LEARNING AND VISCERAL TEMPTATION IN DYNAMIC SAVINGS EXPERIMENTS*

Alexander L. Brown

Colin F. Camerer

Zhikang Eric Chua

Abstract

In models of optimal savings with income uncertainty and habit formation, people should save early to create a buffer stock, to cushion bad income draws and limit the negative internality from habit formation. In experiments in this setting, people save too little initially, but learn to save optimally within four repeated lifecycles, or 1-2 lifecycles with "social learning." Using beverage rewards (cola) to create visceral temptation, thirsty subjects who consume immediately overspend compared to subjects who only drink after time delay. The relative overspending of immediate-consumption subjects is consistent with hyperbolic discounting and dual-self models. Estimates of the present-bias choices are $\hat{\beta}$=0.6-0.7, which are consistent with other studies (albeit over different time horizons).

* 22 March 2007. This research was supported by NSF grant SES-0078911. We thank Chris Carroll, Daniel Houser, Paul Kattman, George Loewenstein, Tanga McDaniel, John Hey, Nat Wilcox, three anonymous referees and an editor for helpful comments. We also thank Julie Malmquist of the SSEL Caltech lab, Chong Juin Kuan (NUS), Hackjin Kim, Tony Bruguier, and especially Min Jeong Kang (who ran several of the beverage-condition subjects herself) for help in doing the experiments.

I. Introduction

Do people generally save the right amount? High rates of consumer debt and personal bankruptcies in the United States, and the drop in consumption upon retirement [Haverman e t al., 2006], are consistent with overspending and undersaving [Angeletos et al., 2001]. Choi et al. [2005] find that many employees do not save in their pension plan even when it is a dominant strategy to do so, and also suggest default options in retirement plans are more responsible for savings than preferences are [Choi et al., 2003]. But Scholz et al. [2006] find that savings in 80% of households in their sample fit their "optimal" model, and only 20% undersave marginally. Lusardi et al. [2004] explain the observed drop in US savings levels can be explained by the increases in government savings plans and stock market accumulation, but also cannot explain the low savings of a small segment of the population. Other economists argue that retirement advisors and traditional savings models have overestimated optimal savings, and when the proper benchmark is used most Americans actually *over*save compared to that benchmark [Darlin, 2007].

As the differences in these studies indicate, it is difficult to conclusively reject or accept the basic premise of lifecycle saving, which is that current saving correctly anticipates future needs and income variation, and smoothes consumption [Browning and Lusardi, 1996; Venti 2006]. The difficulty stems from the fact that econometric tests of the lifecycle model typically depend on many auxiliary assumptions about utility functions, separability across time, income expectations, retirement and other institutional rules, sorting, and credit market constraints. Apparent statistical evidence of undersaving in any particular study might be due to one or more econometric misspecifications or to mismeasurement of capital gains, educational returns or durable consumption flows [Gale et al., 1999].[1]

While there is no doubt that progress is being made in studies using field data, because of the sensitivity of conclusions from field data to possible misspecifications, experiments which can create an economic environment matching the assumptions underlying standard theory precisely could be of some use.[2] Experiments of this type are complements to analyses of field data. Control in the lab enables us to evaluate which theories are on the right track and are good candidates for analysis in field data; and suggestive observations in field data inspire experimental attempts to reproduce those observations in carefully-controlled simple settings.

Of course, complementarity of controlled lab data and field data depends on the working assumption that lifecycle theories apply in both settings. The subjects in our experiments are intelligent college students. They are analytically capable and already have, or will soon be, making important savings decisions-- establishing credit, buying durables, and making retirement-fund investment allocations. It is not unreasonable to assume their behavior in these simple experimental settings will correspond to their later behavior when they begin to save more seriously in a few years.

Our paper tests three basic theories about savings. The first theory is that consumers save optimally relative to the buffer stock savings model. The second theory is that consumers do not understand how to save optimally, but can learn to do so from their own experience or the experiences of others. The third theory is that consumers cannot save optimally because they have a present bias for consumption over savings. We will test these theories with two different experimental studies and a pool of highly skilled subjects.

In these experiments, subjects make saving and spending decisions in a 30-period lifecycle experiment with uncertain income and habit formation. The experimental design implements typical assumptions that buffer-stock models make about income and utility.

Figuring out how to save optimally is cognitively challenging (it requires solving a two-state-variable, 30-period dynamic program). Optimal saving requires subjects to save a lot in early periods to buffer against bad income shocks and to avoid creating an early consumption "internality" that reduces utility from future consumption (because of the habit formation).

If experimental saving is suboptimal in these relatively simple experiments, then critics of the theories as descriptions of actual savings are entitled to be skeptical about whether average consumers save optimally in much more complex natural environments. Those who think people do save optimally, and any undersaving observed in experiments lacks generalizability, should specify what factors explain the difference between experimental and field results. Many of those factors could then be included in future experiments. The experiments are also relevant to repeated consumption decisions which take place over a shorter time scale than savings, such as days or months for highly addictive substances (such as crack cocaine or methamphetamine).

The first study explores accuracy of a computationally-difficult buffer stock model, and the effect of learning, when rewards are monetary. In the first 30-period lifecycle, most subjects save too little, so their lifecycle utility is far from the optimum. However, subjects approximate optimization surprisingly well after learning by doing over repeated lifecycles, or after learning "socially" from good and bad decisions by other subjects in the first one or two lifecycles.

The surprising approximation of optimality observed after a few lifecycles in the first study, led us to investigate what conditions would create suboptimal performance. Since the design of the first study produces near optimal performance with money, it is useful to keep the design across studies. In a second study in which consumption is sips of beverage for thirsty subjects (the first time such rewards have been used in economics experiments), we investigated whether visceral rewards would produce suboptimal performance. The use of beverage is designed to generate possible self-control problems involving other visceral rewards and activities such as food, drugs, exercise, procrastination, and shopping. We designed two conditions for this second study, and the literature suggests only one should produce suboptimal performance. If behavior is near-optimal, similar to the money study, in both visceral conditions, then the tentative conclusion is that near-optimal behavior is robust to whether people are earning abstract money or a visceral reward; if the behavior is far from optimal in the two beverage conditions uniformly then the conclusion is that the reward medium makes a difference, but self-control problems are not present. Only if, we find a difference between the two beverage studies, and the immediate condition is farther from optimal can we say psychological self-control problems may confound savings decisions.[3]

In the second study's "immediate" condition, period t decisions lead to immediate drinking. In the "delayed" condition, because the amount of beverage subjects decide to consume in period t is not delivered until period t+10, any preference for immediacy is muted.

Subjects generally consume too much compared to a total-reward-maximizing optimum when rewards are immediate, and consume less (although slightly more than that optimum) when rewards are delayed. The relative overconsumption in the more tempting immediate condition is consistent with models of hyperbolic discounting [i.e., Ainslie, 1975; Laibson, 1997] and dual-self conflict [e.g., Bernheim and Rangel, 2004; Fudenberg and Levine, 2006; Loewenstein and O'Donoghue, 2004]. Structural parameter estimates of quasi-hyperbolic β-δ discounting models find values comparable to those in other lab and field studies (mean β of 0.6-0.7), albeit over different time horizons.

II. Three Theories of Savings Behavior

The experimental design implements the assumptions of the buffer stock savings model of Carroll, Overland and Weil [2000]. Agents earn income each period, subject to stochastic independent shocks from a distribution they know. In each period their available cash is the previous buffer stock, plus new income. They divide this available cash between consumption and savings. Utility in each period depends upon a ratio of current consumption to a habit index. The habit index is 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 is repeated several times with different income realizations each time. Through these two experiments we are able to test three theories about how people save. These theories are described next.

1. Traditional Economic Theory

Traditional economic theory assumes that people 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. As noted earlier, there is mixed empirical evidence as to whether consumers make their savings decisions correctly. In our experimental design, subjects should save a lot to build up a buffer stock, then spend roughly their average income once their buffer stock is large enough. The buffer stock protects against bad future income draws, and limits the negative "internality" of current spending on future utility, which occurs because of the controlled effect of habit formation. Figure I 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). 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).

2. Bounded Rationality

A second theory is that consumers save suboptimally because their rationality is bounded, and solving for optimal saving in the buffer stock model is computationally difficult. Until the 1990's, computers were not powerful enough to solve exactly for optimal saving in realistic environments. Allen and Carroll [2001] also show that learning by simple reinforcement is far too slow to produce convergence to optimal saving in reasonable time scales. It is possible that consumers simply cannot figure out or learn over time how to save optimally. Their mistakes might be systematic (e.g., undersaving when young) or might be unsystematic, highly-variable deviations around the optimal savings path.

In previous tests of savings in lifecycle models, experimental subjects typically undersave. Kotlikoff et al. [1987] found subjects "overdiscounted" future income in simple environments with no income uncertainty. Other studies with variable lifetimes found evidence of undersaving [Carbone and Hey, 2004; Carbone 2005]. Fehr and Zych [1998] complicate saving by introducing habit formation (as in our study) -- future utilities are lower if previous consumption is higher. Their subjects do not seem to anticipate this "internality" and save too little. Ballinger et al. [2003] divide subjects into three cohorts to study social learning. Each subject in a cohort completes a savings task sequentially; subjects in each cohort give advice to teach their later-cohort successors and observe the advice of their earlier-cohort predecessors. Subjects overspent throughout the experiments. But subjects in the last cohort did significantly better than the first two when there was high income variance (but there was no improvement with low income variance) so there is no social learning. While undersaving is the typical finding, Hey [1988] pointed out that suboptimal savings did not cost subjects that much in his study, and the comparative statics predicted by theory in response to parameter changes are also correct [Hey and Dardanoni, 1988]. Bernasconi and Kirchkamp [2000] is the only study to show oversaving, in a design with uncertainty and overlapping generations.

Our first study will also include social learning. Since there is no widely-accepted theory of how bounded rationally should be modeled formally in these settings,[4] the presence of rationality bounds is inferred indirectly: if subjects make mistakes in the first lifecycle, but learn over time or from the social examples, then we infer that their initial mistakes resulted from bounds on rationality but mistakes are reduced by learning. Development of a theory of rationality bounds and learning is a priority for future research.

Another kind of bound on rationality is that consumers are overoptimistic about future income or underestimate the force of habit formation. Note that these possibilities are ruled out by inducing beliefs about the income process and subjects' understanding of the degree of habit formation. So if we find that subjects save optimally, but believe that Americans do not, then the experiments suggest that misperceptions about income and habit formation could be the culprit in generating suboptimal saving in the field data.

3. Temptation/Dynamically Inconsistent Preferences

A third theory is that consumers know how to save optimally, but cannot resist short-term temptations to consume for some products. To test this explanation for undersavings we converted consumption which is abstract to actual sips of beverage (for thirsty subjects). Comparing immediate and delayed delivery of beverage consumption enables us to study temptation and dynamic inconsistency. In the immediate condition, consumption decisions in one period affect the amount of a beverage the thirsty subjects can drink right away. In the delayed condition, consumption decisions affect the amount of beverage delivered 10 periods later (about 10 minutes, which feels like a substantial time when you are thirsty). Under quasi-hyperbolic discounting, subjects will drink more beverage in the immediate condition than in the delayed condition because the immediacy-preference (or present bias) disappears when rewards are delayed (see our working paper for details). Of course, small amounts of beverage are not as dramatic as temptations like drug addiction, gambling and credit card spending, but they are feasible in the lab and give 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 earn (see our working paper Appendix 1 for details and instructional tables). To avoid demand effects, 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 is $Y_t = P_t\eta_t$, the product of P_t, permanent income that grows at five percent (P_t =(1.05)P_{t+1}, with initial P_1=100) and a multiplicative shock η_t which is lognormally distributed ($\log\eta \sim N(-\frac{1}{2},1)$). There is no interest rate and discount factor, and no borrowing or investment. Period-specific utility depends on consumption and on an accumulated level of habit, according to $u(C_t, H_{t-1}) = k + \frac{\theta}{1-\rho}\left(\frac{C_t + \hat{\varepsilon}}{H_{t-1}^{\gamma}}\right)^{1-\rho}$ with risk-aversion parameter ρ=3 and a habit strength exponent γ=0.6.[5] The habit stock grows according to $H_t = \lambda H_{t-1} + C_t$, where λ =0.7 is a depreciation rate [as in Fehr and Zych, 1998] and the initial habit H_0=10. Notice that larger early consumption builds up the habit level and depreciates current-period utility, creating an "internality" which implies that optimization requires restrained consumption in early periods. The subject's problem is to choose the stream of consumption $\widetilde{C}_S$ in each period t to maximize his expected utility, $E_t\left[\sum_{s=t}^{T} u(\widetilde{C}_S, \widetilde{H}_{S-1})\right]$. Because T=30 in the experiments, the problem can 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 is 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) so we could have "lifetimes" which are 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 .2) in order to deliberately produce more income variation.

The goal of experiments like these is *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, is 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 which is larger than that observed in the modern American economy because higher income variation creates a more analytically-challenging environment. With high income volatility, deviations from rationality are more clearly observed; our design also combines uncertain income and habit formation, because previous experiments have studied each separately and if learning occurs its power is 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 is 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 displays the income obtained, the corresponding cash available, and the habit stock for each consumption choice (see Figure II). The program also calculates and displays the possible points (i.e., utilities) that can be obtained from different levels of spending, and the corresponding savings available for the next period. Participants can experiment by inputting different consumption amounts and see how much utility they will earn, and how much cash they will 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 is 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[6] plus a $5 showup fee. Subjects earned between $7.50 to $65 with an average of $45.

Thirty-six (36) subjects 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 they were also given samples of what three actual subjects had done in the private learning condition.[7] 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 how these three samples were chosen.

There are many ways to implement social learning or imitation [e.g., Ballinger et al., 2003, use direct talking]. Our method mimics 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. Keep in mind that the high-earning role model they are presented with might be a subject who overspent early on (relative to the optimum) but got lucky by receiving high income draws. If subjects copy the "successful" subject too directly they could easily overspend relative to the optimum; so it is not clear whether social learning will actually help, hurt, or have no effect.

Participants were 35 undergraduates from the National University of Singapore (NUS) and 37 undergraduates from California Institute of Technology. These students are 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 had the same income draws (but the draws were different in the two learning conditions).[8] Most participants completed the instruction and seven lifecycles in about 90 minutes.

2. Basic results

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.[9] 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 is well below the unconditional optimum and highly variable across subjects. However, by lifecycle four the average subject earns point totals within 80% of the optimum and the variability across subjects shrinks.

Table I (bottom panel) also shows that social learning brings 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.[10]

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 depends on the participant's *actual* cash-on-hand and accumulated habit stock. A subject who makes some bad decisions in early periods, but then wises up and makes conditionally optimal decisions in later periods, will look bad in Table I but may be close to conditionally optimal overall when a few early mistakes are averaged with smart 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 horizontal axis, 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 are spending significantly more than optimal in early periods, until about period 20 (when they often spend too little). However, the lifecycle 7 conditional deviations are never significantly different from zero, which shows that learning is 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, because the initial performance is so close to optimal.

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=.43) and ethnicity (Chinese=1, mean=.50).[11]

Table II shows the results. The period effect is positive (but nonlinear because the period^2 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 some of the later lifecycles, reflecting overall learning from multiple lifecycles (5 and 7). 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 almost the same as the first study with one large change.[12] 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[13] 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 through spending all available cash immediately. As noted in the introduction, this change was designed to see if savings decisions about abstract money reward are different than viscerally-tempting rewards-- namely, liquid consumption of 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.[14] They also begin by eating some salty snacks. Since it takes them 45 minutes to read the instructions and do two 30-period lifecycles for money before the beverage lifecycle, they are definitely thirsty by the time they reach the beverage lifecycle. It is likely that they do not satiate during the lifecycle because no subject received more than 350 ml of soda (less than a 12 oz. can of Coca-cola) in that lifecycle, and subjects could only drink a maximum of 20 ml/period (0.7 oz), and beverage periods were separated by one minute.[15] Subjects were required to drink their entire beverage consumption in that one minute period.

A syringe pump with three syringes was used to deliver an exact amount of beverage into a cup.[16] 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 is “forgiven” at the end of the beverage lifecycle because we cannot force subjects to “pay back” the debt for future drinking (as we do in the money lifecycles).

There were also 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.[17] Quasi-hyperbolic or present-bias models of time discounting predict that subjects will drink more beverage in the early periods of the immediate condition because delayed rewards are heavily discounted (see our working paper). In the delayed condition, immediate choices do not lead to immediate consumption so the present bias term in discounting disappears. As a result, subjects should drink more overall in the delayed condition.

Subjects were n=52 Caltech students.[18] 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 many 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 predict that subjects in the immediate condition would receive less beverage than in the delayed condition, because they will consume relatively more compared to a total-reward-maximizing optimum in early periods. 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). Though there is substantial variation across subjects, this difference is significant at conventional levels by one-tailed tests (t-test p=0.047, Mann-Whitney rank sum test p=0.015).

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 that could be evidenced by differential performance in the four money lifecycles. To control for these skill differences, we estimate the regression

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

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

Notice that point totals can be negative for the beverage lifecycle, but the total ml of beverage consumed cannot be negative. (Subjects cannot be 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 can be offset by other lifecycles. Subjects are aware of this difference in incentive structures. If a large beverage deficit (>350 ml) occurs in an earlier period, they know that no amount of spending can erase the deficit. As a result, when subjects have large negative point totals they can become indifferent about future decisions (their marginal incentive disappears) and produce 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).[19] In order to reduce the effects of these outliers, two alternative regressions were run. In one specification, each lifetime 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 $p<0.05$. The fact that these results are stronger in significance than the parametric t-tests reported in Table III, suggests 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.

These analyses use the overall point totals in the lifecycle. As in study 1, it is also useful to examine <u>conditional</u> 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 could receive 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, is 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

Figure V shows the average ratios of spending to conditionally optimal spending. (In the first 10 periods the optimal line is now just a flat line at a ratio of 1.) Figure V confirms that even when conditioning on past decisions, the immediate-condition subjects are 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 pushes down their overspending.)[20] Another diagnostic statistic is the average overspending in those periods in which subjects overspend 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),[21] but when they did overspend, the immediate condition subjects spent much more than was optimal (Table III, row 3), which is the key source of greater expected losses.

2.4 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 consume less overall. Because the hyperbolic model is clearly parameterized, we can also estimate best-fitting 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.[22]

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 by searching for a pair of best-fitting values from choices; there are many pairs with low β and high δ values that will fit about equally well, and when this procedure was used it often yields values of β that are close to 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 gives no information about the present bias β, in theory, the delayed-condition data will be used to estimate δ. So we first search for best-fitting values of δ_D which explain delayed-condition subject choices as if they are maximizing discounted expected utility of consumption with a discount rate δ_D and β=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 .904 and standard deviation, across subjects, of .230. This mean value is significantly less than one at the 2% level by a cross-subject t-test (see Table V).

The next challenge is to estimate β_I values in the immediate condition, using reasonable values of δ.[23] The procedure we use first fixes β=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 Table VI; the mean δ_I is 0.85 and the standard deviation is 0.24. Since we are fixing β=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 $\beta_.$ 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 (δ_I*-0.904)/0.23=(δ - 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 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).

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. Table VI shows individual results;[24] 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.[25] The estimates of $\tilde{\beta}$ are in the ballpark of estimates of

Angeletos et al. [2001] ($\tilde{\beta}$=0.55), Della Vigna and Paserman [2005] ($\tilde{\beta}$=0.9), and Tanaka, Camerer and Nguyen [2006] ($\tilde{\beta}$=0.74-0.89) (from macroeconomic calibration, unemployment spells, and experiments in Vietnam, respectively). 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 respectively.

2.5 Myopic loss-aversion

A widely used concept in behavioral economics which might apply here is "myopic loss-aversion." Loss-aversion is the idea that people are disproportionately averse to making decisions that create nominal losses relative to a point of reference [see Kahneman and Tversky, 1979; e.g., Camerer, 2005, Camerer, forthcoming, Goldstein et al., 2006]. Myopic loss-aversion means that people focus on losses only in a small segment of time or a part of a portfolio, neglecting the benefits of decision rules which aggregate losses and gains across choice sets. In our setting, one hypothesis from myopic loss-aversion is that subjects will be unusually reluctant to choose a consumption level that generates a period-specific utility which is negative (assuming zero is a reference point).[26] Figure VI tests this hypothesis using study 1 data, by plotting nominal utility losses in each period from actual consumption on the y-axis, and corresponding losses that would have resulted from conditionally-optimal consumption (for utilities between -50 and +50, n=14,228). There is a sharp visible drop-off in points between the bottom and top halves of the Figure VI scatter plot. It appears that subjects hate to lose a small amount of nominal utility, even when they should take a small loss to build up savings. [as shown in Fehr and Zych, forthcoming][27] A piecewise-linear jackknife regression through the origin gives coefficients in the domain of positive and negative conditionally optimal utilities of 0.79 and 0.15, respectively. The ratio of these two slopes is 5.2. Another way to see the effect is to plot histograms of small actual and optimal utilities (between -10 and +10), across all subjects and periods (see Figure VII). In the actual period-by-period utilities there is a huge spike on the slight positive side. This spike is not at all evident in the corresponding distribution of optimal utilities, so the preference for a small positive utility and aversion to loss is not normative.

Figures VIII-IX show actual-optimal utility scatter plots like Figure VI, for the money periods and beverage periods of study 2. For money (Figure VIII), the jackknife regression gives positive and negative slopes of 0.88 and 0.10 (a ratio of 8.8). For beverage (Figure IX), the slopes are 0.92 and 0.62 (a ratio of 1.49). The difference between money and beverage is consistent with the idea that in the domain of beverage, subjects know that a large loss creates a debt that they must pay off before getting more sips, so they are more willing to accept small losses rather than run up large debts. Johnson et al. [2006] also show variations in loss-aversion across domains.

The myopia underlying Figures VI-IX is surprising. The subjects make 210 separate money decisions in study 1, and 120 decisions in study 2. They know that the utilities in each of those periods will be added up at the end to determine their total money earnings. (The software even updates the total points for each lifecycle every period and shows the total at the bottom of the screen.) There is no good normative reason to avoid a small loss in any single period (as Figure VII showed). These data are a reminder that a complete theory of theory of loss-aversion and its interaction with a myopic focus needs to account for how broadly decisions are bracketed or lumped together [Thaler, 1999; Read et al., 1999].

V. Conclusions

Evidence on savings suggests people are not always saving optimally. However, tests with field data depend sensitively on assumptions about expectations, separability of consumption, and other unobservables. Experiments control for these assumptions, and generally show that experimental subjects save too little.

This paper seeks to establish some boundary conditions under which intelligent subjects at the start of their economic lifecycles make decisions in a complex savings problem which vary from nearly-optimal to highly suboptimal. The experiment environment is difficult because there is income uncertainty and habit formation (i.e., consumption utility depends on previous consumption). We find that subjects save much too little at first, but learn to save close to optimal amounts after three or four lifecycles of experience. Furthermore, subjects who have received social learning, examples of successful, unsuccessful, and average experimental performance, produce savings decisions that are quite close to optimal even in their first lifecycle. Since consumers are limited to one lifetime of private learning, 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?

The fact that subjects can learn to save optimally for money rewards led us to explore whether they save optimally when rewards are more immediate and visceral -- when thirsty subjects' rewards are immediate sips of a cola beverage. The subjects who sip the beverage immediately also overspend, compared to group of subjects who make decisions in one period but do 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 earn less total rewards than those in the delayed condition, and get 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 $\widetilde{\beta}$ (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. There is also strong evidence of myopic loss aversion in our experiment. Subjects in all treatments strongly prefer to make savings choices which result in small positive utility gains, when optimal decisions would lead to small negative utilities.

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 which 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. 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 lifetimes (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. 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 contribute to optimal choice.

Alexander L. Brown

Division of HSS 228-77
Caltech Pasadena CA 91125
abrown@hss.caltech.edu

Colin F. Camerer

Division of HSS 228-77
Caltech Pasadena CA 91125
camerer@hss.caltech.edu

Zhikang Eric Chua

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

References

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

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

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, XV (2001), 47-68.

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

Ballinger, T. Parker, Eric Hudson, Leonie Karkoviata and Nathaniel T. Wilcox, "Saving Performance and Cognitive Abilities," Stephen F. Austin State University, 2005.

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, XCVI (1988), 675-700.

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

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, XLVI (2000), 315-343.

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

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

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

Camerer, Colin F., "Three Cheers - Psychological, Theoretical, Empirical - for Loss Aversion," Journal of Marketing Research, XLII (2005), 129-133.

________, "Behavioral Economics," in R. Blundell, W. Newey and T. Persson, eds., Advances in Economics and Econometrics, Theory and Applications: Ninth World Congress of the Econometric Society. (Cambridge: Cambridge University Press, forthcoming).

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

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

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

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

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

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

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

Chua, Zhikang and Colin F. Camerer, "Experiments on Intertemporal Consumption with Habit Formation and Social Learning.," Caltech, 2004.

Darlin, Damon, "A Contrarian View: Save Less and Still Retire with Enough," The New York Times, January 27, 2007, p. 1

DeGeorge, Francois, Jayendu Patel and Richard Zeckhauser, "Earnings Management to Exceed Thresholds," Journal of Business, LXXII (1999), 1-33.

Della Vigna, Stefano and M. Daniele Paserman, "Job Search and Impatience," Journal of Labor Economics, XXIII (2005), 527-588.

Fehr, Ernst and Peter. K. Zych, "Do Addicts Behave Rationally?," Scandinavian Journal of Economics, C (1998), 643-662.

________, and ________, "Intertemporal Choice under Habit Formation," in C. R. Plott and V. L. Smith, eds., Handbook of Results in Experimental Economics. (forthcoming).

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

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

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

Goldstein, Daniel G., Eric J. Johnson and William F. Sharpe, "Measuring Consumer Risk-Return Tradeoffs," London Business School, 2006.

Gul, Faruk and Wolfgang Pesendorfer, "Temptation and Self-Control," Econometrica, LXIX (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, XLIV (2006), 249-264.

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

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

Johnson, Eric J., Simon Gaechter and Andreas Herrmann, "Exploring the Nature of Loss Aversion," SSRN, 2006.

Kahneman, Daniel and Amos Tversky, "Prospect Theory - Analysis of Decision under Risk," Econometrica, XLVII (1979), 263-291.

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

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

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

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

O'Donoghue, Ted and Matthew Rabin, "Doing It Now or Later," American Economic Review, LXXXIX (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.

Read, Daniel, George Loewenstein and Matthew Rabin, "Choice Bracketing," Journal of Risk and Uncertainty, XIX (1999), 171-197.

Scholtz, John Karl, Ananth Seshadri and Surachai Khitatrakun, "Are Americans Saving “Optimally” for Retirement?," Journal of Political Economy, CXIIII (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, LXXXVII (2002), 342-370.

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

Thaler, Richard H., "Mental Accounting Matters," Journal of Behavioral Decision Making, XII (1999), 183-206.

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

[1] For example, using the economic surprise of German reunification, Fuchs-Schundeln and Schundeln [2005] find that evidence of buffer-stock savings is sensitive to self-selection of risk-averse workers into low-risk professions.
[2] Consumers could undersave in field settings relative to theories based on rational expectations, because they are overconfident about their income, or pessimistic about their longevity. Our experimental setting controls for those beliefs, which are difficult (though not impossible) to observe in field data, because income distributions and longevity are known to subjects.
[3] In the case of the first two conclusions, it is useful to combine data across studies. It is not useful to combine data to make the third.
[4] Ballinger et al. [2005] model bounded rationality as individuals only looking ahead a fixed number of periods. They interpret the results of Ballinger et al. [2003] to suggest most subjects only look ahead two periods.
[5] 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 et al. [2003]. Scaling factors are $\theta = 750$ and $k=40$.
[6] 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 ≈ Sing $1.70).
[7] 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.
[8] 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).
[9] 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.
[10] It should be noted that in both conditions lifetime 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.
[11] See Chua and Camerer [2004] for details. Ethnicity is of interest because Singaporean Chinese have one of the highest savings rates in the world; [cf. 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.
[12] 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 et al, 1994] or if the signals of doing so are noisy [Bodner and Prelec, 2003; Benabou and Tirole, 2004].
[13] 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 middle coauthor) thought fruit juices that were tried were too filling and might induce satiation which complicates the analysis.
[14] 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.
[15] 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.
[16] See our working paper for a diagram of the beverage delivery apparatus.
[17] 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 of the delayed condition but received no rewards. In the last ten periods of that condition subjects received their rewards from periods 21-30 but made no decisions.

[18] 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.
[19] 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" will 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.
[20] 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 v. 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 drank about 57 ml (2 oz) on average after five periods.
[21] Periods in which a subject encountered a deficit of 20 ml or greater were omitted in this analysis.
[22] Additionally, subjects with high enough beverage deficits knew they would not receive liquid again and 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.
[23] 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, $\beta\delta$, $\beta\delta^2$, … $\beta\delta^t$ 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.
[24] A possible, self-revealed, equivalent 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 et al., 1995]. However, these values show little correlation with the individual naïve betas, sophisticated betas, nor subject performance ($-0.1<\text{correlation}<0.1$).
[25] The correlation of β and δ estimates across subjects is around .35 for both specifications of β, so there is no serious identification problem.
[26] Subjects sometimes input a series of consumption levels, trying to find the value that would give a positive utility. Unfortunately, the software did not capture these attempts; data like these would be useful to understand the nature of loss-aversion and its persistence.
[27] The result is reminiscent of DeGeorge, Patel and Zeckhauser's [1999] finding that small negative earnings announcements, and small year-to-year drops, are relatively rare for corporations.

Table I: Summary statistics of actual point outcomes

Lifecycle	1	2	3	4	5	6	7
			Private Learning				
Mean Points $\sum_{t=1}^{T} u(C_t, H_{t-1})$ (average over subjects)	118	-53	224	450	-65	435	440
Std. dev. Points	635	694	498	297	475	255	146
Deviation from Optimum $\sum_{t=1}^{T} u(C_t, H_{t-1}) - \sum_{t=1}^{T} u(C_t^*, H_{t-1}^*)$ (average over subjects)	*-453*	-628	-349	-125	11	-153	-149
Total Income $\sum_{t=s}^{T} X_t$	5471	7083	5215	6235	4300	4571	4789
			Social Learning				
Mean Points	325	586	559	589	309	539	504
Std. dev. Points	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

Table II: Regression of log(absolute conditional deviation) (t-statistics in parentheses)

	Model (3)
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)**
R^2	**0.20**

**p<.05; ** p<.01*

Table III: Summary statistics comparing immediate and delayed conditions in the beverage lifetime

	Immediate	Delayed	Parametric test	Nonparametric
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.

Table IV: Regression of periods and condition on subject performance

	Points	Beverage	Sign-preserved Log Points
Immediate condition dummy *(I)*	-2814.36** (27292.03)	-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 (16.76)	-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

**p<.05 (one-tailed) **p<.01 (one-tailed)*

Table V: Two-stage parameter estimates of $\tilde{\delta}$ and $\tilde{\beta}$

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

*1-tailed cross-subject t-test of parameter < 1: *p<.05 **p<.02 ***p<0.01*

Table VI – Estimated betas of individual subjects in immediate condition

Subject	Best fit $\tilde{\delta}$	Projected $\tilde{\delta}$ Delayed Condition	Soph $\tilde{\beta}$	Naïve $\tilde{\beta}$	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 I: An Optimal Consumption Path

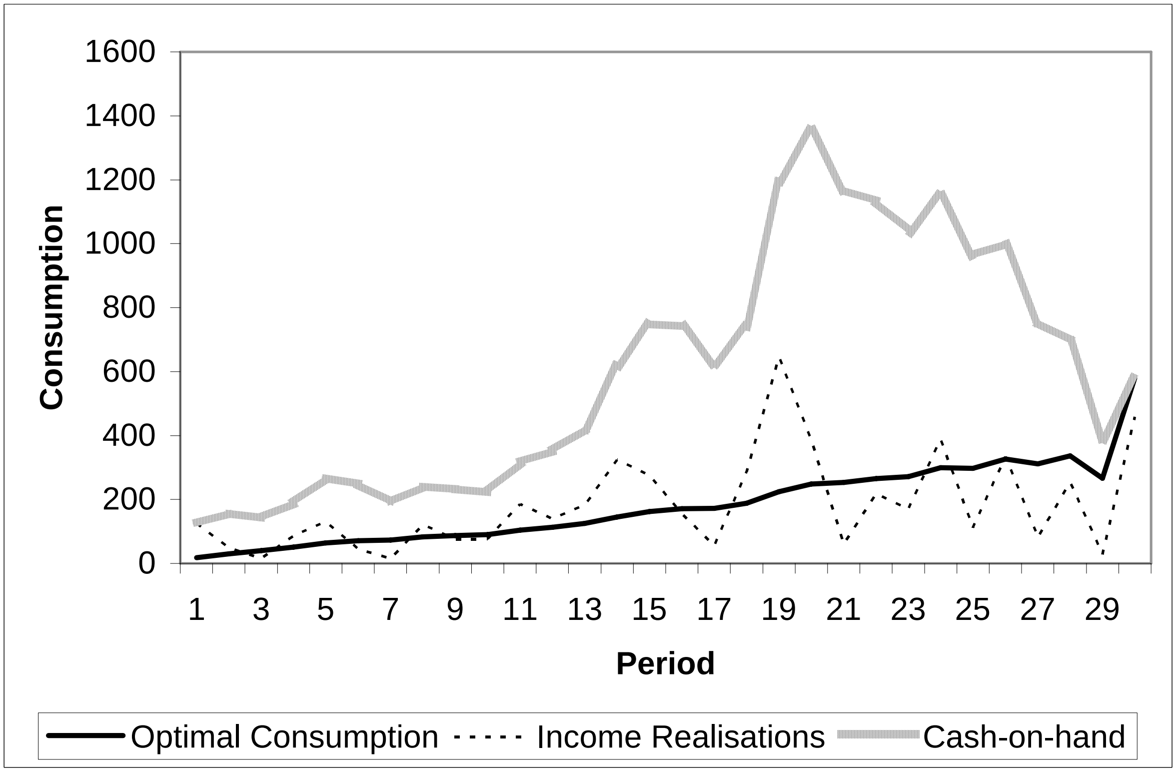

Figure II: Screenshot of Excel interface

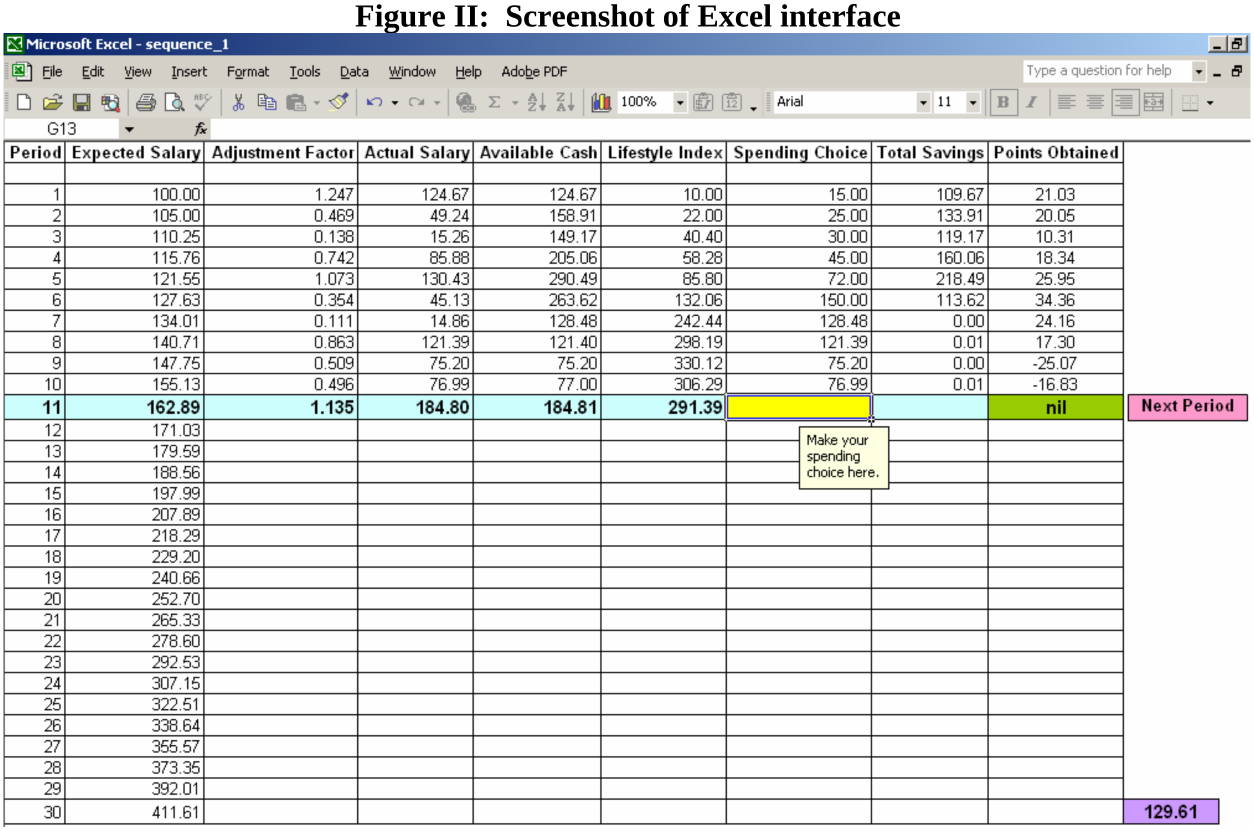

Period	Expected Salary	Adjustment Factor	Actual Salary	Available Cash	Lifestyle Index	Spending Choice	Total Savings	Points Obtained	
1	100.00	1.247	124.67	124.67	10.00	15.00	109.67	21.03	
2	105.00	0.469	49.24	158.91	22.00	25.00	133.91	20.05	
3	110.25	0.138	15.26	149.17	40.40	30.00	119.17	10.31	
4	115.76	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 III: Deviations from conditional optima, lifecycle 1 and 7, private learning

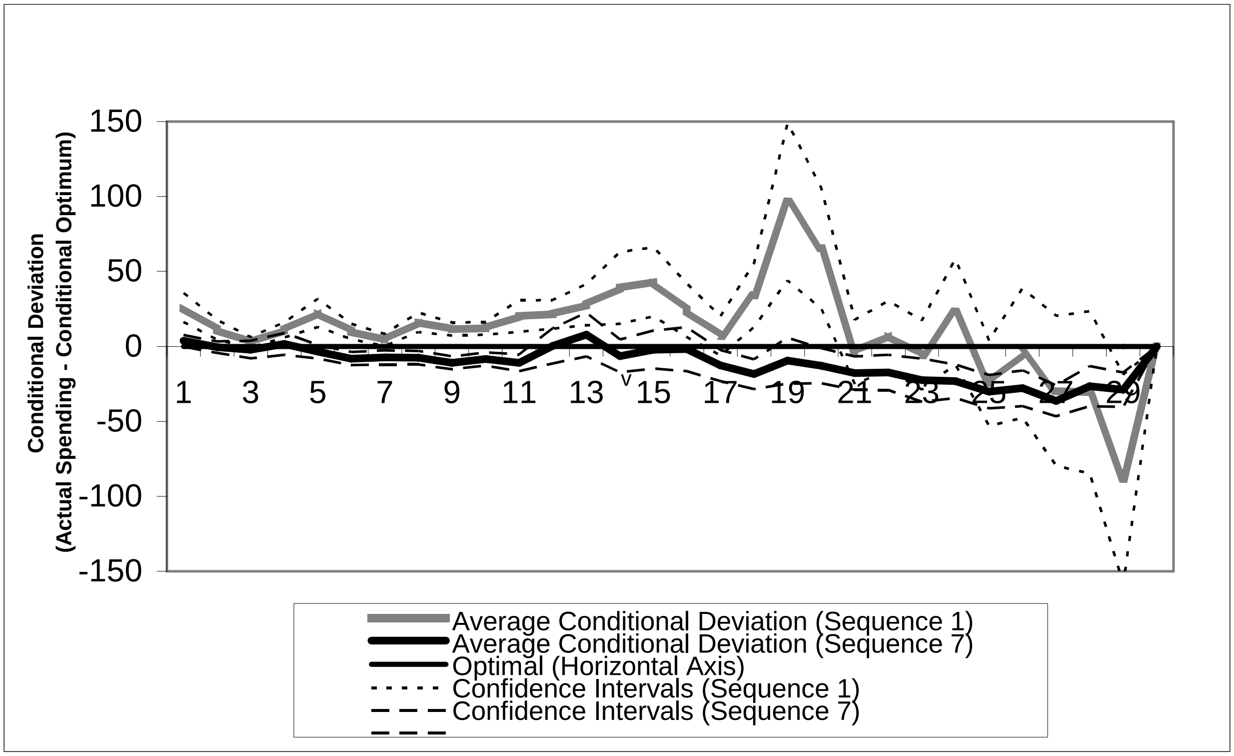

Figure IV: Deviations from conditional optima, lifecycle 1 and 7, social learning

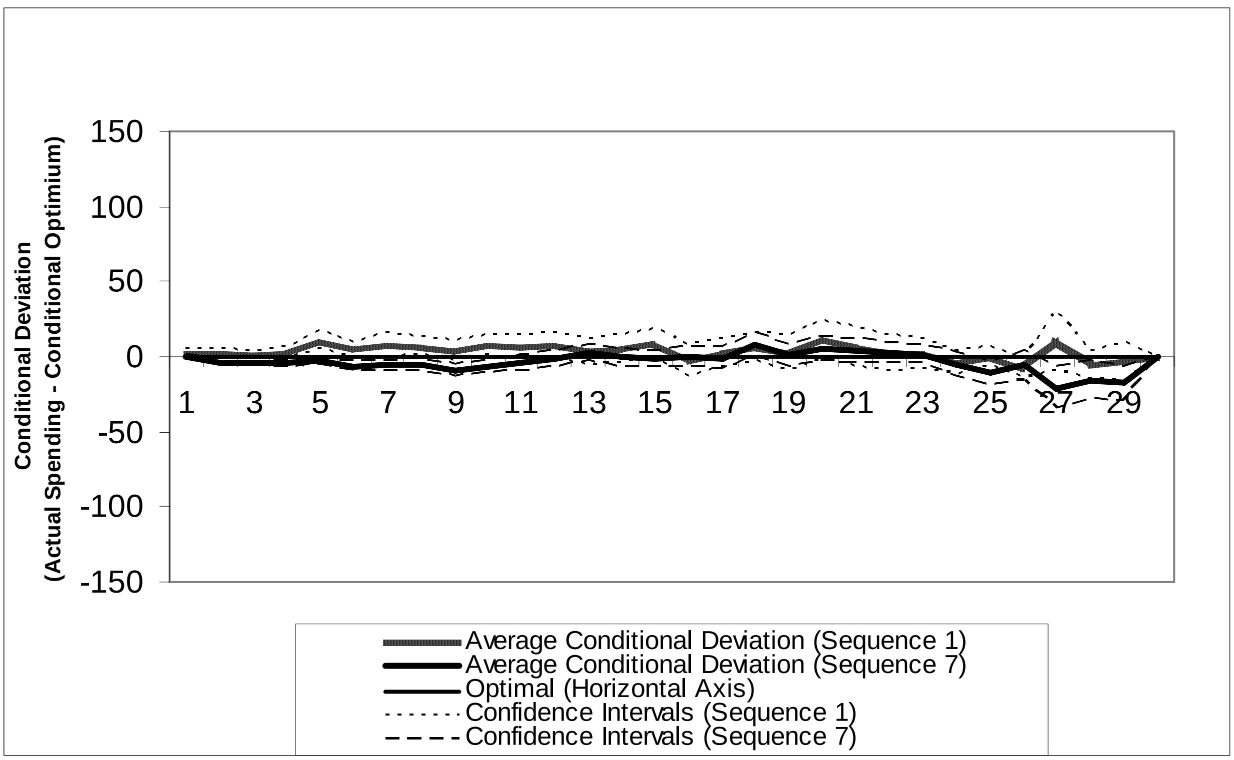

Figure V: Ratio of average consumption to conditional optimal by condition, periods 1-10

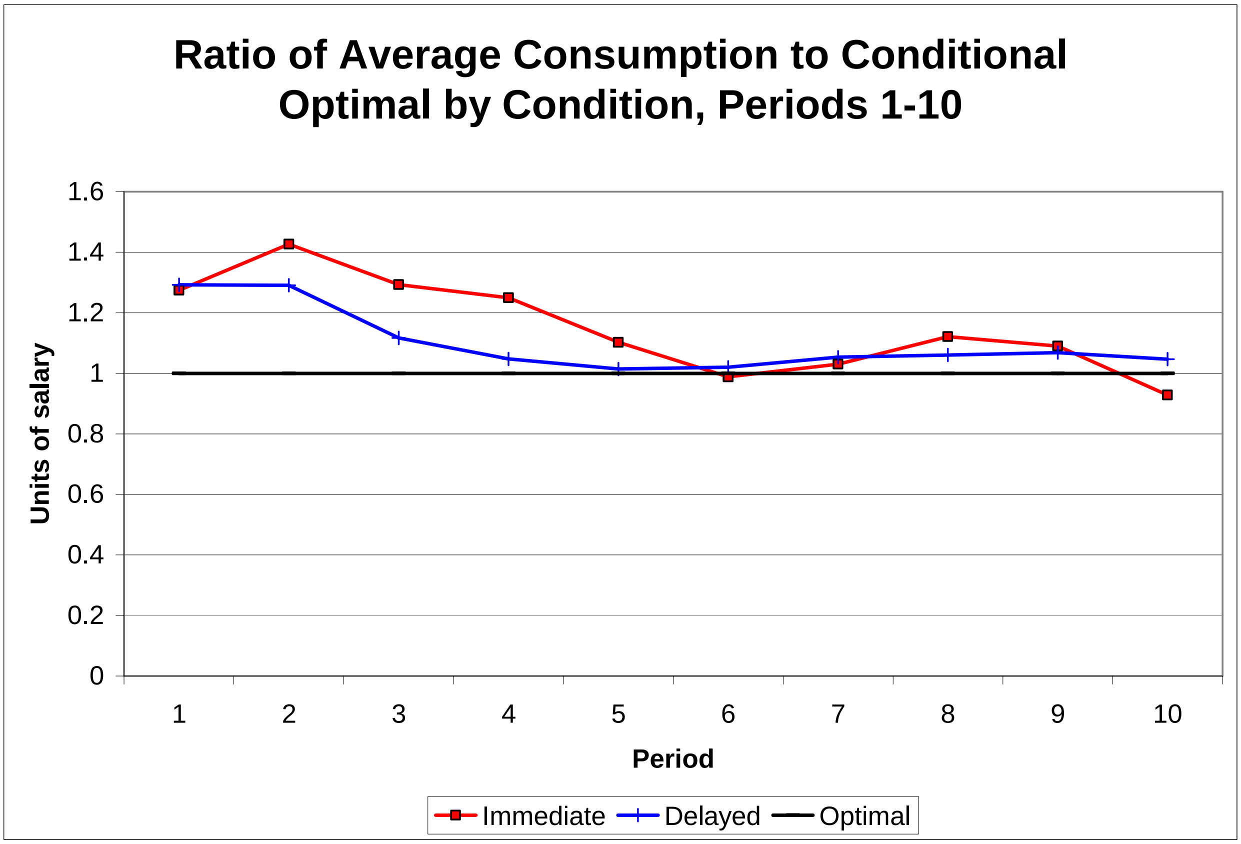

Figure VI: Actual (y) and conditionally optimal (x) utilities, observations between -50 and +50 (n=14,228)

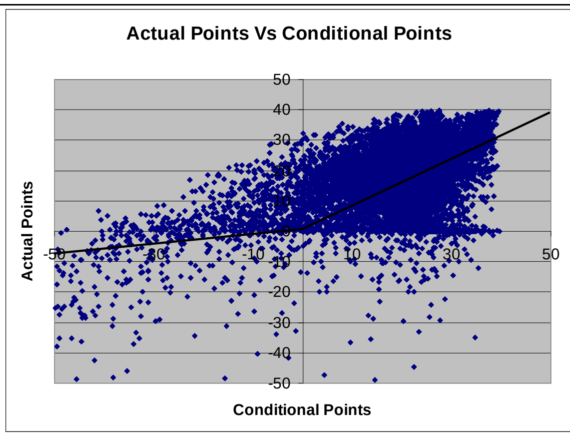

Figure VII: Frequency of actual and conditionally optimal utilities, observations between -10 and +10 (n=14,228)

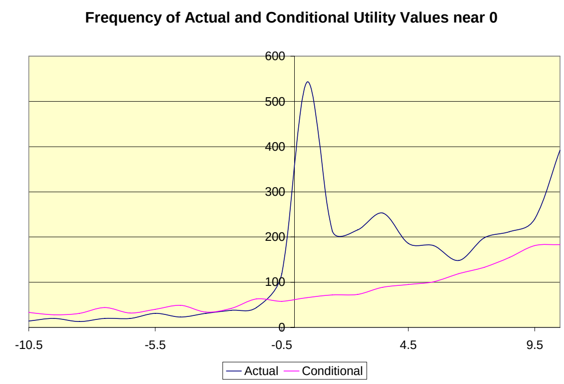

Figure VIII: Actual (y) and conditionally optimal (x) utilities in study 2, observations between -50 and +50 (n=5,840)

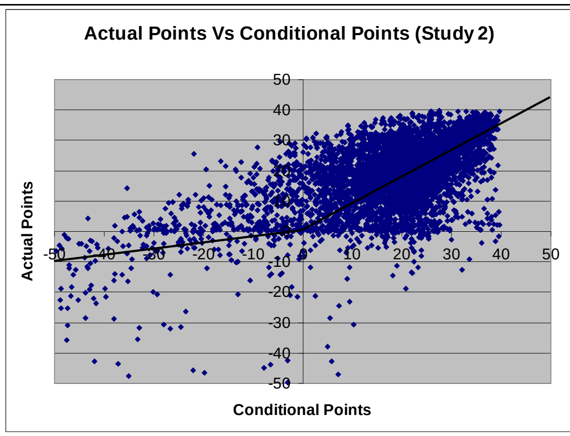

Figure IX: Actual (y) and conditionally optimal (x) ml of beverage, observations between -25 and +25 (n=1346)