." *American Economic Review*. 85:1, pp. 274-83.
- Shea, John. 1995. "Union Contracts and the Life-Cycle/Permanent Income Hypothesis." *American Economic Review*. 85:1, pp.186-200.
- Simmons Market Research Bureau. 1996. *The 1996 Study of Media and Markets*. New York.
- Solnick, Jay V. et al. 1980. "An experimental analysis of impulsivity and impulse control in humans," *Learning and Motivation* 11(1), 61-77.
- Souleles, Nicholas. 1999. "The Response of Household Consumption to Income Tax Refunds." *American Economic Review*.
- Strotz, Robert H. 1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *Review of Economic Studies*. 23, pp. 165-180.
- Thaler, Richard H. 1981. "Some Empirical Evidence on Dynamic Inconsistency." *Economic Letters*. 8, pp. 201-207.

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

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

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

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