LEARNING AND VISCERAL TEMPTATION IN DYNAMIC SAVINGS EXPERIMENTS[*]

Alexander L. Brown

Zhikang Eric Chua

Colin F. Camerer

Abstract

This paper tests two explanations for apparent undersaving in lifecycle models: Bounded rationality; and a preference for immediacy. Each was addressed in a separate experimental study. In first study, subjects saved too little initially—providing evidence for bounded rationality—but learned to save optimally within four repeated lifecycles. In the second study, thirsty subjects that consume beverage sips immediately rather than with a delay show greater relative overspending, consistent with quasi-hyperbolic discounting models. The parameter estimates of overspending obtained from the second study—but not the first—are in range of several empirical studies of savings (with an estimated $\beta=0.6$-0.7).

I. Introduction

Field evidence on whether people save optimally is mixed.[1] Some recent studies suggest saving is optimal;[2] other recent research argues that consumers make fundamental mistakes regarding savings[3] and typically undersave[4] relative to optimal levels. One reason for mixed results is that econometric tests must make several auxiliary assumptions about the underlying model.[5] Laboratory experiments may be of some use in this area of research, because the maintained assumptions of a particular theory (e.g., about utility functions and beliefs about the income process) can be clearly implemented with experimental control.

This paper explores how well people make dynamic optimization[6] decisions in a complex experimental environment. The essential design parameters were taken from models of consumer savings with income uncertainty and habit formation, which created a two-state finite dynamic program. To follow the optimal savings path, subjects must consume very little in early periods to buffer against bad income shocks and to avoid creating an early consumption "internality" from habit that reduces utility from future consumption.

This paper tests two classes of explanations for apparent evidence of undersaving: Bounded rationality; and a preference for immediacy (or present-bias). Each was addressed in a separate experimental study. Reporting the two studies together enables direct judgment of which explanation is generally better (or whether both have some merit) and allows some parametric comparison with field data.

In the first study, subjects had an opportunity to learn privately over seven experimental lifecycles or to learn "socially" from the choices of other subjects. If learning creates movement toward optimal choice, that is *prima facie* evidence of bounded rationality in initial choices (because highly rational subjects would not need to learn). The goal in this study was to see how

close subjects were to optimal, how well they learned, and whether deviations from optimality resemble those in field data (as calibrated by β and δ parameters in quasi-hyperbolic models). The results of the first study were that subjects undersaved at first but were able to learn very quickly from social information and less quickly from private learning.

The approximation to optimality observed after learning was surprisingly close. This near-optimality after learning turned our attention to the second basic explanation for undersaving, a preference for immediacy. Even if people fully understand optimal savings rules, optimal saving might be hard to implement if the choices that generate current utility are viscerally tempting (as in the most extreme case, of addictions).

Therefore, the second study used thirsty subjects who chose how much beverage to consume, as an experimental-scale model of more dramatic types of visceral temptation. Some subjects received beverages immediately (i.e., their period t decisions led to physical consumption in period t) and others received beverages with a 10-period delay (as if ordering from a catalog for later delivery). This immediate-delayed contrast is a simple way to detect the preference for immediacy (since the β term only influenced decisions when consumption was immediate).

In this second experiment, subjects generally consumed more when rewards were immediate than when rewards were delayed, a difference which is consistent with models of hyperbolic discounting (i.e., Ainslie, 1975; Laibson, 1997) and dual-self conflict (e.g., Bernheim and Rangel, 2004; Loewenstein and O'Donoghue, 2004; Fudenberg and Levine, 2006). Structural parameter estimates of β and δ yielded values comparable to those in other lab and field studies (mean β of 0.6-0.7), albeit over very different time horizons.

II. Two Explanations for Undersaving

The experimental design implemented the assumptions of the buffer stock savings model of Carroll, Overland and Weil (2000). Agents earned income each period, subject to stochastic independent shocks from a distribution they knew. In each period their available cash was the previous period's savings (buffer stock), plus new income. In each period they chose how much of this available cash to spend on consumption and the rest was saved. Utility in each period depended upon a ratio of current consumption to a habit index. The habit index was a depreciated sum of previous consumption (as in the pioneering design of Fehr and Zych [1998] based on Becker and Murphy [1988]). An entire 30-period lifecycle was repeated several times with different income realizations each time. Two alternative explanations, which have been hypothesized to explain consumer undersaving in the field, suggested why subjects might have saved too little in this experimental environment.

1. Bounded Rationality

One explanation for apparent undersaving is the bounded rationality of consumers;[7] They may undersave because their rationality is bounded, and solving for optimal saving in the buffer stock model is computationally difficult.[8] Carroll (2001) puts it bluntly

> "One problem is the spectacular contrast between the sophisticated mathematical apparatus required to solve the optimal consumption problem and the mathematical imbecility of most consumers. (p. 41)"

After all, economists were unable to solve the general form of the *same* problem posed in our experiment themselves, before later advances in computing. (Instead, they used an approximation with certainty-equivalent income for many years.) The results of savings experiments agree—most find evidence of undersavings and attribute it to bounded rationality.[9] Optimal saving is unlikely to be reached in buffer stock model through trial-and-error either;

Allen and Carroll (2001) show that learning by simple reinforcement is far too slow to produce convergence to optimal saving in reasonable time scales. Consumers may not be able to figure out or learn over time (with modest experience) how to save optimally.

Since there was no widely accepted theory of how bounded rationally should be modeled formally in these settings,[10] the presence of rationality bounds was inferred indirectly: if subjects made mistakes in the first lifecycle, but learned over time or from the social examples[11] (our first study also included social learning, similar to Ballinger, Palumbo, and Wilcox [2003]), then we would infer that their initial mistakes resulted from bounds on rationality, because subjects with unbounded rationality would not need to learn from experience.

Development of a more precise theory of rationality bounds and learning remains a priority for future research (and is discussed further in the conclusion).

2. Temptation and Dynamically Inconsistent Preferences

A second explanation for undersavings is that consumers know how to save optimally, but cannot resist short-term temptations to spend. For example, the availability of widespread credit can contribute to overspending if, psychologically, credit cards anesthetize the "pain of paying" (Prelec and Loewenstein, 1998). As Carroll (2001) suggests

> "There certainly seems to be strong evidence that American households are now using credit cards in non-optimal ways. The optimal use of credit cards (at least as implied by solving the final optimizing model discussed above) is as an emergency reserve to be drawn on only rarely, in response to a particularly bad shock or series of shocks. However, the median household with at least one credit card holds about $7000 in debt on all cards combined. (p. 42)"

Laibson, Repetto and Tobacman (2003) argue that this pattern is explained by consumers who have a powerful *preference* for immediate consumption, rather than bounded rationality in computing optimal spending. This explanation has been supported by empirical work on savings and neuroeconomics experiments.[12]

To test this explanation for undersavings, in the second experimental study we converted consumption from numbers to actual sips of beverage (for thirsty subjects). Comparing immediate and delayed delivery of beverage consumption enabled us to study the strength of temptation and dynamic inconsistency that might result. Of course, small amounts of beverage are not as dramatic as temptations like drug addiction, gambling, and credit card spending, but they were feasible in the lab and gave us a first contrast between money rewards and visceral temptations that can guide future research.

III. Study 1: Learning with money rewards

1. Experimental design

Participants were carefully instructed about the basic concepts of the experiment, and how their decisions and the random income draws would determine how much money they would earn (see our working paper Appendix 1 for details and instructional tables). To avoid demand effects and to enhance memorability, economic jargon like "income shocks," "habit stock," and "utility," were translated into plainer language—"adjustment factor," "lifestyle index," and "points," respectively.

Subjects chose C_t in each period from cash-on-hand, which is the sum of previous cash plus new income (Y_t). Income in each period was $Y_t = P_t \eta_t$, the product of P_t, permanent income that grew at five percent ($P_t = (1.05) P_{t+1}$, with initial $P_1 = 100$) and a multiplicative shock η_t, which was lognormally distributed ($\log \eta \sim N(-\frac{1}{2}, 1)$). There was no interest rate and discount factor, and no borrowing or investment. Period-specific utility depended on consumption and on an accumulated level of habit, according to $u(C_t, H_{t-1}) = k + \dfrac{\theta}{1-\rho} \left(\dfrac{C_t + \hat{\varepsilon}}{H_{t-1}^{\gamma}} \right)^{1-\rho}$ with risk-aversion parameter $\rho = 3$ and a habit strength exponent $\gamma = 0.6$.[13] The habit stock grew according to

5

$H_t = \lambda H_{t-1} + C_t$, where $\lambda = 0.7$ is a depreciation rate (as in Fehr and Zych, 1998) and the initial habit $H_0 = 10$. Thus, larger early consumption built up the habit level and depreciated future-period utility. This "internality" implied that optimization requires restrained consumption in early periods.

The subject's problem was to choose the stream of consumption $\tilde{C}_S$ in each period t to maximize his expected utility, $E_t \left[\sum_{s=t}^{T} u(\tilde{C}_S, \tilde{H}_{S-1}) \right]$. Because T=30 in the experiments, the problem could be simplified to a dynamic programming problem with two state variables, cash-on-hand C_t and habit H_t (after dividing both variables by the permanent income P_t).

The experimental environment was designed to have some basic empirical features of savings in the modern American economy. The 5% income growth and lognormality of multiplicative shocks are shown by Carroll (1992) to characterize US data. However, we chose T=30 to compress the lifecycles (compared to American annualized lifetimes) in order to create "lifecycles" which were long enough to create a savings challenge and interesting dynamics, but short enough to allow several lifecycles in each experimental session. We also multiplied the standard deviation of multiplicative income shocks η_t by five (creating a standard deviation of 1, rather than Carroll's estimate of 0.2) in order to deliberately produce more income variation.

The goal of experiments like these was *not* to precisely recreate all the empirical properties of naturally occurring decisions in a particular setting. After all, parametric properties of savings problems vary widely across periods of history and across countries so there is no single "real world" to serve as a unique design target. The goal, instead, was to explore a range of environments in which the theory might apply in order to judge when the theory is likely to work and when it is likely to fail. We deliberately chose income shock volatility that is larger

6

than that observed in the modern American economy because higher income variation created a more analytically challenging environment in which deviations from rationality would be more clearly observed. The design also combined uncertain income and habit formation, because (a) previous experiments have already studied each separately and (b) combining them made the problem much more challenging: if learning would occur, then the power of learning would be established with more force.

The instructions explained *all* the details of the structure described above. To make the details easier to understand, we included 30-draw samples from the lognormal distribution to give participants a feel for how much their income could vary and showed the utility functions and habit stock evolution using numerical tables (see our working paper). One table illustrated how the habit stock in each period was determined by the previous period's habit stock and the current spending. A separate table showed how their spending and habit stock in one period determined their utility points in that period. Before participating, subjects took a quiz testing them on how their choices, habit levels, and income shocks would determine utility points. The quiz was designed to satisfy concerns that suboptimal consumption decisions do not arise from confusion about how their decisions map into points (and eventual money earnings).

Consumption decisions were input to an Excel interface, which displayed the income obtained, the corresponding cash available, and the habit stock for each consumption choice (see Figure I). The program also calculated and displayed the possible points (i.e., utilities) that could be obtained from different levels of spending, and the corresponding savings available for the next period. Participants could experiment by inputting different consumption amounts and see how much utility they would earn, and how much cash they would have available at the start of the next period. Most participants tried out several spending choices before making a decision

(especially in the first couple of lifecycles). This process was repeated until the end of the lifecycle of 30 periods. (The program automatically spent all cash in the final period 30.) There were a total of seven lifecycles, to see how rapidly subjects could learn across lifecycles. Each participant's total payoff was a pre-announced linear function of the total points earned in all lifecycles[14] plus a $5 show-up payment. Subjects earned between $7.50 to $65 with an average of $45.

After thirty-six (36) subjects had participated in the private learning condition described above, thirty-six (36) more participated in a "social learning" condition. In the social learning condition, as part of their initial instructions, subjects were given samples of three actual subjects' period-by-period decisions for one lifetime in the private learning condition.[15] The three samples were taken from the highest-earning subject, the lowest-earning subject, and from one subject chosen at random from the private condition in their subject pool. The social learning subjects were told exactly how these three samples were chosen.

There are many ways to implement social learning or imitation (e.g., Ballinger, Palumbo, and Wilcox [2003] use direct talking). Our method mimicked intergenerational imitation in which a parent points out three role models—a great success who retires wealthy, a ne'er-do-well who ends up broke, and a random acquaintance. The high-earning role model might have been a subject who overspent early on (relative to the optimum) but got lucky by receiving high-income draws. In that case, subjects copied the "successful" subject too directly, they would have easily overspent relative to the optimum; so it was not clear whether social learning would have actually helped, hurt, or had no effect.[16]

Participants were 35 undergraduates from the National University of Singapore (NUS) and 37 undergraduates from California Institute of Technology. These students were unusually

adept at analytical thinking so they should represent an upper bound on how well average consumers do in these intertemporal optimization problems. The participants were recruited using the universities' mail servers. Half the participants (18 from each school) did the experiment with private learning and approximately half (17 NUS, 19 Caltech) did the experiment with social learning. Each group had seven lifecycles of 30 periods of income draws. To simplify data analysis, within each condition all participants received the same income draws (but the draws were different in the two learning conditions).[17] Most participants completed the instruction and seven lifecycles in about 90 minutes.

2. Basic results

Under optimality people should act as if they make ex ante optimal savings decisions under uncertainty, discounting future utilities exponentially, given their beliefs about future income and other structural parameters. In our experimental design, subjects should have saved a lot to build up a buffer stock, and then spent roughly their average income once their buffer stock is large enough. The buffer stock would protect against bad future income draws, and high early saving would limit the negative "internality" of current spending on future utility, (which occurs because of the controlled effect of habit formation). Figure II illustrates an optimal path of consumption, and cash-on-hand, given a particular lifecycle of income shocks (based on parameters used in the experiment, described later). Savings is the gap between the black optimal consumption line and the gray cash-on-hand line. In this example, the optimal consumer should spend less than current income in early periods except 6-7 (when income happened to be unusually low and consumers should dip into their savings to earn a reasonable utility from consumption). The optimal cash-on-hand in the example steadily rises to 1500 in period 20, building up a buffer stock which is about six times the annual income at that point. That is,

consumers should brace themselves for a rainy day by saving until about period 20. After period 20, they should start to dissave by spending more than their current income and dipping into their cash-on-hand (i.e., the optimal consumption line is usually above the dotted income line after period 20).

Table I gives summary statistics of actual point outcomes in the two learning conditions. The first and second rows give the average of total lifecycle points in each condition, and the standard deviation across subjects. The third row is the difference between the average point total and the (unconditional) optimal point total.[18] The fifth row is the total income in each lifecycle (which gives an idea of whether deviations from optimality in a particularly lifecycle are due to bad decisions or to bad income luck).

With only private learning, performance in the first three lifecycles was well below the unconditional optimum and highly variable across subjects. However, by lifecycle 4 the average subject earned point totals within 80% of the optimum and the variability across subjects shrank.

Table I (bottom panel) shows that social learning brought point outcomes close to the optimum rapidly. The mean and variation of points in the very first lifecycle with social learning are similar to those statistics from lifecycles 4-7 with only private learning.[19]

3. Behavior relative to conditional optimization

The Table I statistics compare point totals to *unconditional* optimal level of spending in each period. This can be a misleading comparison because *conditional* optimal spending in each period depended on the participant's *actual* cash-on-hand and accumulated habit stock. A subject who had made some bad decisions in early periods, but then wised up and made conditionally optimal decisions in later periods, would look bad in Table I but may have been close to

conditionally optimal overall when those few early mistakes are averaged with the smarter later decisions.

Each subject's average conditional deviation for each period is the difference between their actual spending and the optimum (conditioned on that participant's earlier decisions). Figure III plots the conditional deviation paths for lifecycles 1 and 7 with private learning, along with 95% confidence intervals (dotted lines). Since the optimal conditional path in Figure III is the zero-deviation horizontal line, the reader can judge at a glance whether deviations are significant by seeing whether the confidence interval covers the zero line or is far from it.

Figure III confirms the conclusion from Table I: With only private learning, participants in lifecycle 1 spent significantly more than optimal in early periods, until about period 20 (when they often spent too little). However, the lifecycle 7 conditional deviations are never significantly different from zero, which shows that learning was very effective over the seven lifecycles. In fact, the actual spending path is insignificantly different from the conditional optima by lifecycle 4.

Figure IV shows the analogous data for the social learning condition. These small deviations are deliberately plotted with the same y-axis scale as in Figure III, to show how much smaller the deviations are when there is social learning compared to private learning. Deviations are insignificantly different from zero in most periods. There is also little difference between lifecycles 1 and 7 in the social learning condition. The initial performance is so close to optimal that there is little left to learn over the seven lifestyles.

To measure the effects of private and social learning, we regressed the log of the absolute deviation from the conditional optimum on dummy variables for lifecycles (excluding the first

lifecycle), the period number and its square, and dummy variables for social learning condition, gender (Female=1, mean=0.43) and ethnicity (Chinese=1, mean=0.50).[20]

Table II shows the results. The period effect is positive (but nonlinear because the squared-period effect is negative) because the absolute deviations are larger in later periods, when incomes are larger. The social learning main effect is highly significant (it implies a 24% reduction in conditional deviation), as are the dummy variables for lifecycles, 5 and 7, reflecting learning across lifecycles. There is no significant effect of ethnicity and a small effect of gender (women deviate about 20% more).

IV. Study 2: Beverage rewards and temptation

1. Experimental design

Study 2 was the same as the first study except for one large change.[21] Lifecycles 1-2 and 4-5 (with money rewards) were the same as in study 1. However, in lifecycle 3 subjects received a fixed monetary payment for their participation but did not earn any additional money for decisions. Instead, in each period they drank an amount of a beverage[22] proportional to their consumption decisions each period (1 ml beverage for each 2 points). The Excel interface was modified to show the total milliliters of beverage reward to be obtained, rather than points (utilities). It also displayed the maximum milliliters of beverage reward that could be obtained from spending all available cash immediately. As noted in the introduction, this change was designed to see if savings decisions about abstract money reward were different than viscerally tempting rewards–namely, liquid consumption by thirsty subjects.

To make this reward appealing and limit satiation across the experiment, subjects were asked not to drink for four hours before the experiment began.[23] They also began by eating some salty snacks. Since it took them 45 minutes to read the instructions and to complete two 30-

period lifecycles for money before the beverage lifecycle, they were definitely thirsty by the time they reached the beverage lifecycle. It is likely that they did not satiate during the lifecycle because no subject received more than 350 ml of soda (less than a 12 oz. can) of cola in that lifecycle, subjects would only be able to drink a maximum of 20 ml/period (0.7 oz), and beverage periods were separated by one minute.[24] Subjects were required to drink their entire beverage consumption in that one-minute period (and they always did).

A syringe pump was used to deliver an exact amount of beverage into a cup.[25] If subjects incurred a negative number of points in any period, they incurred a debt of sorts—they would not receive any beverage until that level had been offset by future positive point totals. This debt was "forgiven" at the end of the beverage lifecycle because we could not force subjects to "pay back" the debt by taking away the beverage (as we do in the money lifecycles).

There were two different reward-delivery conditions in the beverage lifecycle. In the immediate condition subjects received their beverage reward right after making their decision. In the delayed condition subjects received their beverage reward (chosen in period t) ten periods after making their decision (in period t+10).[26] Quasi-hyperbolic or present-bias models of time discounting had predicted that subjects would drink more beverage in the early periods of the immediate condition because delayed rewards would be heavily discounted (see our working paper for a more formal explanation). In the delayed condition, immediate choices did not lead to immediate consumption so the present bias term in β-δ discounting would disappear. Intuitively, the delayed condition would provide external self-control that helps β-δ discounters. Subjects should have drunk more overall in the delayed condition if they were quasi-hyperbolic discounters.

Subjects were n=52 Caltech students.[27] Because a single liquid-delivery apparatus was used, experiments were conducted in a single office rather than a computer lab with one subject at a time. As a result, this study was more laborious than most economics experiments (taking about 130 hours of experimenter-subject contact time).

2. Results

2.1 Total beverage awarded

The hyperbolic discounting and dual-self models predicted that subjects in the immediate condition would receive less beverage than in the delayed condition, because they would consume relatively more compared to a total-reward-maximizing optimum in early periods.[28] This prediction is empirically correct (see Table III, row (1)). The immediate-condition subjects drank less total beverage on average (179 ml, std. dev.=84.6) than the delayed-condition subjects (226 ml, std. dev.=79.0). There was substantial variation across subjects, but this difference is significant at conventional levels by one-tailed tests (t-test p=.047, Mann-Whitney rank sum test p=.015).

2.2 Adjusting for skill

Simply comparing total beverages in the immediate and delayed conditions does not control for possible differences in skill or discounting between subjects in those conditions, which could be evidenced by differential performance in the four money lifecycles. To control for these skill differences, we estimate the regression

(1) $\qquad P_{it} = a + b_1 r_1 + b_2 r_2 + b_3 r_4 + b_4 r_5 + b_5 I + e_{it}$,

where P_{it} is the point total for subject i in lifecycle t, r_i is a dummy variable for lifecycle i, and I is a dummy variable for the immediate condition. If immediate consumption triggered overconsumption and poorer savings accumulation, $b_5 < 0$.

14

Notice that point totals can be negative for the beverage lifecycle, but the total ml of beverage consumed cannot be negative. (Subjects could not have been forced to "pay back" liquid once it is consumed.) This constraint is different than for the money rounds because a monetary point debt accumulated in one lifecycle could be offset by other lifecycles (and subjects are aware of this difference in incentive structures). If a large beverage deficit (>350 ml) occurred in an earlier period, subjects would have known that no amount of spending could have erased this deficit. As a result, when subjects had large negative point totals they could have become indifferent about future decisions (their marginal incentive disappears) and produced high negative points. These high deviations occurred disproportionately in the immediate condition, which then greatly overstates b_5 when the dependent variable is points (see Table IV).[29] In order to reduce the effects of these outliers, two alternative regressions were run. In the second specification, each lifecycle money point total was calculated as if it were a beverage lifecycle (i.e., periods with negative utility are ignored). In the third specification, extreme point totals were reduced in magnitude by taking the logarithms of their absolute values with their sign preserved (i.e., the dependent variable is $[|P_{it}|/P_{it}]\ln(|P_{it}|)$.

Table IV shows the results of a random effects regression run on each model. In all three specifications the sign of b_5, the effect of the immediate condition, is negative and significant at p<.05. In fact, these results are stronger in significance than the parametric t-tests reported in Table III, which implies that accounting for individual differences in skill by using the money-lifecycle results actually enhances the significance of the immediate-delayed condition difference (by reducing variation from cross-subject differences in skill or patience.)

These analyses use the overall point totals in the lifecycle. As in study 1, it is also useful to examine *conditional* deviations in each period given decisions in previous periods. For each

period in the beverage lifecycle we calculated the future expected points for that subject resulting from her decision, compared to the future expected points from a conditionally total-reward-maximizing optimal decision in that period. We then converted these amounts to ml of beverage and totaled these values over all thirty periods. Since no subject received more than 350 ml of beverage in the lifecycle or less than 0 ml, we bounded all totals at 350 ml. Row (2) of Table III shows the results. The average total expected beverage loss, in conditional deviation from optimality, was much higher for the immediate condition than for the delayed condition (about twice as high).

2.3 Exploring the time series of overspending in early periods

Appendix I, Figure A.1 shows the average ratios of spending to conditionally optimal spending. The figure confirms that even when conditioning on past decisions, the immediate-condition subjects were spending more in the first five periods. (After that period the higher number of subjects with beverage deficits and large habits in the immediate condition pushed down their overspending.)[30] Another diagnostic statistic is the average overspending in those periods in which subjects overspent compared to the conditional optimum. The immediate group subjects actually made somewhat fewer overspending decisions than the delayed-condition subjects (41% vs. 51% of decisions),[31] but when they had overspent, the immediate condition subjects spent much more than was optimal (Table III, row (3)), which created greater expected losses.

V. Estimating Quasi-Hyperbolic Discounting Parameters

The results presented to this point have supported the basic prediction of the hyperbolic discounting and dual self models, that subjects in the immediate condition would consume less overall. Because the hyperbolic model is clearly parameterized, we can also estimate best-fitting

16

values of the parameters $\tilde{\delta}$ and $\tilde{\beta}$ from savings decisions and compare those values to estimates from other studies. The analysis is restricted to observations when subjects did not encounter beverage deficits. When a subject encountered a beverage deficit, their decision could only be made to receive *future* rewards and so $\tilde{\beta}$, the immediate bias term, should not apply even in the immediate-beverage condition.[32]

In the quasi-hyperbolic model, the weights placed on immediate and future rewards are 1, $\beta\delta$, $\beta\delta^2$, … $\beta\delta^t$ ... If δ is close to one the terms δ^t are close in numerical value, so there will be many combinations of (β, δ) values which produce similar sequences of weights and similar choices. It is therefore difficult to estimate the two parameters separately. When (β, δ) were maximized simultaneously, the analysis often yielded toward 0 or above 1. We therefore use a two-stage procedure to calibrate δ and β for each subject.

Since behavior in the delayed condition gave no information about the present bias β, in theory, the delayed-condition data is used to estimate δ. So we first search for best-fitting values of δ_D which explain delayed-condition subject choices as if they were maximizing discounted expected utility of consumption with a discount rate δ_D and $\beta=1$. These estimates minimize the sum of squared percentage deviations between the actual consumption and the consumption predicted by the model. This estimation gives a distribution of δ_D estimates with a mean of 0.904 and standard deviation, across subjects, of 0.230. This mean value is reasonable but is significantly less than one at the 2% level by a cross-subject t-test (see Table V). (Note that a discount factor around 0.9 is more plausibly interpreted as a reduced-form expression of suboptimal choice rather than true time preference for these short-horizon experiments.)

The next challenge is to estimate β_I values in the immediate condition, using reasonable values of δ.[33] The procedure we use first fixes $\beta=1$ for each immediate-condition subject and then

estimates a best-fitting value of δ_I for each of those subjects. These values are shown in Appendix I, Table A.1; the mean δ_I is 0.85 and the standard deviation is 0.24. Since we are fixing $\beta=1$, but we believe the actual β_I values might be below 1, we need to adjust the δ_I values in some way that permits more precise estimation of β. We do this by projecting the subject-specific values of δ_I onto the value of the distribution of δ_D estimated from the delayed-condition subjects, which has the same standardized deviation. That is, a specific immediate-condition estimate δ_I is adjusted to an estimate δ_I* where $(\delta_I*-0.904)/0.23=(\delta-0.85)/0.24$. This procedure permits individual differences in δ_I values, but yokes their distribution to the distribution of δ_D values to permit better identification of β. Using these adjusted values of δ_I* for each immediate-condition subject, we then estimate β_I for each subject.

There is one further complication. In quasi-hyperbolic models, people can be either sophisticated or naïve (e.g., O'Donoghue and Rabin, 1999). Sophisticated subjects discount delayed payoffs steeply but understand that in the future they will discount steeply too. Naïve subjects discount steeply but believe, mistakenly, that their current discount factors applied to future periods will also be applied to later decisions.

The difference between sophistication and naïveté can be illustrated in a three period example. In the first period, both types of subjects apply weights 1, $\beta\delta$ and $\beta\delta^2$ to the three periods. However, the sophisticated subject knows that the discount rates 1 and $\beta\delta$ will actually be applied to periods 2 and 3 when period 2 decisions are made, and accounts for this weighting in forecasting period 2 and 3 choices. The naïve subject thinks the discount rates $\beta\delta$ and $\beta\delta^2$ will be used in period 2 to weight period 2 and period 3 utilities; since the $\beta\delta$ term will divide out in optimization, the naïve subject therefore thinks the relative weights applied in periods 2 and 3

will be 1 and δ (i.e., the naïve subject thinks he will act like an exponential discounter in the future).

In simple choice experiments these two behavioral assumptions are difficult to distinguish empirically, but our 30-period experiment gives some empirical leverage for distinguishing them. We therefore estimate β values (using the adjustment procedure described above) assuming both sophisticated and naïve forecasting of future behavior (see our working paper for details).

The results are summarized in Table V. The estimates of β in both the sophisticated and naïve models are clustered around 0.6-0.7. Appendix I, Table A.1 shows individual subject results;[34] all but one subject's estimate is below 1 for both specifications, so the hypothesis that there is no present bias (β=1) is strongly rejected.[35] The estimates of $\tilde{\beta}$ are in the ballpark of estimates of Angeletos et al. (2001) ($\tilde{\beta}$=0.55), Fang and Silverman (2004) ($\tilde{\beta}$=0.69), Shui and Ausubel (2004) ($\tilde{\beta} \approx 0.8$), DellaVigna and Paserman (2005) ($\tilde{\beta}$=0.9), Tanaka, Camerer, and Nguyen (2006) ($\tilde{\beta}$=0.74-0.89), and Laibson, Repetto, and Tobacman (2007) ($\tilde{\beta}$=0.7), (from macroeconomic calibration, welfare take-up, credit card use, unemployment spells, experiments in Vietnam, and consumption data, respectively). The values are also close to other experiments with a much different design that used juice and water rewards ($\tilde{\beta}$=0.52, [McClure et al., 2004]).

Measured by the sum of squared deviations, the naïve model fits better in 16 of 26 subjects. Since this structure is not deliberately designed to distinguish the two specifications, this is just a clue that both specifications should be taken seriously as explanations of behavior in future work.

VI. Applying the Quasi-Hyperbolic Model to Bounded Rationality and Social Learning

Study 1 and study 2 both feature decisions made by subjects in the same experimental framework, so it is useful to have a unified approach to compare the results parametrically. In the previous section, we estimated the results in study 2 using the quasi-hyperbolic model because that model has been designed and commonly used to calibrate the tradeoffs between immediate and delayed rewards over time, and used to model temptation. Any reasonable unified model must accommodate the empirical immediate-delayed difference and the β-δ model is one way to do so. However, it has never been used to represent bounded rationality or the effects of private and social learning on decisions

One approach to creating a unified model is to estimate β-δ parameter values for the private and social learning effects. Since learning means, empirically, saving more at the beginning, in the β-δ framework learning is expressed as a change in these preference parameters.[36] While this is clearly a reduced-form approximation, it is the approach we take below. We return to the details after discussing why other approaches are not likely to fit these data any better. Of course, future research should certainly tackle the problem of developing a more sensible and unified approach, and designing the best experiments to test it.

Two other ways to potentially model private and social learning involve limited planning horizons and rules of thumb.

Limited planning horizons: Ballinger et al. (2006) find that their data are reasonably explained by a model in which subjects tend to only think ahead three periods. A more general version of this approach is a model in which people optimize but act as if only K periods remain (Ballinger et al. estimate K=3). Note that full optimality is K≥30 (in these 30-period experiments) and in beginning period T-K subjects fully optimize (conditionally). Empirically,

this model cannot explain all our data. The period at which consumption becomes conditionally optimal can be used to approximate K, because in period T-K people will begin to (conditionally) optimize. Figure III suggests K is around 10 because their decisions are conditionally optimal beginning around period 20. But a model with a horizon K=10 predicts that subjects will consume everything in the delayed condition of study 2 (since consumption is delivered 10 periods later in the delayed condition). This appealing model cannot easily account parametrically for both the oversavings in study 1 (which implies K around 10) and the fact that there is a limit on consumption in the 10-period-delayed condition of study 2. To be clear, our view is that the truncated-horizon model is a very plausible one, but it just does not do well in explaining the central empirical features of both of our studies.

Rules of thumb: Another approach to model bounded rationality and observed learning is that consumers use a rule of thumb which is adjusted by experience (e.g., Cochrane, 1989). Two plausible rules are consuming a constant fraction of current income or a constant fraction of accumulated cash-on-hand. Neither model fits our experimental data especially well.

The actual consumption-to-income ratio does not exhibit a trend across periods, but fluctuates wildly across periods (see Appendix II for details). For example, in the first lifecycle of the private learning condition, in half the periods the propensity to consume out of current income is less than 1, but in six of 30 periods it is above 2 (i.e., subjects spend all the current income and also dip into savings, because current income is too low to produce an adequate consumption utility). Subjects seem to have some intuitive ability, even in the first lifecycle, to adjust spending from current income to smooth consumption across periods so a simple rule-of-thumb model is strongly rejected.[37]

The actual consumption-to-cash-on-hand ratio also exhibits little trend before learning takes place, but also fluctuates substantially (in six periods it is around 0.4, and in seven periods it is around 0.7, in the first private learning lifecycle). Learning does change this ratio so it looks like a quadratic polynomial across periods after learning takes place.

The most promising approach to modeling bounded rationality is a propensity to consume policy function, which is a low-order polynomial in the state variables (including cash-on-hand and the number of periods remaining), which adjusts with experience (e.g. Houser, Keane, and McCabe, 2004). Adjusting from experience is not so straightforward, however, because each lifecycle only provides one observation on performance of a particular cross-period policy. It is not clear how to adjust a polynomial policy across 30 periods from a single observation on the entire dynamic policy's performance rapidly enough to match the human learning we observe (see the well-known "credit assignment" problem in learning of dynamic policies, e.g., Holland [1985]).

Furthermore, even if we had an ideal model of variation and learning across rule-of-thumb policies that could explain the observed learning in experiment 1, these models are not likely to explain the immediate-delayed condition effect in study 2.

Therefore, we use the β-δ approach as a benchmark unified model (since it is the most natural way to explain the immediate-delayed difference). We treat estimated changes in those parameters as very reduced-form expressions of learning as expressed through these parameters.

Because of the problem in separately identifying β and δ in this design, *ad hoc* methods are used to first identify δ, then estimate β given the estimates of δ. We use the same technique as in study 2.

First consider private learning. We assume that β=1 in lifecycle 7, estimate δ values from those subjects in lifecycle 7 (the mean estimate is 1.00), then apply those estimates of δ (using the same standardization procedure as before) to estimate β in lifecycles 1 and 4. These numbers indicate the strength of learning, if learning is assumed to only change β and not affect δ. These results are shown in Table VI. The estimates of first- and fourth-lifecycle β are 0.415 and 0.778 for the naïve model and 0.273 and 0.585 for the sophisticated model. These numbers are a crude indication of the size of the suboptimality in consumption; they suggest that about half the gap between the first lifecycle β and the value of one is closed by lifecycle 4.

For social learning, we use the same procedure. First β=1 is assumed in lifecycle 7 and an estimate of δ is derived from the lifecycle 7 data (the mean estimate is 1.03), which are then applied when estimating β in lifecycles 1 and 4. The resulting estimates are 0.671 and 0.974 for the naïve model and 0.421 and 1.025 for the sophisticated model.

Together, all these figures give us a simple parametric index of the strength of learning, when learning is parameterized by a change in the immediacy preference β–more appropriately, a reduced-form proxy for undersaving. Assuming β=1 in the last lifecycle (to permit identification), social learning increases β estimates substantially in the first lifecycle compared to private learning (from 0.415 to 0.671, assuming naiveté, or 0.273 to 0.421 assuming sophistication). The learning is apparently much more rapid in social learning, as well, since the estimates of β are very close to one even in lifecycle 4 but are still far from one in private learning.

VI. Conclusions

Dynamic choice models in which current choices influence future constraints or utilities are computationally difficult. Savings in the presence of income uncertainty and habit formation are an example of choice models in this class, which are especially relevant in the economy.

Empirical evidence on savings suggests people are not always saving optimally (though many studies are consistent with some features of optimal savings). However, tests with field data depend sensitively on assumptions about expectations, separability of consumption, and other unobservables. Experiments control for these assumptions. Simple experiments done by others generally show that experimental subjects save too little. Our goal in this paper was to extend this research to a more complicated (and lifelike) environment that combines income uncertainty and habit formation.

The paper's two studies examine two explanations for undersavings in dynamic optimization problems. The first is bounded rationality and the second is an immediacy preference (i.e., even if people knew the optimal saving rules they could not execute them). Both explanations are tested in two experimental studies using a common design and pool of highly skilled subjects. The β-δ model measures undersaving across both studies.

We find that subjects saved much too little at first, but learned to save close to optimal amounts after three or four lifecycles of direct experience (private learning). Furthermore, subjects who have received social learning, examples of successful, unsuccessful, and average experimental performance, produced savings decisions that are quite close to optimal even in their first lifecycle. Since consumers are limited to one lifecycle of private learning (absent reincarnation with memory), it would be interesting to know what types of social learning are more effective. Does social learning work better when it comes from family and friends, from

total strangers, from financial planners, or from training and education? Our data suggest that one type of social learning works well but invites consideration of other forms which can be tested in future experiments and in field data.

The fact that subjects could learn to save optimally for money rewards led us to explore whether they saved optimally when rewards are more immediate and visceral–when thirsty subjects' rewards were immediate sips of a cola beverage. The subjects who sipped the beverage immediately also overspent (i.e., overdrank), compared to group of subjects who made decisions in one period but did not get to sip that period's beverage amount until ten periods later. As a result of their overspending, subjects in the immediate-reward condition earned less total rewards than those in the delayed condition, and received less than the theoretical, total-reward-maximizing optimum. This unique feature of our second study provides a model for future studies of highly tempting decisions like addiction, overeating, and perhaps spending splurges.

The difference in the performance of the immediate and delayed conditions is consistent with the predictions of both the quasi-hyperbolic and dual self models, and is not consistent with the standard exponential model. When parameters of the quasi-hyperbolic model are calibrated from subject decisions in the immediate condition, the mean best-fitting $\tilde{\beta}$ (the degree of present bias) is 0.62 for the sophisticated case and 0.72 for the naïve case. These values are close to values observed in some other studies using both calibrations to aggregate data and direct experimental measurement. Parameter estimates using beta as a representation of inexperience are much lower than those observed in field data and experiments, which allow immediacy preference. Although this model was not intended to be applied to measure bounded rationality in a reduced-form way, the lower β suggest that if consumers were inexperienced and did not know how much to consume, they would be much more impatient than has been inferred by

Angeletos et al. (2001) and Laibson, Repetto, and Tobacman (2007) from aggregate savings and investment.

There are many directions for future research. The experimental paradigm could also be extended by adding more lifelike features, such as stochastic mortality, retirement, and supply-side advice that either tempts subjects more or gives them good advice. The fact that subjects in the delayed condition are able to resist temptation better (and drink more total beverage as a result) corroborates the conclusion of models like Bernheim and Rangel's (2004), that creating a time wedge between "ordering" and consuming may be helpful to people. This observation suggests an experimental way to measure demand for external self-control. The immediate-condition subjects are making a mistake, but they can't help doing so. If they had access to external commitment, sophisticated hyperbolics would seek external commitment. Future experiments could allow subjects in beverage studies the choice between whether they want to participate in the immediate or delayed condition; sophisticated subjects should opt for the imposed delay. Naïve hyperbolics and exponential discounters would be indifferent about both conditions. An alternative theory (Gul and Pesendorfer, 2001) suggests that agents might prefer the delayed condition if it reduces disutility from temptation.

The natural question about experiments of this type is how well their results generalize to naturally occurring savings by different groups of people. The experiments inevitably reflect the classic trade-off in generalizing from stylized lab experiments to naturally occurring choice: Experiments have high "internal validity" because the maintained assumptions of a particular theory (e.g., about utility functions and beliefs about the income process) can be clearly implemented with experimental control. Experimental comparison of different treatments can also shed some light on competing explanations (that is the potential advantage of the

immediate-delayed consumption comparison in our second study). However, the generalizability of the experiments to most actual savings decisions is debatable because experiments necessarily take place over a very short horizon and the saving lifecycle is long. These experiments, much like a time-lapse photograph, show a process in a short amount of time (a few hours) that usually spans a much longer time interval (many years). Therefore, while we discuss the results in terms of their relevance to debates about lifecycle savings, readers are entitled to think of the data as more generalizable about much shorter-term dynamic decisions like consumption of addictive substances, or taking up an exercise regimen or diet.

Nonetheless, consumers may be making some lifecycle decisions in a way that is similar to the experimental protocol. Financial planning software exists (e.g., www.financialfate.com), similar to that used in this experiment, which allows consumers to plan their savings over a short-term horizon. The software allows them to experiment with different assumptions and see what results occur. So people using such software might be making decisions in short periods of time, which are similar in time scale to the experimental decisions, even though the software-guided decisions have long-run consequences for many years.

Additionally, while economic agents cannot experience more than one lifecycle, they can learn from the savings success and mistakes of others. Retirement advisors may exist because individuals are unable to make retirement decisions in one lifecycle, but can make good decisions after observing multiple lifecycles (and those histories are bottled and sold by advisors) and with formal tools to analyze and explain what to do. The market may have solved the cognitive problem in savings models by producing a supply of helpful retirement advisors. Alternatively, retirement advisors may offer products allowing a cross-subsidization between inexperienced and experienced consumers, causing the decisions of the most inexperienced

consumers to become more suboptimal (an idea discussed in Campbell [2006]). These phenomena can be studied in experiments too, by allowing markets for advice and group-level decisions (e.g. household saving) to see whether these institutions lead toward or away from optimal choice.

Alexander L. Brown

Graduate Student
Division of Humanities and Social Sciences
California Institute of Technology
Pasadena, CA 91125
abrown@hss.caltech.edu

Zhikang Eric Chua

Singapore Public Service Commission Scholar
eczk@singnet.com.sg

Colin F. Camerer

Robert Kirby Professor of Behavioral Economics
Division of Humanities and Social Sciences
California Institute of Technology
Pasadena, CA 91125
camerer@hss.caltech.edu

References

Ainslie, George, "Specious Reward - Behavioral Theory of Impulsiveness and Impulse Control," *Psychological Bulletin*, 82 (1975), 463-496.

Allen, Todd W., and Christopher D. Carroll, "Individual Learning about Consumption," *Macroeconomic Dynamics*, 5 (2001), 255-271.

Anderson, Lisa R., and Charles A. Holt, "Information Cascades in the Laboratory," *The American Economic Review*, 87 (1997), 847-862.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 15 (2001), 47-68.

Ballinger, T. Parker, Eric Hudson, Leonie Karkoviata, and Nathaniel T. Wilcox, "Saving Performance and Cognitive Abilities," Stephen F. Austin State University, 2006.

Ballinger, T. Parker, Michael G. Palumbo, and Nathaniel T. Wilcox, "Precautionary Saving and Social Learning across Generations: An Experiment," *Economic Journal*, 113 (2003), 920-947.

Baumeister, Roy F., Todd F. Heatherton, and Dianne M. Tice, *Losing Control: How and Why People Fail at Self-Regulation* (San Diego: Academic Press, 1994).

Becker, Gary S., and Kevin M. Murphy, "A Theory of Rational Addiction," *The Journal of Political Economy*, 96 (1988), 675-700.

Benabou, Roland, and Jean Tirole, "Willpower and Personal Rules," *Journal of Political Economy*, 112 (2004), 848-886.

Benartzi, Shlomo, and Richard H. Thaler, "Heuristics and Biases in Retirement Savings Behavior," *Journal of Economic Perspectives*, 21 (2007), 81-104.

Bernasconi, Michele, and Oliver Kirchkamp, "Why Do Monetary Policies Matter? An Experimental Study of Saving and Inflation in an Overlapping Generations Model," *Journal of Monetary Economics*, 46 (2000), 315-343.

Bernheim, B. Douglas, and Antonio Rangel, "Addiction and Cue-Triggered Decision Processes," *American Economic Review*, 94 (2004), 1558-1590.

Bodner, Ronit, and Drazen Prelec, "The Diagnostic Value of Actions in a Self-Signaling Model," in *The Psychology of Economic Decisions,* I. Brocas, and J.D. Carillo, eds. (USA: Oxford University Press, 2003).

Brown, Alexander L., "Endogenous Time Preference and Personal Rules" in *Investigating Psychology-Influenced Economic Models in Lab, Field, and Theory*, Ph.D. Dissertation, California Institute of Technology, 2008.

Browning, Martin, and Annamaria Lusardi, "Household Saving: Micro Theories and Micro Facts," *Journal of Economic Literature*, 34 (1996), 1797-1855.

Campbell, John Y., "Household Finance," *The Journal of Finance*, 61 (2006), 1553-1604.

Carbone, Enrica, "Demographics and Behaviour," *Experimental Economics*, 8 (2005), 217-232.

Carbone, Enrica, and John D. Hey, "The Effect of Unemployment on Consumption: An Experimental Analysis," *The Economic Journal*, 114 (2004), 660-683.

Carroll, Christopher D., "The Buffer-Stock Theory of Saving - Some Macroeconomic Evidence," *Brookings Papers on Economic Activity*, 1992 (1992), 61-135.

---, "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112 (1997), 1-55.

---, "A Theory of the Consumption Function, with and without Liquidity Constraints," *Journal of Economic Perspectives*, 15 (2001), 23-45.

Carroll, Christopher D., Judy Overland, and David N. Weil, "Saving and Growth with Habit Formation," *American Economic Review*, 90 (2000), 341-355.

Carroll, Christopher D., Changyong Rhee, and Byungkun K. Rhee, "Does Cultural Origin Affect Saving Behavior? Evidence from Immigrants," *Economic Development and Cultural Change*, 48 (1999), 33-50.

Choi, James J., David Laibson, and Brigitte C. Madrian, "$100 Bills on the Sidewalk: Suboptimal Saving in 401(k) Plans," Harvard, 2005.

Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Optimal Defaults," *American Economic Review Papers and Proceedings*, 93 (2003), 180-185.

Chua, Zhikang, and Colin F. Camerer, "Experiments on Intertemporal Consumption with Habit Formation and Social Learning," California Institute of Technology, 2004.

Cochrane, John H., "The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives," *The American Economic Review*, 79 (1989), 319-337.

Darlin, Damon, "A Contrarian View: Save Less and Still Save Enough for Retirement." *The New York Times*, January 27.

DellaVigna, Stefano, and M. Daniele Paserman, "Job Search and Impatience," *Journal of Labor Economics*, 23 (2005), 527-588.

Fang, Hanming, and Dan Silverman, "On the Compassion of Time-Limited Welfare Programs," *Journal of Public Economics*, 88 (2004), 1445-1470.

Fehr, Ernst, and Peter. K. Zych, "Do addicts behave rationally?," *Scandinavian Journal of Economics*, 100 (1998), 643-662.

Friedman, Milton, and National Bureau of Economic Research., *A Theory of the Consumption Function* (Princeton,: Princeton University Press, 1957).

Fuchs-Schundeln, Nicola, and Matthias Schundeln, "Precautionary Savings and Self-Selection: Evidence from the German Reunification "Experiment"," *The Quarterly Journal of Economics*, 120 (2005), 1085-1120.

Fudenberg, Drew, and David Levine, "A Dual Self Model of Impulse Control," *American Economic Review*, 95 (2006), 1449-1476.

Gale, William G., John Sabelhaus, and Robert E. Hall, "Perspectives on the Household Saving Rate," *Brookings Papers on Economic Activity*, 1999 (1999), 181-224.

Gul, Faruk, and Wolfgang Pesendorfer, "Temptation and Self-Control," *Econometrica*, 69 (2001), 1403-1435.

Haveman, Robert, Karen Holden, Barbara Wolfe, and Shane Sherlund, "Do Newly Retired Workers in the United States Have Sufficient Resources to Maintain Well-Being?," *Economic Inquiry*, 44 (2006), 249-264.

Hey, John D., "A Pilot Experimental Investigation into Optimal Consumption under Uncertainty," in *Applied Behavioural Economics,* Shlomo Maital, ed. (New York: New York University Press, 1988).

Hey, John D., and Valentino Dardanoni, "Optimal Consumption Under Uncertainty: An Experimental Investigation," *The Economic Journal*, 98 (1988), 105-116.

Holland, John H., "Properties of the Bucket Brigade Algorithm," in *Proceedings of the First International Conference on Genetic Algorithms and their Applications,* John J. Grefenstette, ed. (Hillsdale, NJ: Lawrence Erlbaum, 1985).

Houser, Daniel, Michael Keane, and Kevin McCabe, "Behavior in a Dynamic Decision Problem: An Analysis of Experimental Evidence Using a Bayesian Type Classification Algorithm," *Econometrica*, 72 (2004).

Kotlikoff, Laurence J., Stephen Johnson, and William Samuelson, "Can People Compute? An Experimental Test of the Life-Cycle Consumption Model," in *Essays on Savings, Bequests, Altruism and Life-Cycle Planning,* Laurence J. Kotlikoff, ed. (Cambridge, MA: MIT Press, 2001).

Kable, J. W., and P. W. Glimcher, "The neural correlates of subjective value during intertemporal choice," *Nat Neurosci*, 10 (2007), 1625-1633.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 112 (1997), 443-477.

Laibson, David, Andrea Repetto, and Jeremy Tobacman, "A Debt Puzzle," in *Knowledge, Information, and Expectations in Modern Macroeconomics : in honor of Edmund S. Phelps,* Edmund S. Phelps, and Philippe Aghion, eds. (Princeton: Princeton University Press, 2003).

---, "Estimating Discount Functions with Consumption Choices over the Lifecycle," Harvard University, 2007.

Loewenstein, George, and Ted O'Donoghue, "Animal Spirits: Affective and Deliberative Influences on Economic Behavior," Carnegie Mellon University, 2004.

Lusardi, Annamaria, and Olivia S. Mitchell, " Financial Literacy and Retirement Preparedness: Evidence and Implications for Financial Education," *Business Economics*, 42 (2007), 35-44.

---, "Baby Boomer retirement security: The roles of planning, financial literacy, and housing wealth," *Journal of Monetary Economics*, 54 (2007), 205-224.

Lusardi, Annamaria, Jonathan Skinner, and Steven Venti, "Savings Puzzles and Savings Policies in the United States," *Oxford Review of Economic Policy*, 17 (2001), 95-115.

McClure, Samuel. M., Keith. M. Ericson, David I. Laibson, George Loewenstein, and Johnathan D. Cohen, "Time Discounting for Primary Rewards," *Journal of Neuroscience*, 27 (2007), 5796-5804.

McClure, Samuel M., David I. Laibson, George Loewenstein, and Jonathan D. Cohen, "Separate Neural Systems Value Immediate and Delayed Monetary Rewards," *Science*, 306 (2004), 503-507.

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

Patton, Jim H., Matthew S. Stanford, and Ernest S. Barratt, "Factor Structure of the Barratt Impulsiveness Scale," *Journal of Clinical Psychology*, 51 (1995), 768-774.

Plott, Charles R., and Shyam Sunder, "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models," *The Journal of Political Economy*, 90 (1982), 663-698.

---, "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," *Econometrica*, 56 (1988), 1085-1118.

Prelec, Drazen, and George Loewenstein, "The Red and the Black: Mental Accounting of Savings and Debt," *Marketing Science*, 17 (1998), 4-28.

Rick, Scott I., Cynthia E. Cryder, and George Loewenstein, "Tightwads and Spendthrifts," *Journal of Consumer Research*, 34 (2008), 767-782.

Scholz, John Karl, and Ananth Seshadri, "Children and Household Wealth," SSRN, 2007.

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun, "Are Americans Saving "Optimally" for Retirement?," *Journal of Political Economy*, 114 (2006), 607-643.

Shiv, Baba, and Alexander Fedorikhin, "Spontaneous Versus Controlled Influences of Stimulus-Based Affect on Choice Behavior," *Organizational Behavior and Human Decision Processes*, 87 (2002), 342-370.

Shui, Haiyan, and Lawrence M. Ausubel, "Time Inconsistency in the Credit Card Market," SSRN, 2004.

Skinner, Jonathan, "Are You Sure You're Saving Enough for Retirement?" *Journal of Economic Perspectives*, 21 (2007), 59-80.

Strahan, Erin J., Steven J. Spencer, and Mark P. Zanna, "Subliminal Priming and Persuasion: Striking While the Iron Is Hot," *Journal of Experimental Social Psychology*, 38 (2002), 556--568.

Tanaka, Tamomi, Colin F. Camerer, and Quang Nguyen, "Poverty, Politics, and Preferences: Experimental and Survey Data from Vietnam," California Institute of Technology, 2006.

Venti, Steven, "Choice, Behavior and Retirement Saving," in *Oxford Handbook of Pensions and Retirement Income,* G. Clark, A. Munnell, and M. Orzsag, eds. (Oxford: Oxford University Press, 2006).

Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *The Quarterly Journal of Economics*, 104 (1989), 275-298.

[13] Since $\rho=3$, the term k is the upper asymptote of utility. θ is a scaling parameter, and $\hat{\varepsilon}$ bounds the utility function from below. In the experiments, $\hat{\varepsilon}$ = 2.7, similar to Ballinger, Palumbo, and Wilcox (2003). Scaling factors are $\theta = 750$ and $k=40$.

[14] The exchange rates were US \$1.50 for every 100 experimental points in Caltech, and US \$2.50 in Singapore (using an exchange rate of US \$1 $\approx$ Sing \$1.70).

[15] The tables looked like the screens the participants had, showing income each period, cash-on-hand, spending decisions, and points from each period of a 30-period lifecycle.

[16] Ex post we know that the highest scoring subject underconsumed for the first ten periods, but then overconsumed for the remainder. The subject was lucky to draw a high-income realization in the later periods so he could still produce positive utility under a high level of lifestyle habit.

[17] The income realizations were different so that the social learning subjects would never have a lifecycle that matched exactly the income realizations seen by the role model subjects (drawn from the private learning condition).

[18] Note that in some cases, the average subject does *better* than the unconditional or conditional optimum (i.e., the deviation from optimality is positive). This can happen if participants overspend (underspend) but get lucky (unlucky) and have good (bad) income shocks in later periods.

[19] It should be noted that in both conditions lifecycle 5 featured the lowest total income (the harshest income draws). In condition 1 it managed to cause the subjects and the ex ante optimal path to have negative utility. In condition 2 it only reduced the utility of the subjects.

[20] See Chua and Camerer (2004) for details. Ethnicity is of interest because Singaporean Chinese have one of the highest savings rates in the world (see Carroll, Rhee and Rhee, 1999). Participant random effects were also included to control for individual differences, which are substantial. In a broader specification a Caltech dummy variable was also included but is insignificant and is dropped. The Chinese dummy variable is correlated with subject pool, but not strongly. There are many ethnic Chinese students at Caltech, and Singaporean students are not exclusively Chinese.

[21] One reason to keep the complex design with habit formation and stochastic income was because behavioral research suggests that higher cognitive loads make people more likely to succumb to visceral temptation (Shiv and Fedorikhin, 2002). Additionally, subjects are more likely to succumb to temptation if they are unaware they are doing so (Baumeister, Heatherton, and Tice, 1994) or if the signals of doing so are noisy (Bodner and Prelec, 2003; Benabou and Tirole, 2004).

[22] Subjects were given their preference of Coke or Pepsi, and could substitute Diet Coke or Diet Pepsi if they requested it. We used these beverages because they are widely valued, water was as motivating as colas, and because pilot subjects (including the senior coauthor who was a pilot subject) thought fruit juices that were tried were likely to induce satiation, which complicates the analysis.

[23] There is no way to know whether all subjects obeyed our request to show up thirsty. However, because assignment to the immediate and delayed conditions did not depend upon apparent thirst, uncontrolled and unmeasured differences in pre-experiment thirst are sources of sampling error in comparing the two groups which lower the power of the test and bias the test against finding a difference between the immediate and delayed conditions.

[24] The concavity of utility and properties of the buffer stock savings model ensure that no subject could earn more than 700 points in any beverage or monetary lifecycle.

[25] See our working paper for a diagram of the beverage delivery apparatus.

[26] To standardize both conditions completely there were forty periods of one minute each in lifecycle 3. In the immediate condition, subjects did nothing in the last ten periods. In the delayed condition, subjects made decisions in the first ten periods, but received no rewards. In the last ten periods of that condition subjects received their rewards from periods 21-30 but made no decisions.

[27] The first 44 subjects were run from April 21 to July 27, 2005. After that, 11 more subjects were run from February 7-16, 2006 to enlarge the sample and check robustness of the result. Two subjects refused to drink during the beverage period and were dropped from the analysis. Another subject's data were lost by mistake.

[28] An alternative explanation is that the first taste of cola primed subjects to consume more. Since immediate subjects first received cola after period 1 and delayed subjects first received cola after period 11, this priming could be responsible for the difference in total rewards. While thirst priming has been studied in psychology through subliminal means (e.g., Strahan et al., 2002), to our knowledge no psychological work has studied or found evidence of this specific type of priming. Further, the data finds an average consumption increase (11.7 vs. 14.2 ml) between periods 1 and 2 for the immediate condition, but an average consumption decrease (10.4 vs. 8.2 ml) for the delayed.

The jump between periods 1 and 2 in the immediate condition is most likely explained by a very high income draw in period 2 (3.56) that on average, sextupled cash-on-hand (76.6 vs. 429.1).

[29] Since subjects know they will not be forced to pay back previously consumed beverage, it is conceivable that they exploit this design property by deliberately overconsuming in early periods and then running up point debts they do not have to pay. We do not estimate such a model because the period-specific maximum of liquid consumption is 20 ml per period, so the marginal beverage value of increased consumption falls sharply. As a result, subjects who are trying to optimize total liquid would smooth consumption and would never deliberately run up a debt. In terms of our estimation below, a deliberate strategy of overconsuming because of anticipated "bankruptcy" would be misclassified as a low value of the discounting parameter δ. There is no a priori reason to think this pattern will be more common in the immediate and delayed conditions if both types of subjects have similar discounting patterns.

[30] Immediate subjects have more beverage deficits (4 subjects vs. 1 in period 6; 15 vs. 8 by period 10) and higher average habit levels accumulated (218 vs. 185 in period 6) than the delayed condition. It is not the case that the immediate subjects have satiated on soft drinks compared to the delayed group, because the immediate subjects have only drunk about 57 ml (2 oz) on average after five periods.

[31] Periods in which a subject encountered a deficit of 20 ml or greater were omitted in this analysis.

[32] Subjects with beverage deficits that were large enough knew they would not receive liquid again and therefore have no incentive to choose one spending decision over another. While some subjects never encountered a beverage deficit, and others encountered them early, each subject was given a single parameter value and the results were analyzed so that each subject's value counts as much as any other.

[33] Using the mean of the delayed-condition estimates δ_D and estimating subject-specific β_I works poorly because differences in δ values for those subjects from the mean δ_D leads to implausible variation in estimates of β_I. The problem with using the delayed-condition mean δ_D for the immediate-condition subjects is the following: suppose an immediate-condition subject's δ is smaller than the mean δ_D. Then the best-fitting sequence of weights $1, \beta\delta, \beta\delta^2, \ldots \beta\delta^t$ will overestimate β because the β parameter is forced to pick up the slack for the under-estimated δ. Similarly, if the immediate-condition δ is below the mean δ_D, β will be underestimated. Indeed, when we tried this procedure the estimate of β tends to bifurcate to the lower and upper bounds placed on β.

[34] A possible correlate of individual β values are subject values on the Barratt Impulsivity Scale. After subjects had completed their experimental session, they answered a survey measuring their total "impulsivity" on the BIS 11 Barratt Impulsivity scale (Patton, Stanford, and Barratt, 1995). However, these values show little correlation with the individual naïve betas, sophisticated betas, and subject performance (correlations smaller than .1 in absolute value).

[35] The correlation of β and δ estimates across subjects is around .35 for both specifications of β, so there is no serious identification problem.

[36] One major reason why the quasi-hyperbolic model has not been used to model bounded rationality is that β-δ are assumed to be static parameters. Any change to bounded rationality (i.e., learning) implies that one or more of the parameters must change, violating that assumption. Relaxing this assumption in a theoretical model of temptation, Brown (2008) finds that a dynamic β term can explain motivations behind preserving internal precedents and personal rules.

[37] In our working paper we regressed consumption against conditionally optimal consumption and rule-of-thumb spending of a constant percentage of current income. The latter term has essentially no statistical weight. Our interpretation is that while subjects are not exactly optimizing (they clearly undersave in early periods with only private learning), the variation across the 30 periods is much better picked up by variation in optimal consumption than by a constant rule of thumb.

Table I: Summary statistics of actual point outcomes across 7 lifecycles, study 1

lifecycle	1	2	3	4	5	6	7
private learning							
mean points	118	-53	224	450	-65	435	440
standard deviation	635	694	498	297	475	255	146
deviation from optimum	*-453*	-628	-349	-125	11	-153	-149
total income	5471	7083	5215	6235	4300	4571	4789
social learning							
mean points	325	586	559	589	309	539	504
standard deviation	238	54	93	62	255	73	47
deviation from optimum	-215	-68	-69	-66	-220	-66	-49
total income	4342	5416	5224	5901	4193	5344	5050

Note: This table provides the average performance measures for each condition and lifecycle in the first study of the experiment. "Mean Points" is the subject-averaged total number of utility points ($\sum_{t=1}^{T} u(C_t, H_{t-1})$) received in a given lifetime. "Deviation from Optimum" ($\sum_{t=1}^{T} u(C_t, H_{t-1}) - \sum_{t=1}^{T} u(C_t^*, H_{t-1}^*)$) is the subject-averaged "Mean Points" minus the total utility that would have been earned following the ex ante optimal path and "Total Income" is the total income received over all thirty periods of the lifecycle ($\sum_{t=s}^{T} X_t$); it indicates the severity of each lifecycle's income shocks.

Table II: Regression of log(absolute deviation from conditional optimum)

	Dependent variable: Logged absolute deviation from conditional optimum
social learning	-0.24* (-2.51)
lifecycle 2	0.092* (2.30)
lifecycle 3	-0.027 (-0.67)
lifecycle 4	0.075 (1.86)
lifecycle 5	-0.43** (-10.69)
lifecycle 6	-0.063 (-1.58)
lifecycle 7	-0.17** (-4.21)
period	0.084** (15.91)
period squared	-0.00034* (-2.01)
female	0.19* (1.99)
Chinese	0.0006 (0.01)
constant	0.77** (16.39)
R²	0.20

Note: This table displays a regression of log(absolute deviation from conditional optimum), which is the absolute difference in subject spending and each period's conditional optimum. Standard errors are given in parenthesis. Independent variables include "social learning," which is 1 if the condition is social learning and 0 otherwise. Dummy variables for each lifecycle (lifecycle 1 is omitted so a constant may be included) and period and period squared show the effects of learning over time and having fewer future periods to plan, respectively. Dummy variables for demographics (Chinese=1, female=1) are also included. Asterisks indicate statistical significance (*p<.05; ** p<.01).

Table III: Summary statistics of results in immediate and delayed conditions in the beverage lifecycle, study 2

	Immediate	Delayed	Parametric test	Nonparametric test
total beverage received	176.78 (81.31)	215.65 (82.89)	t=1.71 p=0.047	z=2.09 p=0.018
total expected losses from optimal (bounded at 350 ml)	171.91 (128.13)	96.98 (104.04)	t=2.35 p=0.011	z=2.34 p=0.010
average expected loss from overspending	18.36 (28.78)	6.40 (10.91)	t=1.92 p=0.031	z=1.77 p=0.038

Note: Sample standard deviations are in parentheses below means. All p-values are one-tailed. The nonparametric test used is the Mann-Whitney test.

Table IV: Regression of lifecycle number and condition on three measures of subject performance, study 2

	Points	Beverage	Sign-preserved Log Points
immediate condition (I)	-28137.36** (7284.97)	-39.58* (20.39)	-4.21** (1.43)
r_1	-171.61 (6264.08)	-20.78 (16.77)	-2.30 (1.20)
r_2	-458.17 (6264.08)	-57.27** (16.77)	-3.73** (1.20)
r_4	105.10 (6264.08)	44.28** (16.77)	1.14 (1.20)
r_5	-656.71 (6264.08)	-26.28 (16.77)	-6.93** (1.20)
constant	282.30 (5151.29)	215.82** (15.14)	4.60** (1.02)
R^2	0.09	0.16	0.22
N	268	268	268

Note: Three measures of subject performance are points (the total utility obtained in each lifecycle, which is linearly converted to cash in lifecycles 1,2, 4, and 5), beverage (the ml of beverage that would be obtained in each lifecycle if utility were converted to beverage as in lifecycle 3), and sign-preserved points (the log of absolute points obtained, keeping negative point totals negative). I is a dummy variable for the immediate condition, r_i is a dummy variable for lifecycle i (lifecycle 3 is excluded because a constant is used). Standard errors are in parenthesizes. Asterisks indicate one-tailed statistical significance (*p<.05; ** p<.01).

Table V: Two-stage parameter estimates of $\widetilde{\delta}$ and $\widetilde{\beta}$ (N=26) for beverage lifecycle only, study 2

Model	Standard	Sophisticated	Naïve
mean $\widetilde{\delta}$ of Delayed (std deviation)	1 n/a	0.904** (0.230)	
mean $\widetilde{\beta}$ of Immediate (std deviation)	1 n/a	0.619*** (0.211)	0.721*** (0.134)
average squared deviations per period, before deficits, immediate subjects only, using mean $\widetilde{\beta}$ and $\widetilde{\delta}$	0.230	0.189	0.193

Note: The value "average squared deviations…" is calculated from the sum of the squared difference between a subject's consumption (measured as a proportion of permanent income) and theoretical consumption (for the mean $\widetilde{\beta}$ and $\widetilde{\delta}$) in each period, before a beverage deficit is encountered (if applicable). It is divided by the number of periods before the beverage deficit to standardize the value over all subjects, and then averaged over all subjects. It provides a fit of each line. Asterisks indicate one-tailed statistical significance testing the null hypothesis that the given parameter is equal to one (*p<.05; ** p<.02, ***p<0.01).

Table VI: Two-stage parameter estimates of $\widetilde{\beta}$ by learning condition (N=36), study 1

Condition	Lifetime 1		Lifetime 4		Lifetime 7
	Naïve	Sophisticated	naïve	sophisticated	
private learning	0.415 (0.339)	0.273 (0.352)	0.778 (0.323)	0.585 (0.450)	1 n/a
social learning	0.671 (0.230)	0.421 (0.424)	0.974 (0.192)	1.025 (0.398)	1 n/a

Note: Standard errors are in parenthesizes. The value $\widetilde{\beta}$ is fixed at 1 for lifetime 7. In that lifecycle, the mean value $\widetilde{\delta}$ =1.00 (0.16) for private learning and mean $\widetilde{\delta}$ =1.03 (0.07) for social learning. The relationship $\widetilde{\beta}$ <1 indicates a tendency to overconsume relative to the reward maximizing optimum.

Figure I: Screenshot of Excel interface

| Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	100.00	1.247	124.67	124.67	10.00	15.00	109.67	21.03	
2	105.00	0.469	49.24	158.91	22.00	25.00	133.91	20.05	
3	110.25	0.138	15.26	149.17	40.40	30.00	119.17	10.31	
4	115.76	0.742	85.88	205.06	58.28	45.00	160.06	18.34	
5	121.55	1.073	130.43	290.49	85.80	72.00	218.49	25.95	
6	127.63	0.354	45.13	263.62	132.06	150.00	113.62	34.36	
7	134.01	0.111	14.86	128.48	242.44	128.48	0.00	24.16	
8	140.71	0.863	121.39	121.40	298.19	121.39	0.01	17.30	
9	147.75	0.509	75.20	75.20	330.12	75.20	0.00	-25.07	
10	155.13	0.496	76.99	77.00	306.29	76.99	0.01	-16.83	
11	162.89	1.135	184.80	184.81	291.39			nil	Next Period
12	171.03								
13	179.59								
14	188.56								
15	197.99								
16	207.89								
17	218.29								
18	229.20								
19	240.66								
20	252.70								
21	265.33								
22	278.60								
23	292.53								
24	307.15								
25	322.51								
26	338.64								
27	355.57								
28	373.35								
29	392.01								
30	411.61								129.61

Figure II: An optimal consumption path
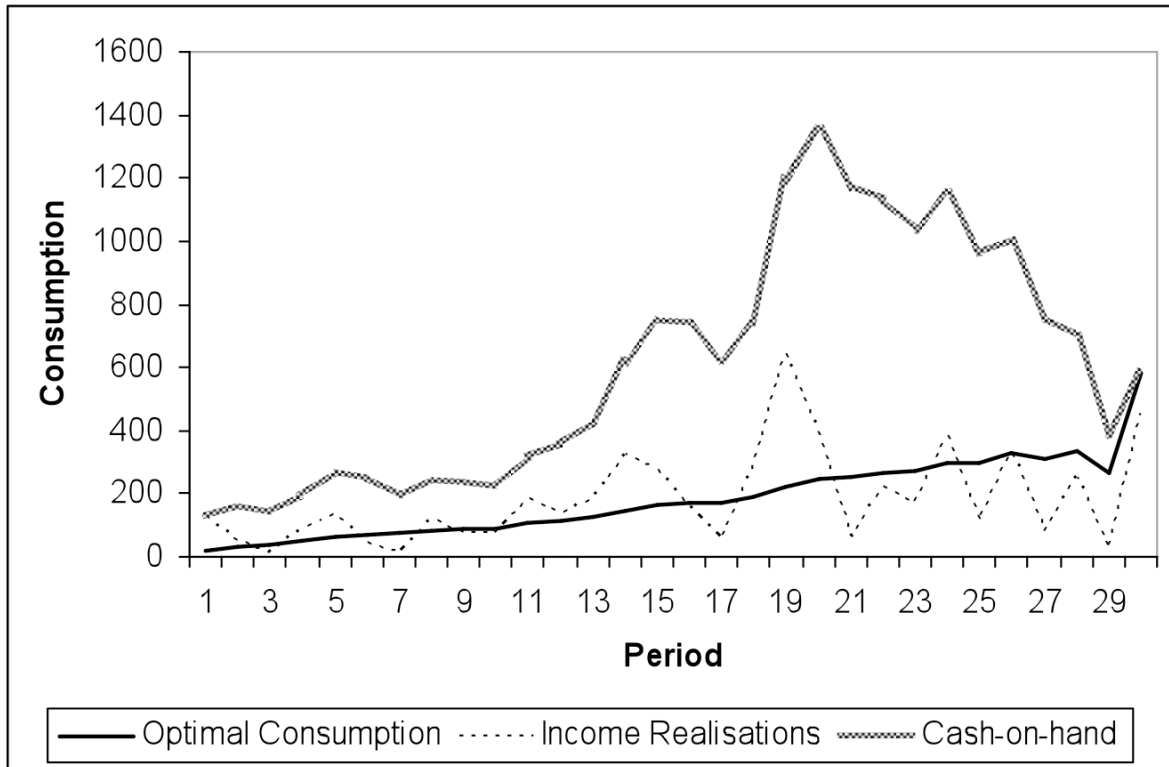
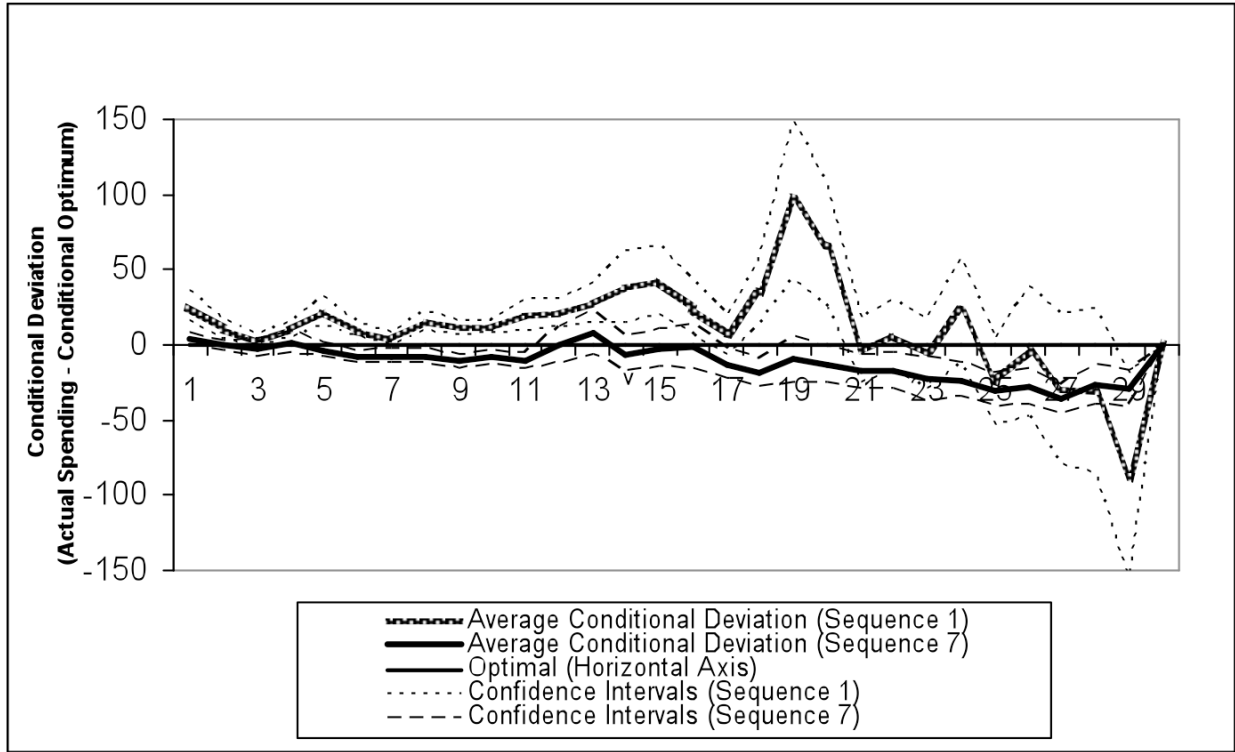Figure III: Deviations from conditional optima, lifecycle 1 and 7, private learning
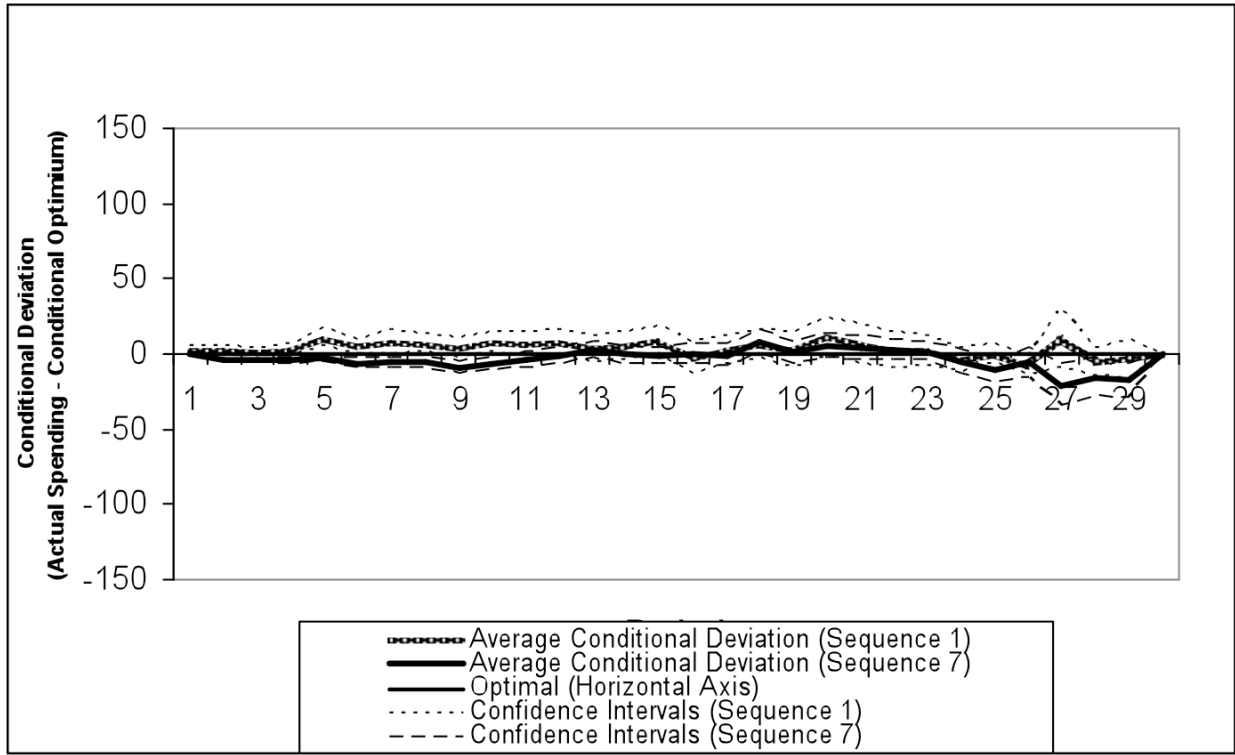

Figure IV: Deviations from conditional optima, lifecycle 1 and 7, social learning

Appendix I: Study 2 supplemental figures and tables

Figure A.1 shows the average ratios of spending to conditionally optimal spending. The figure confirms that even when conditioning on past decisions, the immediate-condition subjects were spending more in the first five periods. (After that period the higher number of subjects with beverage deficits and large habits in the immediate condition pushed down their overspending.)

Table A.1 shows individual subject results for the immediate condition. Subject values were first best fit to a specific $\widetilde{\delta}$ (column (2)), then projected onto to delayed condition results to standardize the values (column (3)). With the standardized $\widetilde{\delta}$, best-fitting $\widetilde{\beta}$ were estimated for both naïve (column (4)) and sophisticated (column (5)) quasi-hyperbolic models. All but one subject's estimate is below 1 for both specifications, so the hypothesis that there is no present bias ($\beta=1$) is strongly rejected.

Appendix II: Spending as proportions of current income and cash-on-hand

These figures show that simple rule-of-thumb models of constant consumption across periods are badly rejected by actual behavior. Figures A.2 and A.3 show that consumption as a proportion of current income fluctuates a lot across periods. Figures A.4 and A.5 show that consumption as a proportion of cash-on-hand fluctuate a lot, as well. However, those Figures show that learning could conceivably be modeled as a shift in the consumption function from one that exhibits little trend across periods (in lifetime 1) to one which is clearly increasing across periods and jumps up sharply in the last couple of periods. (Keep in mind that the software automatically spends all available cash-on-hand in the last period, so the consumption/cash-on-hand ratio is necessarily equal to 1 in the last period 30.)

Table A.1: Estimated β's of individual subjects in immediate condition

Subject	Best fit $\widetilde{\delta}$	Projected $\widetilde{\delta}$ Delayed Condition	Soph $\widetilde{\beta}$	Naïve $\widetilde{\beta}$	Soph SS fit per period	Naïve SS fit per period	Soph fit-Naïve fit
i1	1.07	1.11	1.10	1.19	0.074	0.082	0.008
i2	0.83	0.88	0.32	0.86	0.014	0.022	0.008
i3	0.41	0.48	0.79	0.63	0.393	0.364	-0.029
i4	1.05	1.09	0.89	0.69	0.051	0.051	0.000
i5	1.06	1.10	0.54	0.58	0.005	0.001	-0.004
i6	0.51	0.58	0.49	0.73	0.012	0.010	-0.002
i7	0.87	0.92	0.70	0.74	0.004	0.002	-0.002
i8	1.02	1.07	0.70	0.58	0.492	0.467	-0.025
i9	1.04	1.08	0.57	0.71	0.008	0.012	0.004
i10	0.83	0.88	0.67	0.76	0.004	0.004	0.000
i11	1.01	1.05	0.40	0.82	0.017	0.023	0.006
i12	0.12	0.21	0.08	0.44	0.014	0.001	-0.012
i13	0.70	0.75	0.69	0.75	0.027	0.023	-0.004
i14	0.91	0.96	0.73	0.69	0.021	0.013	-0.008
i15	1.10	1.13	0.61	0.70	0.005	0.010	0.004
i16	1.04	1.08	0.67	0.58	0.063	0.045	-0.018
i17	1.04	1.08	0.77	0.86	0.027	0.033	0.006
i18	0.89	0.94	0.82	0.69	0.056	0.042	-0.014
i19	0.99	1.04	0.19	0.61	0.042	0.038	-0.005
i20	0.97	1.01	0.61	0.69	0.002	0.001	-0.001
i21	0.87	0.92	0.65	0.73	0.077	0.074	-0.004
i22	0.99	1.03	0.54	0.79	0.018	0.023	0.004
i23	0.72	0.78	0.71	0.78	0.016	0.014	-0.002
i24	0.94	0.99	0.76	0.66	0.049	0.036	-0.013
i25	0.45	0.52	0.58	0.72	0.064	0.058	-0.006
i26	0.79	0.85	0.51	0.77	0.007	0.011	0.003
Mean	0.85	0.90	0.62	0.72	0.060	0.056	-0.004
Median	0.93	0.97	0.66	0.71	0.020	0.023	-0.002
St Dev	0.24	0.23	0.21	0.13	0.116	0.109	0.010

Figure A.1: Ratio of average consumption to conditional optimal by condition, periods 1-10
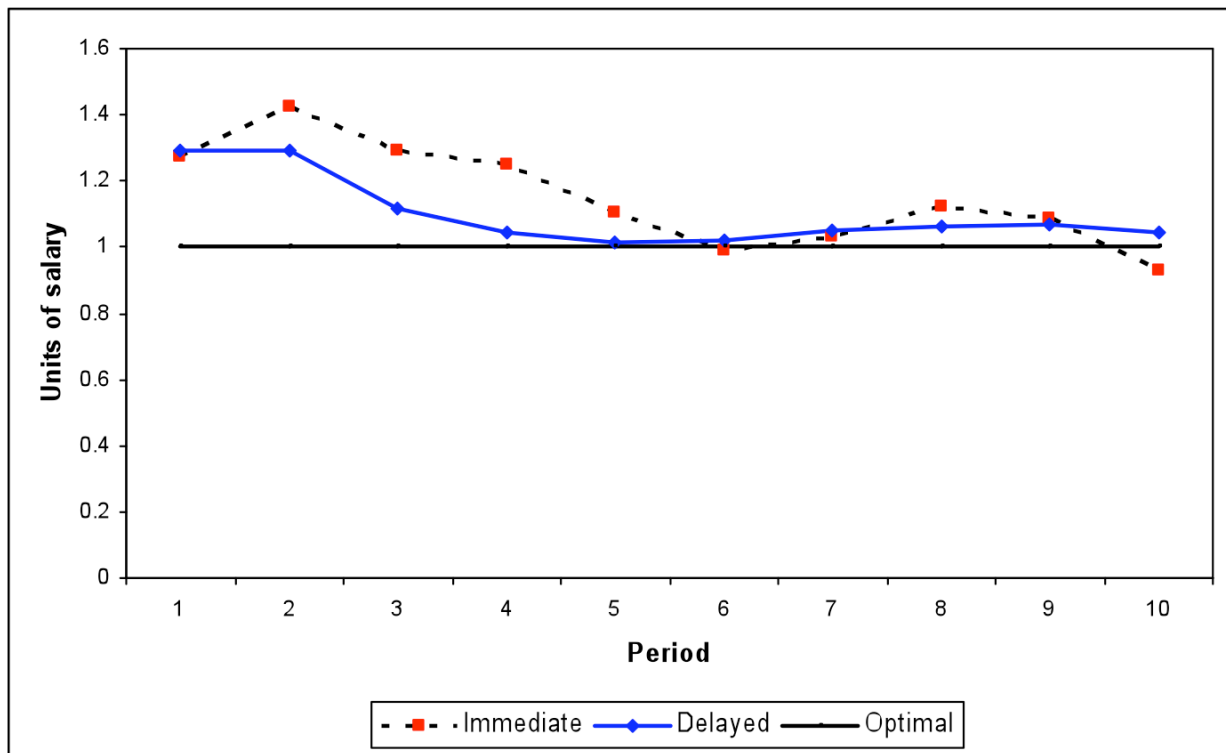
Figure A.2: Spending as a proportion of actual income each period, lifetimes 1 & 7, private learning condition (N=36)
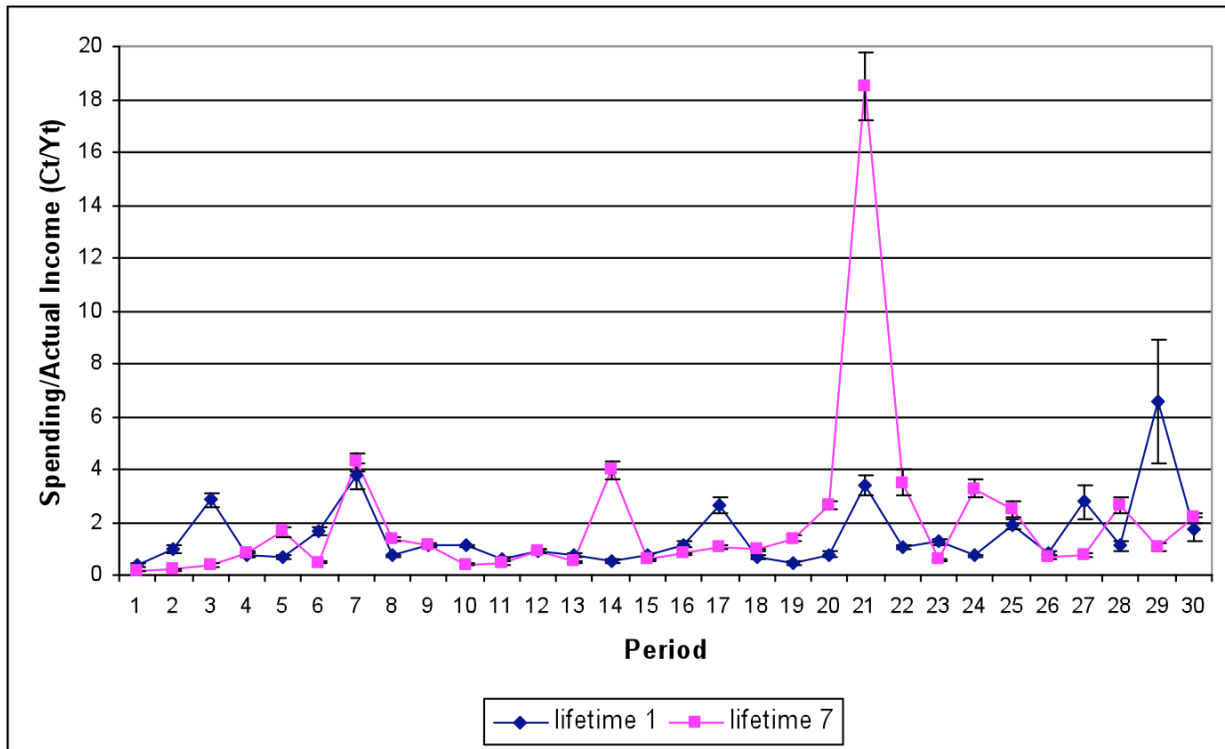

Figure A.3: Spending as a proportion of actual income each period, lifetimes 1 & 7, social learning condition (N=36)
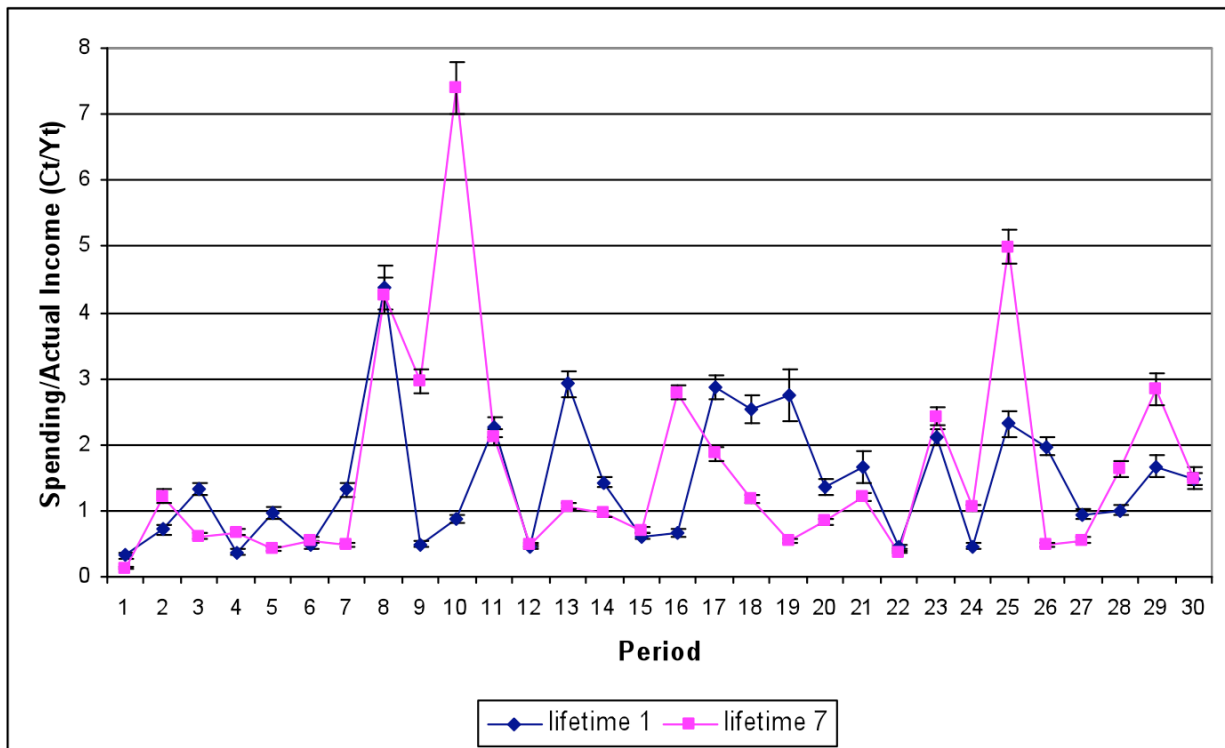
Figure A.4: Proportion of available cash spent each period, lifetimes 1 & 7, private learning condition (N=36)
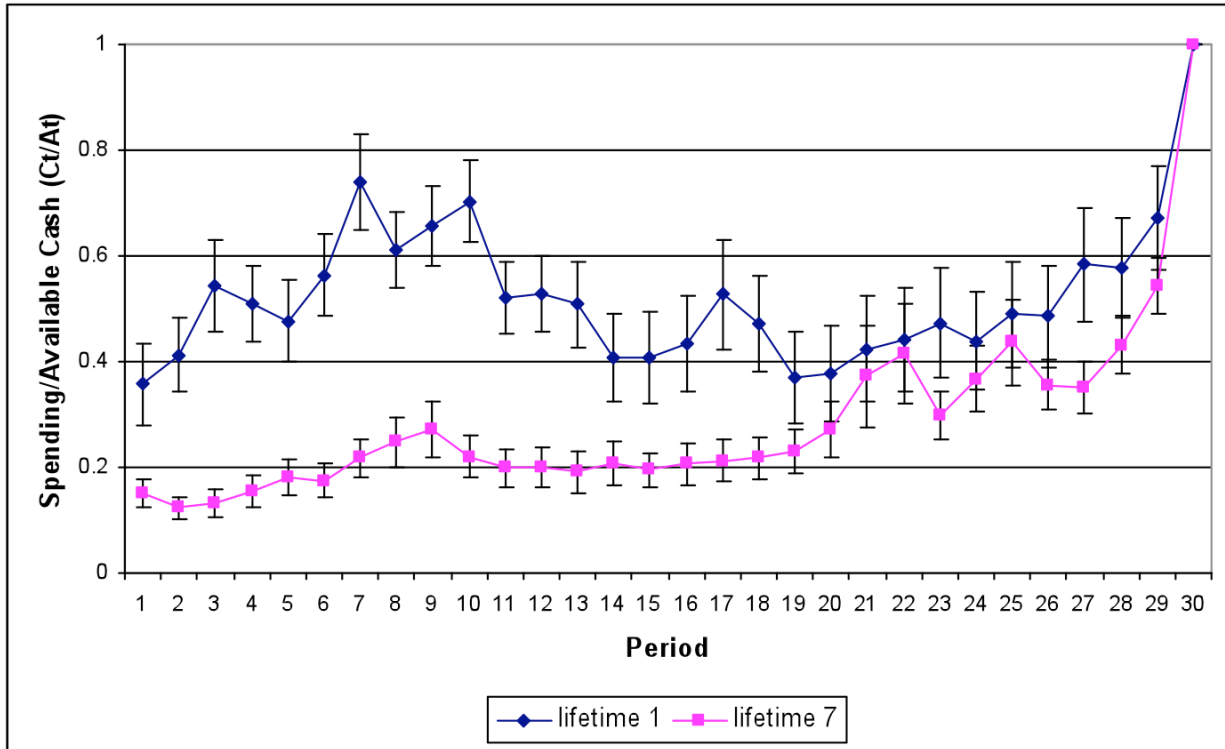

Figure A.5: Proportion of available cash spent each period, lifetimes 1 & 7, social learning condition (N=36)